

On the Inefficiency of Equilibria in Linear Bottleneck Congestion Games

Bart de Keijzer, Guido Schäfer, and Orestis A. Telelis*

Center for Mathematics and Computer Science (CWI)
Amsterdam, The Netherlands

`{B.de.Keijzer,G.Schaefer,O.Telelis}@cwi.nl`

Abstract. We study the inefficiency of equilibrium outcomes in *bottleneck congestion games*. These games model situations in which strategic players compete for a limited number of facilities. Each player allocates his weight to a (feasible) subset of the facilities with the goal to minimize the maximum (weight-dependent) latency that he experiences on any of these facilities. We derive upper and (asymptotically) matching lower bounds on the (strong) price of anarchy of linear bottleneck congestion games for a natural load balancing social cost objective (i.e., minimize the maximum latency of a facility). We restrict our studies to linear latency functions. Linear bottleneck congestion games still constitute a rich class of games and generalize, for example, load balancing games with identical or uniformly related machines with or without restricted assignments.

1 Introduction

Load balancing games constitute an important class of strategic games that capture many applications of practical relevance. These games model situations in which a set of strategically acting players (or jobs) compete for a limited number of resources (or machines). Every player chooses one of the resources available to him and assigns his weight (or load) to this resource. The latency of a resource depends on the total weight of the players using it. The goal of each player is to select a resource such that the latency that he experiences on this resource is minimized.

The study of load balancing games is motivated by the need for quantifying the inefficiency caused by selfish behavior of a set of autonomous players that utilize distributed processors upon which a system is built. The social cost objective of an assignment of loads to processors is measured by the *makespan*, i.e., the completion time of the most loaded machine, which reflects the distance from equi-distribution (balancing) of the load to the machines. Load balancing games have recently been studied extensively for a variety of different machine

* This work was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme.

environments, including identical [15], uniformly related [9, 11, 14, 15], restricted assignment [5, 11], and unrelated machines [2].

A natural extension of load balancing games are the *bottleneck congestion games (BCGs)* [6, 12]. Here, every player chooses a subset of the resources (also called *facilities* in this context) from a set of feasible facility allocations and assigns his weight to each of these facilities. The goal of each player is to select a subset of the facilities such that the maximum latency over the chosen facilities is minimized. Bottleneck congestion games generalize, for example, load balancing games and network routing games, and have several applications in practice. Despite their importance, bottleneck congestion games have received only very little attention in the literature and are far from being well-understood. In this paper we study the inefficiency of stable outcomes in bottleneck congestion games.

Bottleneck congestion games essentially generalize the context of load balancing games by modeling the activity of each selfish player upon complexes of interrelated resources. This generalization brings the model closer to practice, as in most large scale computing systems the workload of a player occupies different components of the system simultaneously. For example, instantiations of such games emerge if the components form paths in networks, or if they correspond to parallel processors, etc. It is natural to assume that each player wants to balance his load across the different components available to him and hence attempts to minimize the maximum latency of a facility that he uses.

One of the most prominent solution concepts for the prediction of outcomes of rational behavior in strategic games is the *Nash equilibrium* concept. It describes outcomes that are resilient to unilateral player deviations. Throughout this paper we will focus exclusively on pure Nash equilibria. A more general solution concept is the *strong equilibrium* concept introduced by Aumann [3]. It describes outcomes of strategic games that are stable with respect to pure deviations of player subsets (also called *coalitions*). More precisely, an outcome of a strategic game is a strong equilibrium if no coalition of the players can deviate such that every member of the coalition strictly benefits. An outcome is said to be a *k-strong equilibrium* if this property holds for all coalitions of size at most k . Strong equilibria thus generalize the pure Nash equilibrium concept ($k = 1$). Very recently, Harks, Klimm and Möhring [12] showed that (under rather general assumptions) bottleneck congestion games always admit strong equilibria.

It is well known that equilibrium outcomes might be inefficient in the sense that they are suboptimal with respect to some socially desirable objective function. The *price of anarchy (PoA)* [15–17] has become the standard measure to assess the inefficiency of equilibrium outcomes. It is defined as the worst-case ratio (over all instances) of the maximum cost of a Nash equilibrium outcome and the cost of a socially optimal outcome. The *strong price of anarchy (SPoA)* and the *k-strong price of anarchy (k-SPoA)* [2] refer to the natural adaptations of this measure to strong and k -strong equilibrium outcomes, respectively.

Contribution. We study the inefficiency of both pure Nash equilibria and strong equilibria of BCGs, under the natural assumption that the *social cost* of an outcome refers to the maximum latency of a facility. We restrict our studies to

	id. facilities		arb. facilities (SPoA)	
	k -SPoA (lower)	SPoA	id. players	arb. players
symmetric	$\max \left\{ 2, \lfloor \frac{m}{2k} \rfloor + 1 \right\}$	2	2	$O(m)$
asymmetric	$\max \left\{ \sqrt{2m + \frac{1}{4}} - \frac{1}{2}, \lceil \frac{m}{k-1} \rceil - 1 \right\}$	$\Theta(\sqrt{m})$	$O(\sqrt{n})$	$\Theta(m)$

Table 1: Summary of the bounds obtained for the SPoA and the k -SPoA of linear BCGs. The PoA of linear BCGs is at most $2m - 1$ and there is an asymptotically matching lower bound showing $\text{SPoA} \geq m - 1$.

linear bottleneck congestion games, where the latency of each facility is a linear function of the total weight assigned to it. These games still constitute a rich class of games and generalize, for example, load balancing games with identical or uniformly related machines with or without restricted assignments. We provide upper and lower bounds on the (strong) price of anarchy for symmetric and asymmetric linear BCGs (definitions will be given below). A summary of the results that we obtain in this paper is given in Table 1. Here, we use n and m to refer to the number of players and facilities, respectively.

1. We show that both the PoA and the SPoA of linear BCGs is $\Theta(m)$. More precisely, we show that $m \leq \text{PoA} \leq 2m - 1$ and $m - 1 \leq \text{SPoA} \leq m$.
2. We derive better bounds for identically weighted players. We prove that $\text{SPoA} = 2$ for symmetric linear BCGs and at most $O(\sqrt{n})$ and $O(\sqrt{m\gamma^*})$ for asymmetric linear BCGs, where γ^* refers to the cost of a socially optimal outcome.
3. We consider the case of identical facilities, i.e., all facilities have identical linear latency functions, and show that $\text{SPoA} = \Theta(\sqrt{m})$.
4. We also give elaborate lower bounds on the k -SPoA for symmetric and asymmetric BCGs with identical facilities (see Table 1).

We remark that we also provide asymptotically tight worst-case examples for (directed) network congestion games (definitions will be given below).

Related Work. Network BCGs were considered first by Banner and Orda in [6]. The authors showed existence of pure Nash equilibria and provided an $\Theta(m)$ bound on the PoA for identical network links. Busch and Magdon-Ismail studied in [7] the PoA of network BCGs for identically weighted players. Very recently, Harks, Klimm and Möhring introduced general bottleneck congestion games and showed that strong equilibria are guaranteed to exist in these games.

As mentioned above, bottleneck congestion games generalize load balancing games, which have been studied intensively in recent years. Load balancing games were first studied by Koutsoupias and Papadimitriou [15]. Among other results, the authors provided a lower bound on the PoA of *mixed* Nash equilibria for the case of identical machines. Koutsoupias, Mavronicolas and Spirakis [14] and, independently, Czumaj and Vöcking [9], proved a matching upper bound. Czumaj and Vöcking also proved that $\text{PoA} = \Theta(\log m / \log \log m)$ for pure Nash equilibria. The same bound on the PoA was shown by Awerbuch et al. [5] for

restricted assignments and identical machines. Gairing et al. [11] obtained independently the same bounds and proved $m - 1 \leq \text{PoA} \leq m$ for restricted assignments and uniformly related machines.

Andelman, Feldman and Mansour [2] were the first to study strong and k -strong equilibria in the context of load balancing games. They proved that $m \leq \text{SPoA} \leq 2m - 1$ for the case of unrelated machines, which was tightened to exactly m by Fiat et al. [10]. In this latter work it was also shown that the SPoA of strong equilibria for uniformly related machines is exactly $\Theta(\log m / (\log \log m)^2)$. For results in the context of more general scheduling games and associated scheduling policies (termed *coordination mechanisms*), the interested reader is referred to [13] and the references therein.

Bottleneck congestion games owe their name to their similarity to *congestion games*, which were introduced by Rosenthal [18]. In these games, the latency on each facility depends on the number of players using it (i.e., players have unit weights). The goal of each player is to minimize his cost which is defined as the sum (as opposed to the maximum for BCGs) of the latencies over the facilities used by the player. Rosenthal [18] proved the existence of pure Nash equilibria in congestion games. The price of anarchy of pure Nash equilibria for congestion games was resolved by Christodoulou and Koutsoupias [8] and, independently, by Awerbuch, Azar and Epstein [4]. It is shown in [8] that $\text{PoA} = \Theta(\sqrt{n})$ for asymmetric linear congestion games and the social cost being the maximum over the players' cost, and $\text{PoA} = \frac{5}{2}$ for (symmetric and asymmetric) linear congestion games and the social cost being the sum of the players' costs. Bounds for polynomial latency functions were also derived in [8]. Exact bounds for polynomial latencies and also for weighted players were developed in [1].

2 Preliminaries

In a *bottleneck congestion game*, we are given a set $N = [n]$ of n players that want to utilize (non-cooperatively) a set $E = [m]$ of m resources, which we also call *facilities*.¹ Every player $i \in N$ has a positive weight (or load) $w_i > 0$ and a strategy set $\Sigma_i \subseteq 2^E$ of feasible facility subsets from which he can choose. If player i chooses facility subset $S_i \in \Sigma_i$, he allocates his entire weight w_i to each facility $e \in S_i$. Let $\Sigma = (\Sigma_1, \dots, \Sigma_n)$ be the set of all possible strategy choices of the players. A *strategy profile* $S = (S_1, \dots, S_n) \in \Sigma$ specifies for each player $i \in N$ a strategy $S_i \in \Sigma_i$ that he has chosen. We define $N_e(S)$ as the set of players that have chosen facility $e \in E$ under S , i.e., $N_e(S) = \{i \in N \mid e \in S_i\}$. The total weight of facility $e \in E$ with respect to S is defined as $w_e(S) = \sum_{i \in N_e(S)} w_i$.

Every facility $e \in E$ has a latency function $l_e : \Sigma \rightarrow \mathcal{R}^+$ which satisfies the following three properties (see also [12]):

1. **Non-negativity:** $l_e(S) \geq 0$ for all $S \in \Sigma$.
2. **Independence of irrelevant alternatives:** $l_e(S) = l_e(S')$ for all $S, S' \in \Sigma$ with $N_e(S) = N_e(S')$.

¹ We use notation $[k]$ to refer to the set $\{1, \dots, k\}$ for some positive integer k .

3. Monotonicity: $l_e(S) \geq l_e(S')$ for all $S, S' \in \Sigma$ with $N_e(S) \supseteq N_e(S')$.

Given a strategy profile $S \in \Sigma$, every player $i \in N$ experiences an individual cost $c_i(S)$ equal to the latency of the most loaded facility that he uses, i.e., $c_i(S) = \max_{e \in S_i} l_e(S)$. We assume that every player $i \in N$ acts strategically and chooses his strategy $S_i \in \Sigma_i$ in order to minimize his own individual cost $c_i(S)$.

Aumann [3] introduced the notion of a *strong equilibrium*. Here we consider the refined notion of *k-strong equilibrium*. We use the standard notation S_{-i} to refer to $(S_1, \dots, S_{i-1}, S_{i+1}, \dots, S_n)$. Similarly, we use S_I and S_{-I} to refer to the strategy profiles of S induced by the players in I and $N \setminus I$, respectively.

Definition 1. A strategy profile $S \in \Sigma$ is a *k-strong equilibrium* if for every non-empty player set $I \subseteq N$ with $|I| \leq k$ and every possible joint deviation S'_I of I there is at least one player $i \in I$ whose cost with respect to $S' = (S_{-I}, S'_I)$ is not better than with respect to S , i.e., $c_i(S_{-I}, S'_I) \geq c_i(S)$.

With this definition, a *strong equilibrium* is a *k-strong equilibrium* with $k = n$, and a *pure Nash equilibrium* is a *k-strong equilibrium* with $k = 1$. Very recently, Harks, Klimm and Möhring [12] showed that strong equilibria always exist in BCGs satisfying Properties 1–3 above.

We are interested in characterizing the inefficiency of *k-strong equilibria* for BCGs. We assess the efficiency of a strategy profile S by the maximum load of a facility under S . That is, the *social cost* $C(S)$ of a strategy profile $S \in \Sigma$ is defined as the maximum latency over all facilities, which is equivalent to the maximum cost over all players, i.e., $C(S) = \max_{e \in E} l_e(S) = \max_{i \in N} c_i(S)$. We will use S^* to refer to an optimal strategy profile that minimizes $C(S)$ and denote its cost by $\gamma^* = C(S^*)$.

The *k-strong price of anarchy (k-SPoA)* [2, 15] refers to the worst-case ratio over all possible input instances of the maximum cost of a *k-strong equilibrium* and the cost γ^* of the social optimum. We will simply refer to the *price of anarchy (PoA)* and *strong price of anarchy (SPoA)* for the 1-SPoA and the *n-SPoA*, respectively. One can easily make an example to show that the SPoA is unbounded in general. This motivates our studies of *linear BCGs*: We assume that the latency function l_e of each facility $e \in E$ is a linear function of the total weight assigned to it, i.e., $l_e(S) = a_e w_e(S)$ for some $a_e \geq 0$. Linear BCGs constitute an important class of BCGs because they generalize, for example, various load balancing games as outlined in the Introduction.

A BCG is called a *network BCG* if there exists a directed graph $G = (V, E)$ such that every player $i \in N$ is associated with a source $s_i \in V$ and a sink $t_i \in V$ and i 's strategy set Σ_i refers to the set of all directed paths from s_i to t_i in G . We call a game *symmetric* if all players have the same strategy set, i.e., $\Sigma_i = \Sigma_j$ for all $i, j \in N$; we call a game *asymmetric* otherwise. Observe that the above example corresponds to a network BCG, but is not symmetric.

Unless stated otherwise, we assume subsequently that all player weights are at least one, i.e., $w_i \geq 1$ for every $i \in N$, and that the coefficient of each latency function is at least one, i.e., $l_e(S) = a_e w_e(S)$ with $a_e \geq 1$ for every $e \in E$. These

assumptions are without loss of generality as we can always enforce them by scaling the weights and coefficients appropriately.

3 Arbitrary Facilities

In this section, we derive bounds on the PoA and SPoA of linear BCGs. We consider both the general and the identical player case.

3.1 Arbitrary Players

We first consider the most general case of arbitrary linear latency functions and arbitrary player weights. We show that the PoA is at most $2m - 1$ in this case. We obtain a better bound of m on the SPoA and present an almost tight lower bound.

Theorem 1. *The price of anarchy of linear BCGs is at most $2m - 1$ and at least m .*

Proof. Let S be a pure Nash equilibrium with cost $C(S) = \alpha\gamma^*$ for some $\alpha \geq 1$. We prove by induction that for every integer k , $1 \leq k < \frac{\alpha+1}{2} + 1$, there is a set E_k of k distinct facilities such that for every $e \in E_k$, $l_e(S) \geq (\alpha - k + 1)\gamma^*$.

The claim holds true for $k = 1$ because there must exist a facility $e \in E$ with latency $l_e(S) = \alpha\gamma^*$. Suppose that the induction hypothesis holds true for $k < \frac{\alpha+1}{2}$. We will prove that there exists a set E_{k+1} of $k + 1$ distinct facilities such that $l_e(S) \geq (\alpha - k)\gamma^*$ for every $e \in E_{k+1}$. Choose from E_k a facility \hat{e} with smallest a_e , i.e., $\hat{e} = \arg \min_{e \in E_k} a_e$. By the induction hypothesis, we have $l_{\hat{e}}(S) \geq (\alpha - k + 1)\gamma^* > k\gamma^*$. Let $I_{\hat{e}} = N_{\hat{e}}(S)$ be the set of players choosing \hat{e} under S . Note that $w_{\hat{e}}(S) \geq l_{\hat{e}}(S)/a_{\hat{e}} > k\gamma^*/a_{\hat{e}}$. Consider the strategies that the players in $I_{\hat{e}}$ choose under S^* and suppose for the sake of a contradiction that for every $i \in I_{\hat{e}}$, $S_i^* \cap E_k \neq \emptyset$. Then there is a facility $e \in E_k$ with $w_e(S^*) \geq w_{\hat{e}}(S)/k > \gamma^*/a_{\hat{e}}$. By the choice of \hat{e} , we have $l_e(S^*) = a_e w_e(S^*) > \gamma^*$, which is a contradiction to the definition of γ^* . Thus there is a player $j \in I_{\hat{e}}$ that chooses a strategy S_j^* that is disjoint from E_k . Note that for every $e \in S_j^*$ we have $a_e w_j \leq \gamma^*$. Since S is a pure Nash equilibrium, player j cannot decrease his cost by deviating to S_j^* and thus there is some facility $e' \in S_j^*$ such that:

$$l_{e'}(S) = (a_{e'} w_{e'}(S) + a_{e'} w_j) - a_{e'} w_j \geq c_i(S) - a_{e'} w_j \geq l_{\hat{e}}(S) - \gamma^* \geq (\alpha - k)\gamma^*$$

The inductive step follows by setting $E_{k+1} = E_k \cup \{e'\}$. By choosing $k = \lceil \frac{\alpha+1}{2} \rceil < \frac{\alpha+1}{2} + 1$, we obtain that there is a set $E_k \subseteq E$ with $|E_k| \geq k$ and thus $m \geq |E_k| \geq k \geq \frac{\alpha+1}{2}$. We conclude that $\text{PoA} = \alpha \leq 2m - 1$.

The following instance shows that $\text{PoA} \geq m$, even for symmetric BCGs with identical facilities and identical players. Consider a BCG with player set $N = [n]$ and facility set $E = [m]$ with $m = n$. Every player $i \in N$ has unit weight $w_i = 1$ and the latency function $l_e(S)$ of every $e \in E$ is the identity function,

i.e., $l_e(S) = w_e(S)$. Suppose that each player $i \in N$ has strategy set $\Sigma_i = 2^E$. If every player chooses a distinct facility we obtain an optimal strategy profile S^* with $\gamma^* = 1$. On the other hand, consider the strategy profile S in which every player allocates all facilities in E . This is a pure Nash equilibrium of cost $C(S) = m$. \square

We derive a better upper bound on the SPoA for linear BCGs. The following key lemma will be used several times in the paper.

Lemma 1. *Let S be a strong equilibrium and let $I_\lambda \subseteq I$ be a non-empty subset of the players such that for every $i \in I_\lambda$ we have $c_i(S) \geq \lambda\gamma^*$, for some $\lambda \geq 1$.*

1. *Then there is a player $i \in I_\lambda$ and a facility $e \in S_i^*$ such that $l_e(S_{-I_\lambda}) \geq (\lambda - 1)\gamma^*$.*
2. *Suppose that I_λ is maximal. Then there is a player set $T_\lambda \subseteq N \setminus I_\lambda$ with $w(T_\lambda) \geq \lambda - 1$ and for every $i \in T_\lambda$ we have $(\lambda - 1)\gamma^* \leq c_i(S) < \lambda\gamma^*$.*

Proof. We first prove the first part of the lemma. Note that for every player $i \in I_\lambda$ and every $e \in S_i^*$ we have

$$l_e(S_{I_\lambda}^*) \leq l_e(S^*) \leq \gamma^*. \quad (1)$$

Suppose for the sake of a contradiction that for every player $i \in I_\lambda$ and for every $e \in S_i^*$ it holds that $l_e(S_{-I_\lambda}) < (\lambda - 1)\gamma^*$. Consider the strategy profile $S' = (S_{-I_\lambda}, S_{I_\lambda}^*)$ in which the players in I_λ deviate to their optimal strategies in S^* . Using (1), we obtain for every $i \in I_\lambda$ and for every $e \in S_i^*$:

$$l_e(S') = l_e(S_{I_\lambda}^*) + l_e(S_{-I_\lambda}) < \gamma^* + (\lambda - 1)\gamma^* = \lambda\gamma^*. \quad (2)$$

Thus, for every $i \in I_\lambda$, $c_i(S') = \max_{e \in S_i^*} l_e(S') < \lambda\gamma^*$, which is a contradiction to S being a strong equilibrium.

We next prove the second part of the lemma. Let $i \in I_\lambda$ be a player and $e \in S_i^*$ be a facility satisfying $l_e(S_{-I_\lambda}) \geq (\lambda - 1)\gamma^*$. Define T_λ as the set of players that choose e under S but are not contained in I_λ , i.e., $T_\lambda = N_e(S) \setminus I_\lambda \subseteq N \setminus I_\lambda$. We have

$$a_e w(T_\lambda) = l_e(S_{T_\lambda}) = l_e(S_{-I_\lambda}) \geq (\lambda - 1)\gamma^*. \quad (3)$$

Since $e \in S_i^*$ and $w_i \geq 1$ for every $i \in N$, we have $a_e \leq \gamma^*$. Thus, $w(T_\lambda) \geq \lambda - 1$. Consider an arbitrary player $i \in T_\lambda$. By the above we have, $c_i(S) \geq l_e(S) \geq l_e(S_{T_\lambda}) \geq (\lambda - 1)\gamma^*$. Moreover, by the maximality of I_λ and since $i \notin I_\lambda$, we have $c_i(S) < \lambda\gamma^*$. \square

Remark 1. Observe that in the above proof we exploit the linearity of the latency functions only in (2). In fact, we can draw exactly the same conclusion if all latency functions are *sub-additive*, i.e., for every $e \in E$, $l_e(x + y) \leq l_e(x) + l_e(y)$ for every $x, y \in \mathcal{R}^+$. As a consequence, all our upper bounds on the SPoA (which exploit Lemma 1) hold for sub-additive latency functions.

Theorem 2. *The strong price of anarchy of linear BCGs is at most m .*

Proof. Let S be a strong equilibrium with cost $C(S) = \alpha\gamma^*$ for some $\alpha > 1$. For an arbitrary real value $1 < \lambda \leq \alpha$, let I_λ be the maximal non-empty set of players $I_\lambda = \{i \in N \mid c_i(S) \geq \lambda\gamma^*\}$. Applying Lemma 1, we obtain a player set T_λ such that for every $i \in T_\lambda$ we have $(\lambda - 1)\gamma^* \leq c_i(S) < \lambda\gamma^*$. Moreover, $w(T_\lambda) \geq \lambda - 1 > 0$ because $\lambda > 1$ and thus T_λ is non-empty. We can thus identify a family $F = \{T_\alpha, T_{\alpha-1}, \dots, T_{\alpha-k}\}$ of $k + 1$ player sets that are non-empty and pairwise disjoint, where k is the largest integer satisfying $\alpha - k > 1$. Every set $T_\lambda \in F$ identifies at least one distinct facility $e \in E$ with $(\lambda - 1)\gamma^* \leq l_e(S) < \lambda\gamma^*$. Moreover, there is one facility $e \in E$ with $l_e(S) = \alpha\gamma^*$. We conclude that $m \geq |F| + 1 = k + 2 \geq \alpha$ and thus $\text{SPoA} = \alpha \leq m$. \square

Theorem 3. *The strong price of anarchy is at least $m - 1$ in general linear BCGs and at least $\frac{m+1}{3}$ in single-sink linear network BCGs.*

The proof of this result is deferred to the full version. The lower bound of $m - 1$ can also be derived by a construction in [11].

3.2 Identical Players

We next derive an upper bound on the SPoA for linear BCGs if the weights of all players are identical. In this subsection, we assume without loss of generality that the weight of each player $i \in N$ is $w_i = 1$.

Theorem 4. *The strong price of anarchy is at most $O(\min\{\sqrt{n}, \sqrt{m\gamma^*}\})$ for linear BCGs with identical players and 2 for linear symmetric BCGs with identical players.*

Proof. We prove the first part of the theorem. Let S be a strong equilibrium with cost $C(S) = \alpha\gamma^*$ for some $\alpha > 1$. As in the proof of Theorem 2, we can apply Lemma 1 to identify a family $F = \{T_\alpha, T_{\alpha-1}, \dots, T_{\alpha-k}\}$ of $k + 1$ player sets that are non-empty and pairwise disjoint, where k is the largest integer satisfying $\alpha - k > 1$. Each such set $T_\lambda \in F$ contains at least $\lambda - 1$ players, i.e., $|T_\lambda| \geq \lceil \lambda - 1 \rceil$ for every $\alpha - k \leq \lambda \leq \alpha$. Moreover, there is at least one player that experiences a congestion of $\alpha\gamma^*$. Thus

$$n \geq 1 + \sum_{\lambda=1}^{\lceil \alpha-1 \rceil} \lambda \geq 1 + \frac{\alpha(\alpha-1)}{2}.$$

Solving for α we obtain $\alpha \leq \frac{1}{2} + \sqrt{2n - 3/2}$. Recall that we assume without loss of generality that $a_e \geq 1$ for every $e \in E$ and thus $\gamma^* \geq n/m$. We therefore also obtain $\alpha \leq \frac{1}{2} + \sqrt{m\gamma^* - 3/2}$. Thus $\text{SPoA} \leq \alpha = O(\min\{\sqrt{n}, \sqrt{m\gamma^*}\})$.

We next prove the second part of the theorem. In a strong equilibrium S , at least one player $i \in N$ must have cost $c_i(s) \leq \gamma^*$ since otherwise the grand coalition could deviate to the socially optimal strategy profile. Suppose there is a player $j \in N$ whose cost is more than two times larger than the cost of i .

Consider the deviation $S' = (S_{-j}, S_j)$ where player j deviates to the strategy of player i . Then $c_j(S') \leq \max_{e \in S_i} a_e(w_e(S) + 1) \leq \max_{e \in S_i} 2a_e w_e(S) \leq 2c_i(S)$, which is a contradiction to S being a strong equilibrium.

The following example establishes the tightness of this bound: Let $N = [3]$ and $E = [6]$. The strategy set of every player is $\{\sigma_1 = \{1\}, \sigma_2 = \{2, 3\}, \sigma_3 = \{4, 5\}, \sigma_4 = \{2, 5, 6\}\}$. The social optimum is $S_i^* = \sigma_i$ for every player $i \in [3]$ with $\gamma^* = 1$. A strong equilibrium is given by $S_1 = \sigma_4$ and $S_2 = S_3 = \sigma_1$. The cost of S is $C(S) = 2$. It is easy to see that this example is a network BCG. \square

4 Identical Facilities

In this section, we study the SPoA for the case of linear BCGs with identical facilities, i.e., the latency function of every facility $e \in E$ is $l_e(S) = w_e(S)$.

Theorem 5. *The strong price of anarchy of linear BCGs with identical facilities is at most $-\frac{1}{2} + \sqrt{2m + \frac{1}{4}}$ in general and exactly 2 in case of symmetric games.*

Proof. For the symmetric case we claim that in any strong equilibrium configuration S , there is at least one player i_0 with $c_{i_0}(S) \leq \gamma^*$. Indeed, if $c_i(S) > \gamma^*$ for all players, then the grand coalition would deviate to S^* . Now for any player i we have $\gamma^* \geq w_i$. Let i be any player with $e \in S_i$ such that $c_i(S) = l_e(S) = C(S)$. Consider unilateral deviation $S'_i = S_{i_0}$ of i . Then, because S is also a pure Nash equilibrium, $C(S) = c_i(S) \leq c_{i_0}(S) + w_i \leq 2\gamma^*$. A tight lower bound has already been presented in Theorem 4.

For the asymmetric case let the cost of a strong equilibrium S be $C(S) = \alpha\gamma^*$, for some $\alpha > 1$. Similar to the proof of Theorem 2, let I_λ be the maximal non-empty set of players $I_\lambda = \{i \in N \mid c_i(S) \geq \lambda\gamma^*\}$ for some $1 < \lambda \leq \alpha$. By Lemma 1, we obtain a player set T_λ such that for every $i \in T_\lambda$ we have $(\lambda - 1)\gamma^* \leq c_i(S) < \lambda\gamma^*$. We can refine the argument given in the proof of Lemma 1 to bound the weight of T_λ for identical facilities as follows: By inequality (3), we have $w(T_\lambda) \geq (\lambda - 1)\gamma^*/a_e = (\lambda - 1)\gamma^*$, where the last equality holds because for identical facilities $a_e = 1$ for every $e \in E$. Moreover, $w(T_\lambda) \geq (\lambda - 1)\gamma^* > 0$ because $\lambda > 1$ and thus T_λ is non-empty. That is, we can identify a family $F = \{T_\alpha, T_{\alpha-1}, \dots, T_{\alpha-k}\}$ of $k + 1$ player sets that are non-empty and pairwise disjoint, where k is the largest integer satisfying $\alpha - k > 1$. Moreover, by construction we have $I_\alpha \cap T_\lambda = \emptyset$ for every $T_\lambda \in F$ and $w(I_\alpha) \geq \alpha\gamma^*$ since facilities are identical. The total weight $w(N)$ is then:

$$w(N) \geq \alpha\gamma^* + \sum_{\lambda=\alpha-k}^{\alpha} w(T_\lambda) \geq \alpha\gamma^* + \sum_{\lambda=\alpha-k}^{\alpha} (\lambda - 1)\gamma^* \geq \alpha\gamma^* + \sum_{\lambda=0}^{\alpha-1} \lambda\gamma^*$$

The latter equals $\frac{1}{2}\alpha\gamma^*(1 + \alpha)$. Observe that $\gamma^* \geq w(N)/m$ because facilities are identical. We obtain $2m \geq \alpha(1 + \alpha)$ or equivalently $\alpha \leq -\frac{1}{2} + \sqrt{2m + 1/4}$. Since $\text{SPoA} \leq \alpha$ the claim follows. \square

Theorem 6. *The strong price of anarchy of linear BCGs with identical players and identical facilities is at least $-\frac{1}{2} + \sqrt{2m + \frac{1}{4}}$ in general and at least $-\frac{1}{4} + \frac{1}{2}\sqrt{2 + 2m}$ in single-sink network BCGs .*

Proof. We give a family of instances with m facilities and $n = \Theta(m)$ unweighted players, which we turn into a family of network instances subsequently. Consider a partition of the set of players N into q subsets, $N = \bigcup_{j=1}^q P_j$, where $|P_j| = j$, $j \in [q]$. Denote players in P_j by p_{ji} , $i \in [j]$. For each subset P_j make a new set of j distinct facilities $E_j = \{e_1^j, \dots, e_j^j\}$. Define $E_{q+1} = E_1$. For every player $p_{ji} \in P_j$, $i \in [j]$, set the strategy space of p_{ji} to:

$$\Sigma_{p_{ji}} = \left\{ \{e\} \mid e \in E_j \right\} \cup \{E_{j+1}\}$$

For the socially optimal configuration set $S_{p_{ji}}^* = \{e_i^j\}$. Then $C(s^*) = 1$. Now consider the configuration S where $S_{p_{ji}} = E_{j+1}$ for $i \in [j]$, $j \in [q]$. The cost of S is defined by the latency of the unique facility $e = e_1^1 \in E_1$ and is $C(S) = l_e(S) = |P_q| = q$. For every player $p \in P_j$, we have $c_p(S) = j$. We claim that S is a strong equilibrium. Consider any deviation of any coalition $I \subseteq N$. Denote by S'_p the novel strategy that any player $p \in I$ adopts and let S' denote the resulting configuration. Notice that for the unique player $p \in P_1$ we have $c_p(S) = 1$, hence no deviation may lessen his cost and $P_1 \cap I = \emptyset$.

Let $j = \min\{j' \mid P_{j'} \cap I \neq \emptyset\}$; then $j \geq 2$, and $S'_j \cap E_j \neq \emptyset$. For all $j-1$ players $p_{j-1,i} \in P_{j-1}$ it holds that $S_{p_{j-1,i}} = E_j$, because $I \cap P_{j-1} = \emptyset$. Hence, $c_j(S') = j-1 + 1 = j = c_j(S)$. In any deviation of any coalition I , at least one player does not have incentive to deviate jointly with I and hence $\text{SPoA} \geq q$. Now for q we have $m = |\bigcup_j E_j| = \sum_{j=1}^q j = \frac{q(q+1)}{2}$, which yields $q \geq -\frac{1}{2} + \sqrt{2m + 1/4}$.

We convert the example into a network BCG . To grant access to players in P_{j-1} to facilities in E_j , we make a path of length 3, $\{(s_j, u_{ji}), (u_{ji}, v_{ji}), (v_{ji}, t)\}$, for every facility $e_i^j \in E_j$, $i \leq j-1$ and a length-2 path $\{(s_j, u_{jj}), (u_{jj}, t)\}$ for e_j^j . Let A_j be the set of arcs in these paths. Node s_j is the source of all players in P_j and t is a common sink for all players. Now we add auxiliary arcs $A'_j = \{(v_{ji}, u_{j,i+1}) \mid i \in [j-1]\}$. And, finally, an arc (s_{j-1}, u_{j1}) , $j \in \{2, \dots, q\}$, by which players P_{j-1} gain access to A_j . For the last group of players we add an arc (s_q, t) . Let us illustrate the analog of configuration S on the constructed network. All players in $p_{ji} \in P_j$, $i \in [j]$, play the same path strategy:

$$\begin{aligned} S_{ji} &= \{(s_j, u_{j+1,1})\} \\ &\cup \{(u_{j+1,r}, v_{j+1,r}), (v_{j+1,r}, u_{j+1,r+1}) \mid r \in [j-1]\} \\ &\cup \{(u_{j+1,j}, v_{j+1,j}), (v_{j+1,j}, t)\} \end{aligned}$$

and $S_{iq} = (s_q, t)$ for $i \in [q]$. See Fig. 1a for an example with $q = 4$. The proof that S is strong is analogous to the proof given for the non-network example. For the optimal configuration we set $S_{ji}^* = \{(s_j, u_{ji}), (u_{ji}, v_{ji}), (v_{ji}, t)\}$, for each player $p_{ij} \in P_j$, $i < j$, and $S_{jj} = \{(s_j, u_{jj}), (u_{jj}, t)\}$. The number of links m is:




Fig. 1: Lower bound constructions for Strong and 2-Strong Equilibria on identical links.

$$m = \sum_{j=1}^q (|A_j| + |A'_j|) + q = \sum_{j=1}^q (3j - 1 + (j - 1)) + q - 1 = 2q^2 + q - 1$$

which yields $q \geq -\frac{1}{4} + \frac{1}{2}\sqrt{2 + 2m}$. □

4.1 Lower Bounds On k -Strong Equilibria

For the k -SPoA of symmetric and general BCGs with identical facilities we show:

Theorem 7. *The k -strong price of anarchy of linear BCGs is at least:*

1. $\lfloor \frac{m}{2k} \rfloor + 1$ for symmetric BCGs and $\lceil \frac{m+2}{6k} \rceil$ for symmetric network BCGs, when $2 \leq k \leq \frac{m}{2}$.
2. $\lceil \frac{m}{k-1} \rceil - 1$ in general, when $2 \leq k \leq \frac{3}{4} + \frac{1}{2}\sqrt{\frac{1}{4} + 2m}$.

The proofs of these results are deferred to the full version. Figure 1b presents a 2-strong equilibrium for 6 identical players and 34 identical links. The maximum latency over all links under this configuration is 3. The social optimum has cost 1 and emerges when all players use link-disjoint paths to reach t from s .

References

1. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact Price of Anarchy for Polynomial Congestion Games. In: Proceedings of the 23rd Annual Symposium on Theoretical Aspects of Computer Science (STACS), LNCS 3884. pp. 218–229. Springer (2006)

2. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 189–198. ACM Press (2007)
3. Aumann, R.J.: Acceptable points in games of perfect information. *Pacific Journal of Mathematics* 10, 381–417 (1960)
4. Awerbuch, B., Azar, Y., Epstein, A.: Large the price of routing unsplittable flow. In: Proceedings of the 37th ACM Symposium on Theory of Computing (STOC). pp. 57–66 (2005)
5. Awerbuch, B., Azar, Y., Richter, Y., Tsur, D.: Tradeoffs in worst-case equilibria. *Theoretical Computer Science* 361(2–3), 200–209 (2006)
6. Banner, R., Orda, A.: Bottleneck Routing Games in Communication Networks. *IEEE Journal on Selected Areas in Communication* 25(6), 1173–1179 (2007)
7. Busch, C., Magdon-Ismail, M.: Atomic Routing Games on Maximum Congestion. *Theoretical Computer Science* 410(36), 3337–3347 (2009)
8. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the 37th ACM Symposium on Theory of Computing (STOC). pp. 67–73. ACM Press (2005)
9. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms* 3(1) (2007)
10. Fiat, A., Kaplan, H., Levy, M., Olonetsky, S.: Strong Price of Anarchy for Machine Load Balancing. In: Proceedings of the 34th International Colloquium on Automata, Languages and Programming (ICALP), LNCS 4596. pp. 583–594. Springer (2007)
11. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B.: The price of Anarchy for Restricted Parallel Links. *Parallel Processing Letters* 16(1), 117–132 (2006)
12. Harks, T., Klimm, M., Möhring, R.: Strong nash equilibria in games with the lexicographical improvement property. In: Proceedings of the 5th Workshop on Internet and Network Economics (WINE), LNCS 5929. pp. 463–470. Springer (2009)
13. Immorlica, N., Li, E., Mirrokni, V.S., Schulz, A.S.: Coordination mechanisms for selfish scheduling. *Theoretical Computer Science* 410(17), 1589–1598 (2009)
14. Koutsoupias, E., Mavronicolas, M., Spirakis, P.G.: Approximate Equilibria and Ball Fusion. *Theory of Computing Systems* 36(6), 683–693 (2003)
15. Koutsoupias, E., Papadimitriou, C.H.: Worst-case Equilibria. In: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS), Springer LNCS 1563. pp. 404–413 (1999)
16. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. *Computer Science Review* 3(2), 65–69 (2009)
17. Papadimitriou, C.H.: Algorithms, Games, and the Internet. In: Proceedings of the 33rd ACM Symposium on Theory of Computing (STOC). pp. 749–753. ACM Press (2001)
18. Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2(1), 65–67 (1973)