

The Strong Price of Anarchy of Linear Bottleneck Congestion Games

Bart de Keijzer · Guido Schäfer ·
Orestis Telelis

the date of receipt and acceptance should be inserted later

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 analyze the (strong) Price of Anarchy of these games for a natural load balancing social cost objective, i.e., minimize the maximum latency of a facility. In our studies, we focus on Bottleneck Congestion games with *linear* latency functions. 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 derive upper and (asymptotically) matching lower bounds on the (strong) Price of Anarchy of these games. We also derive more refined bounds for several special cases of these games, including the cases of identical player weights, identical latency functions and symmetric strategy sets. Further, we provide lower bounds on the Price of Anarchy for k -strong equilibria.

Keywords algorithmic game theory · congestion games · inefficiency of equilibria · Price of Anarchy · strong equilibria

A preliminary version of this work appeared in [12].

Bart de Keijzer
Sapienza University of Rome. DIAG, Web Algorithmics and Data Mining Group. 00185 Rome, Italy. E-mail: dekeijzer@dis.uniroma1.it

Guido Schäfer
Centre for Mathematics and Computer Science (CWI). Networks and Optimization Group. Science Park 123, 1098XG Amsterdam, The Netherlands. E-mail: G.Schaefer@cwi.nl.
VU University Amsterdam. Department of Econometrics and Operations Research. De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands.

Orestis Telelis
Athens University of Economics and Business. Department of Informatics. 10434 Athens, Greece. E-mail: telelis@gmail.com.

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 the 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 environments, including identical [20], uniformly related [11, 15, 19, 20], restricted assignment [5, 15], and unrelated machines [2].

A natural generalization of Load Balancing games are *Bottleneck Congestion games* (BCGs) [6, 17]. Here, every player chooses a subset of the resources (also referred to as *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 that he experiences over the chosen facilities is minimized. Load Balancing games constitute a special case of BCGs in that the strategy space of each player contains only *singleton* subsets of facilities (machines), i.e., the strategy of each player is always a single facility. Another interesting special case of BCGs are *Bottleneck Network Routing games* [6]. In these games, the facilities are identified with the links of an underlying network and the players' strategies correspond to paths in the network. Despite their importance, BCGs have received relatively little attention in the literature and are far from being well-understood. In this paper we study the inefficiency of stable outcomes in these 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 reality because 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 interconnected 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 strategy deviations. 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 coordinated pure deviations of player subsets (also referred to as *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 ($k = 1$). Harks, Klimm and Möhring [17] showed that (under quite loose assumptions) BCGs 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. A natural social cost objective function that has been studied intensively in the context of Load Balancing games is to minimize the maximum latency of a facility. We will use it also here for our investigations of the inefficiency of equilibria for BCGs.

The *Price of Anarchy (PoA)* [20,21,23] 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.

1.1 Our Contributions

We study the inefficiency of both pure Nash equilibria and strong equilibria of BCGs. In our studies, we focus on BCGs with *linear* latency functions, 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). An overview of the main results obtained in this paper is given in Table 1. Here, n and m refer to the number of players and facilities, respectively.

Our main contributions are as follows:

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 $m - 1 \leq \text{SPoA} \leq m$. These results hold for linear BCGs in general.

For comparison, Gairing et al. [15] proved that $m - 1 \leq \text{PoA} \leq m$ for the case of Load Balancing games with uniformly related machines and restricted assignments (which is a special case of the linear BCGs considered here). Further, Banner and Orda in [6] showed that $\text{PoA} = \Theta(m)$ for *network* BCGs with *identical* latency functions of the form x^p (where p is a constant).

2. We also derive bounds for linear BCGs with *identical players*, i.e., when all players have the same weight. In general, for asymmetric BCGs, we establish an upper bound of $\text{SPoA} = O(\sqrt{n})$. For symmetric BCGs, i.e., when all players have the same strategy space, we derive an (exact) bound of $\text{SPoA} = 2$.

Previously, Gairing et al. [15] proved an upper bound $\text{PoA} = O(\log n / \log \log n)$ for Load Balancing games with uniformly related machines and restricted assignments (which is a special case of asymmetric linear BCGs).

BCG	id. linear latency		arb. linear latency (SPoA)	
	k -SPoA (lower)	SPoA	id. players	arb. players
sym.	$\max\{2, \lceil m/(2k) \rceil\}$	2	2	$O(m)$
asym.	$\max\left\{\left\lfloor -\frac{1}{2} + \sqrt{2m + \frac{1}{4}} \right\rfloor, \left\lfloor \frac{m}{k-1} \right\rfloor - 1 \right\}$	$\Theta(\sqrt{m})$	$O(\sqrt{n})$	$\Theta(m)$

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 $\text{SPoA} \geq m - 1$.

3. We also consider linear BCGs with *identical facilities*, i.e., when the latency functions of all facilities are the same. For asymmetric BCGs, we show that $\text{SPoA} = \Theta(\sqrt{m})$. Further, for symmetric BCGs, we show that $\text{SPoA} = 2$. In fact, the lower bound of 2 on the SPoA for symmetric BCGs with identical players mentioned above also holds for identical facilities. We also give elaborate lower bounds on the k -SPoA for symmetric and asymmetric BCGs (see Table 1 for details).

We note that Load Balancing games with *identical machines* constitute a special case of symmetric linear BCGs with identical facilities. For these games, Andelman, Feldman and Mansour [2] showed a lower bound of $2m/(m+1)$ on the SPoA, which approaches 2 as the number m of machines (which correspond to the facilities) goes to infinity. Our bound is exact and we show that the worst-case occurs even for a constant number of facilities and players. In [2] an exact bound of 2 was shown for the SPoA of Load Balancing games with *unrelated machines*; in this setting, the weights of the players may vary, depending on the machine they choose. Thus, the lower bound of 2 in this case is incomparable to our lower bound.

4. Finally, we also provide asymptotically tight worst-case examples for (directed) network linear BCGs (definitions will be given below).

1.2 Related Work

One of the earliest works concerning Network BCGs is by Caragiannis *et al.* [9]. The authors devised polynomial time algorithms for computing pure Nash equilibria, on single sink and/or single source networks, with linear latency functions. They also provided bounds on the PoA of equilibria for such networks. Network BCGs were also considered by Banner and Orda [6]. The authors showed existence of pure Nash equilibria and provided an $\Theta(m)$ bound on the PoA for identical network links with latency functions of the form x^p , where p is some constant. Busch and Magdon-Ismael [8] studied the PoA of network BCGs with identical players. Harks *et al.* [17] introduced general BCGs and showed that strong equilibria are guaranteed to exist in these games. In [16], the problem of computing pure Nash equilibria and strong equilibria is addressed. The paper shows several hardness results and proposes polynomial time algorithms for special cases. In a very recent work, Busch and Kannan [7] study the PoA of BCGs under the assumption that

the players have strategy spaces with bounded *stretch*; here the stretch constitutes a measure of variation in resource utilization.

As mentioned above, BCGs generalize Load Balancing games, which have been studied intensively in recent years. Load Balancing games were first studied by Koutsoupias and Papadimitriou in their seminal work [20], where they introduce the Price of Anarchy notion. Among other results, they prove for Load Balancing games with m identical machines a lower bound of $\Omega(\log m / \log \log m)$ on the PoA for *mixed* Nash equilibria. Czumaj and Vöcking [11] and Koutsoupias, Mavronicolas and Spirakis [19] proved a matching upper bound. For Load Balancing games with *uniformly related* machines, Czumaj and Vöcking [11] show that $\text{PoA} = \Theta(\log m / \log \log m)$ for mixed Nash equilibria.

Concerning the PoA of *pure* Nash equilibria, an upper bound of $\frac{2m}{m+1}$ for m identical machines follows from an early approximation analysis of a local search heuristic for scheduling by Finn and Horowitz [14]. Schuurman and Vredeveld [25] gave a matching lower bound example. Czumaj and Vöcking [11] proved that the pure PoA for Load Balancing games on uniformly related machines is $\Theta(\log m / \log \log m)$. Awerbuch et al. [5] showed that the pure PoA of Load Balancing games with m identical machines and restricted assignments is also $\Theta(\log m / \log \log m)$. Independently, Gairing et al. [15] obtained the same bounds. Further, the authors prove that for m uniformly related machines with restricted assignments, $m - 1 \leq \text{PoA} \leq m$. For a detailed coverage of these results we refer the reader to [26].

Andelman, Feldman and Mansour [2] were the first to study strong equilibria and k -strong equilibria in the context of Load Balancing games. For the case of m *unrelated* machines, they proved that $m \leq \text{SPoA} \leq 2m - 1$. Later, Fiat et al. [13] improved these bounds and showed that $\text{SPoA} = m$ for these games. The authors also proved that for uniformly related machines the SPoA is $\Theta(\log m / (\log \log m)^2)$. For results in the context of more general scheduling games and associated scheduling policies (called *coordination mechanisms*), the interested reader is referred to [18] and the references therein.

Bottleneck Congestion games owe their name to their similarity to *congestion games*, which were introduced by Rosenthal [24]. 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 [24] proved the existence of pure Nash equilibria in congestion games, by usage of potential function arguments. Monderer and Shapley proved in [22] that the class of congestion games coincides isomorphically with the class of *finite potential games*. The Price of Anarchy of pure Nash equilibria for congestion games was resolved by Christodoulou and Koutsoupias [10] and, independently, by Awerbuch, Azar and Epstein [4]. It is shown in [10] that $\text{PoA} = \Theta(\sqrt{n})$ for asymmetric linear congestion games and the social cost being the maximum over the players' costs, 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 [10]. 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 refer to as *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 which he can choose from. 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 $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$ 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 $N_e(S)$ to be 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 $e \in E$ is associated with a latency function $l_e : \Sigma \rightarrow \mathbb{R}^+$, which satisfies the following three properties (see also [17]):

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

In this paper, we focus on *linear* latency (thus, also cost) functions; we motivate our choice in Example 1 below. 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. Under a 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 facility in S_i with the highest latency, 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} for $(S_1, \dots, S_{i-1}, S_{i+1}, \dots, S_n)$. Similarly, for a subset of players $I \subseteq N$ we use S_I and S_{-I} to refer to the strategy profiles induced by the strategies of players in I and $N \setminus I$, respectively. For every facility $e \in E$ and every subset of players $I \subseteq N$, we use the notation $l_e(S_I)$ to denote the latency of e under the strategy profile S_I , induced *only* by the players in I (i.e., as if all players in $N \setminus I$ are absent).

Definition 1 A strategy profile $S \in \Sigma$ is a *k-strong equilibrium* if for every non-empty player subset $I \subseteq N$ with $|I| \leq k$ and for 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 strictly 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 [17] 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 (linear latency) 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

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

$S \in \Sigma$ 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, 20] 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.

We next give an example showing that the SPoA is unbounded, even for symmetric BCGs with arbitrary latency functions on the facilities satisfying Properties 1–3 above.

Example 1 We construct a BCG with $n = 2$ players and $m = 3$ facilities, denoted e_1, e_2 and e_3 . Let $\Sigma_1 = \{\{e_1\}, \{e_2\}\}$ and $\Sigma_2 = \{\{e_2\}, \{e_3\}\}$. The latency functions are defined as follows. Let $M \gg \varepsilon > 0$.

$$l_{e_1}(S) = \varepsilon, \quad l_{e_3}(S) = M, \quad \text{and} \quad l_{e_2}(S) = \begin{cases} M & \text{if } N_{e_2}(S) = \{1, 2\} \\ \varepsilon & \text{otherwise.} \end{cases}$$

The socially optimal strategy profile is $S^* = (S_1^*, S_2^*)$ with $S_1^* = \{e_1\}$ and $S_2^* = \{e_2\}$. We have $C(S^*) = \varepsilon$. Consider the strategy profile $S = (S_1, S_2)$, where $S_1 = \{e_2\}$ and $S_2 = \{e_3\}$. Clearly, $C(S) = M$. We claim that S is a strong equilibrium. Indeed, the only alternative strategy for player 1 is $\{e_1\}$, which will incur him an individual cost of $\varepsilon = c_1(S)$. This rules out a unilateral deviation of player 1, but also a joint deviation along with player 2. The only alternative strategy for player 2 is $\{e_2\}$, which will incur him a cost of $M = c_2(S)$, if he deviated unilaterally. The SPoA in this example is at least M/ε .

We can easily modify the above instance to make it a symmetric BCG: we include strategy $\{e_1\}$ in Σ_2 and define $l_{e_1}(S) = \infty$ if $2 \in N_{e_1}(S)$ (thus, $e_1 \in S_2$). Similarly, we include strategy $\{e_3\}$ in Σ_1 and define $l_{e_3}(S) = \infty$ if $1 \in N_{e_3}(S)$ (thus, $e_3 \in S_1$). Note that the latency functions of the resulting symmetric BCG satisfy Properties 1–3 above.

A final comment is in order. Structurally, the BCG defined in Example 1 appears to be identical to a Load Balancing game (such as the ones studied in [2]) with 2 players and 3 machines. The difference, however, is in the definition of the latency functions. In Load Balancing games the latency of each machine (facility) is additive in the players' loads (weights), whereas in our example the latencies are *super-additive*; this is especially visible in the definition of l_{e_2} above.

This example motivates our study 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. In particular, Load Balancing games with *uniformly related* [11] or *identical* machines, also involving *restricted assignments* (see [15]), constitute important special cases of linear BCGs.

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 .

Observe that the above example corresponds to a network BCG, with parallel links connecting a single source node to a single sink (each facility corresponds to a distinct link).

Unless stated otherwise, we assume subsequently that all player weights are at least one, i.e., $w_i \geq 1$ for every $i \in N$; this assumption is without loss of generality as we can always enforce it, by scaling the weights appropriately. Moreover, when studying *identical facilities*, we assume that the latency functions are identities, i.e., $l_e(S) = w_e(S)$ for every $e \in E$.

3 Facilities with Arbitrary Linear Latencies

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 Arbitrarily Weighted 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 Bottleneck Congestion games 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 the 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 $w_{\hat{e}}(S) = l_{\hat{e}}(S)/a_{\hat{e}} > k\gamma^*/a_{\hat{e}}$. Consider the strategies that the players in $I_{\hat{e}}$ choose under S^* and suppose for the sake of deriving a contradiction that for every $i \in I_{\hat{e}}$, $S_i^* \cap E_k \neq \emptyset$. Then, by the pigeonhole principle, 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 $e' \in S_j^* \setminus S_j$ such that:

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

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 utilizes all the 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 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 subset of players $T_\lambda \subseteq N \setminus I_\lambda$ 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 $x, y \in \mathbb{R}^+$. As a consequence, all our upper bounds on the SPoA that exploit the first part of Lemma 1 hold for sub-additive latency functions.

Theorem 2 *The strong Price of Anarchy of linear Bottleneck Congestion games 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 Bottleneck Congestion games and at least $\frac{m+1}{3}$ in single-sink linear network Bottleneck Congestion games.*

Proof We describe a directed network BCG; the general case will follow directly. Define $0! = 1$ and let player $i \in [n]$ have weight $w_i = 1/(i-1)!$ and a source vertex s_i . There is a single sink t for all players and $n-1$ auxiliary vertices v_i , $i = 2, \dots, n$. The set of arcs is $E = E_1 \cup E_2 \cup \{(s_1, t), (s_n, t)\}$, where:

$$\begin{aligned} E_1 &= \{(s_i, v_{i+1}) \mid i = 1, \dots, n-1\} \cup \{(s_i, v_i) \mid i = 2, \dots, n\} \\ E_2 &= \{(v_i, t) \mid i = 2, \dots, n\} \end{aligned}$$

Then, $m = |E| = |E_1| + |E_2| + 2 = 2(n-1) + (n-1) + 2 = 3n-1$. For each arc $e \in E_1$ set $a_e = 1$. For $i = 2, \dots, n-1$ let $a_{(v_i, t)} = (i-1)!$. Also, set $a_{(s_1, t)} = 1$ and $a_{(s_n, t)} = n!$. An example for the case $n = 4$ appears in Figure 1. Each player has two strategies, an *upper* path $\{(s_i, v_{i+1}), (v_{i+1}, t)\}$, $i = 1, \dots, n-1$, and a *lower* path $\{(s_i, v_i), (v_i, t)\}$, $i = 2, \dots, n$. The upper path of player n is $\{(s_n, t)\}$ and the lower path of player 1 is $\{(s_1, t)\}$.

Under configuration S , wherein all players play their *upper* paths we have $c_i(S) = i!/(i-1)! = i$, thus $C(S) = n = \frac{m+1}{3}$. We claim S is a strong equilibrium. Consider deviation of any coalition $I \subseteq N$ and call S' the resulting profile. Let i be a player of *minimum* index in I and assume first $i \geq 2$. Then, under S and S' , $i-1$ plays his upper path, $\{(s_{i-1}, v_i), (v_i, t)\}$. The only deviation available to i is $\{(s_i, v_i), (v_i, t)\}$. Then:

$$\begin{aligned} c_i(S') &= \max\left\{l_{(s_i, v_i)}(S'), l_{(v_i, t)}(S')\right\} \\ &= \max\left\{\frac{1}{(i-1)!}, (i-1)! \cdot \left(\frac{1}{(i-1)!} + \frac{1}{(i-2)!}\right)\right\} \\ &= \max\left\{\frac{1}{(i-1)!}, 1 + (i-1)\right\} = i = c_i(S) \end{aligned}$$

Player 1 will not participate in any coalition, because a cost of 1 is incurred to him under S and when playing his lower path. Then S is a strong equilibrium. In the socially optimum configuration every player plays his lower path alone and has cost $c_i(S^*) = (i-1)!/(i-1)! = 1$. Thus $\text{SPoA} \geq \frac{m+1}{3}$. For the general case we simply regard links in $\{(s_1, t)\} \cup E_2 \cup \{(s_n, t)\}$ as $m = n+1$ facilities e_j , $j \in [n+1]$ and restrict every player's i strategy space to $\{e_i, e_{i+1}\}$. If every player plays e_{i+1} we obtain a strong equilibrium similar to S described above. The social




Fig. 1 An example of a single-sink linear network BCG used to prove the lower bound of Theorem 3. There are $n = 4$ players and $m = 11$ links. Player i has weight w_i and source node s_i . Each link is labeled with its latency factor. If every player plays the “lower” path available to him, the maximum latency on any facility is 1. If every player plays the “upper” path available to him, the maximum latency on any facility is $4 = (m + 1)/3$.

optimum occurs when player i plays e_i . We omit a detailed analysis because a similar construction appeared in [15]. \square

3.2 Identically Weighted 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 of linear Bottleneck Congestion games with identically weighted players is, in general, at most $\frac{1}{2} + \sqrt{2n - \frac{3}{2}}$; it is exactly 2 for symmetric such games.*

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 $\alpha \leq 1/2 + \sqrt{2n - 3/2}$.




Fig. 2 A symmetric network linear BCG instance with identical links as used in the proof of Theorem 4. The strong Price of Anarchy is at least 2 on this instance.

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

The following example establishes the tightness of this bound. Let $N = [3]$ and $E = [6]$. Every player $i \in N$ has weight $w_i = 1$ and every facility $e \in E$ incurs latency $a_e = 1$ per unit of load. The strategy set of every player is:

$$\{\sigma_1 = \{e_1\}, \sigma_2 = \{e_2, e_3\}, \sigma_3 = \{e_4, e_5\}, \sigma_4 = \{e_2, e_5, e_6\}\}.$$

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 (symmetric) network BCG, as depicted in Figure 2; define the directed arcs $e_1 = (s, t)$, $e_2 = (s, u)$, $e_3 = (u, t)$, $e_4 = (s, v)$, $e_5 = (v, t)$, and $e_6 = (u, v)$. Then each strategy corresponds to an (directed) $s - t$ path. \square

4 Facilities with Identical Linear Latencies

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 Bottleneck Congestion games with identical facilities is at most $-\frac{1}{2} + \sqrt{2m + \frac{1}{4}}$, in general, and exactly 2 in the 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 i , 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 (with a directly corresponding network example in Figure 2).

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:

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

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, $\alpha \leq -1/2 + \sqrt{2m + 1/4}$. Since $\text{SPoA} \leq \alpha$ the claim follows. \square

Theorem 6 *The strong Price of Anarchy of linear Bottleneck Congestion games with identical players and identical facilities is at least $\left[-\frac{1}{2} + \sqrt{2m + \frac{1}{4}}\right]$. For single-sink linear network Bottleneck Congestion games, it is at least $\left[-\frac{1}{4} + \frac{1}{2}\sqrt{2 + 2m}\right]$.*

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:

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

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,




Fig. 3 An example for the SPoA on a single-sink network BCG with 35 identical links and 10 identical players. The strong equilibrium configuration consists of the dotted paths for all players with the same source node and incurs a maximum latency of 4 on the uppermost link in the figure. The social optimum has a cost of 1 and occurs if each player utilizes a distinct (lowermost possible) path from his source node to t . Then, $\text{SPoA} = 4 \geq -\frac{1}{4} + \frac{1}{2}\sqrt{2} + 2m$ (Theorem 6).

$c_j(S') = j - 1 + 1 = j = c_j(S)$. In any deviation of any coalition I , at least one player does not have an incentive to deviate jointly with I and, hence, $\text{SPoA} \geq q$. Now q is the largest integer satisfying $m = |\cup_j E_j| \geq \sum_{j=1}^q j = q(q+1)/2$, which yields $q = \lfloor -1/2 + \sqrt{2m+1}/4 \rfloor$.

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:

$$S_{ji} = \left\{ (s_j, u_{j+1,1}) \right\} \\ \cup \left\{ (u_{j+1,r}, v_{j+1,r}), (v_{j+1,r}, u_{j+1,r+1}) \mid r \in [j] \right\} \\ \cup \left\{ (u_{j+1,j+1}, t) \right\}$$

and $S_{iq} = (s_q, t)$ for $i \in [q]$. See Figure 3 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 $S_{jj} = \{(s_j, u_{jj}), (u_{jj}, t)\}$. We identify q as the largest integer satisfying for the total number of links m :

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

which yields $q = \lfloor -\frac{1}{4} + \frac{1}{2}\sqrt{2 + 2m} \rfloor$. \square

4.1 k -Strong Equilibria

In this subsection, we derive more refined lower bounds for the k -SPoA for asymmetric and symmetric BCGs with identical facilities. Particularly, in the symmetric case, we exhibit a lower bound for the k -SPoA of symmetric *network* linear BCGs which is of the same order as for general symmetric linear BCGs (Theorem 8 below).

Theorem 7 *For any $k \geq 2$, the k -strong Price of Anarchy of linear Bottleneck Congestion games is at least $\max \left\{ \lfloor \frac{m}{k-1} \rfloor - 1, \left\lfloor -\frac{1}{2} + \sqrt{2m + \frac{1}{4}} \right\rfloor \right\}$.*

Proof Take an instance with a set E of m identical facilities and a set N of $n = |N| = m$ identical players (of unit weights). For $k \leq \lfloor \sqrt{m} + 1 \rfloor$, we divide the facilities and the players into $p = k - 1$ groups $\langle E_0, N_0 \rangle, \langle E_1, N_1 \rangle, \dots, \langle E_{p-1}, N_{p-1} \rangle$. For each group $\langle E_r, N_r \rangle$, we place $\lfloor \frac{m}{k-1} \rfloor$ facilities in E_r and $\lfloor \frac{m}{k-1} \rfloor$ players in N_r . The remaining at most $k - 2$ players and $k - 2$ facilities are distributed into the existent groups (one player and facility per group). These players and facilities we call *residual*. We shall not be concerned with residual players and facilities in our construction; for a residual player $i \in N_r$, we will have the corresponding residual facility, $e \in E_r$, to be the player's only available strategy. Moreover, the residual facility e will not be part of any other player's strategy. Thus, we may ignore existence of residual players and facilities, as they will not affect the strategic choices of other players, or be affected by them. Then, define $q = \lfloor \frac{m}{k-1} \rfloor$ and, without loss of generality, we assume $q = |E_r| = |N_r|$, for $r = 0, \dots, p - 1$; we name the facilities in E_r by e_0^r, \dots, e_{q-1}^r . Let $[\cdot]_q$ denote $(\cdot) \bmod q$. For $r = 0, \dots, p - 1$, we define over E_r two kinds of strategies, called s_j^r and σ_j^r , for every $j \in \{0, \dots, q - 1\}$:

$$s_j^r = \left\{ e_{j+1}^r, \dots, e_{[j+q-1]_q}^r \right\}, \quad \sigma_j^r = \{ e_j^r \}$$

Now any specific subset N_r contains q players. Assume w.l.o.g. that they are indexed by $i \in \{0, \dots, q - 1\}$. For a socially optimal configuration, S^* , we set $S_i^* = \sigma_i^{[r+1]_p}$ for $i = 0, \dots, q - 1$. Then, each player uses exactly one facility under S^* and $C(S^*) = 1$.

Now consider the configuration S , where for each $i \in N_r$, $S_i = s_i^r$. We make this strategy, along with $\sigma_i^{[r+1]_p}$, to be the only two available strategies for i . Under S , every player from N_r plays $|E_r| - 1 = q - 1$ distinct facilities from E_r and each different player has a different facility of E_r left out of his strategy, S_i . Thus, $C(S) = q - 1 = \lfloor \frac{m}{k-1} \rfloor - 1$. An example of a 3-strong such configuration S with $m = 10$ facilities and $n = m = 10$ players is described in Figure 4, along with the previously described socially optimal configuration.




Fig. 4 An example for the 3-SPoA of asymmetric linear BCGs with identical facilities (Theorem 7), involving $m = 10$ facilities (depicted as boxes) and $n = m = 10$ identical players. Each facility is labeled (inside the box) with the index of a single player using it in the socially optimal configuration. Player indices on the top denote players using each facility at 3-strong equilibrium.

We argue that S is a k -strong equilibrium. Consider any coalition, $I \subseteq N$, of players that attempt to switch to their socially optimal strategies. We will show that $|I| \geq k + 1$. First, we show that I must contain at least one player from every group. Indeed, let $i_r \in I \subseteq N_r$ for any $r = 0, \dots, p - 1$. Then, i_r has an optimum singleton strategy in $E_{[r+1]_p}$. Thus, at least two players from $N_{[r+1]_p}$ (using under S the single optimum facility, e_i^r , of i_r) must belong to I , so that it is beneficial for i_r to deviate to his (only alternative) optimum strategy. Then, we deduce inductively that I must contain players from each of the $p = k - 1$ distinct groups. Secondly, we observe that, from each group, there must exist in I at least 2 players from each group, so that a beneficial (socially optimal) strategy is created for some other player of I (outside the group), to switch to. Otherwise, some player will not be able to strictly improve his cost by deviating jointly within I . Thus, we obtain $|I| \geq k + 1$.

For an illustration of the argument, consider the example of Figure 4 which depicts a 3-strong equilibrium. Only a joint deviation of 4 (or more) players is possible: each player in $I = \{0, 1, 7, 8\}$ can switch jointly to their optimal strategies and improve their cost by 1.

Combination of the lower bound $C(S)/C(S^*)$ obtained above, with the one shown earlier in Theorem 6 for strong equilibria (that are also k -strong for any value of k), yields the result. \square

For the social inefficiency of k -strong equilibria in *symmetric* linear BCGs we obtain the following result.

Theorem 8 *The k -strong Price of Anarchy of symmetric linear Bottleneck Congestion games is at least $\max\{2, \lceil m/(2k) \rceil\}$. It is at least $\max\{2, \lceil (m+2)/(6k) \rceil\}$ for symmetric network linear Bottleneck Congestion games.*

Proof Take n identical players (of unit weight) and a set of $m = 2n$ identical facilities $E \cup F$, where $E = \{e_j \mid j = 0, \dots, n-1\}$, $F = \{f_j \mid j = 0, \dots, n-1\}$. Without loss of generality, we index the players by $i = 0, \dots, n-1$. Let us first define the strategy space of player i . A strategy of player i consists of any contiguous ordered sequence of facilities e_j, \dots, e_t , followed by any contiguous ordered sequence of facilities f_t, \dots, f_i . We allow sequences of facilities in E or F , that “wrap around” modulo n . In the socially optimal configuration player i plays $\{e_i, f_i\}$ and the social cost is 1.

Now we demonstrate a k -strong equilibrium configuration. Let $q = \lceil \frac{n}{k} \rceil - 1$. Set the strategy of player i , to $S_i = \{\{e_i, \dots, e_{[i+q]_n}\}, \{f_{[i+q]_n}, \dots, f_{[i+2q]_n}\}\}$, where for any $i = 0, \dots, n-1$, we define $[i+q]_n = (i+q) \bmod n$. This way each player

uses $2(q + 1)$ facilities, $q + 1$ from set E and $q + 1$ from set F . Now we note that every facility in E or F is the first for some player, the second for another and, continuing in the same manner, the $(q + 1)$ -th for some player (distinct from the previous ones). Hence every facility is used by exactly $q + 1$ players, which yields $C(S) = \lceil \frac{n}{k} \rceil = \lceil \frac{m}{2k} \rceil$.

We claim that S is a k -strong equilibrium. Consider deviation of any coalition $I \subseteq N$, $|I| \leq k$. We examine the number of valid beneficial strategies that emerge under S_{-I} , for members of I to adopt in the new configuration. First we notice that removal of any pair of players i_1, i_2 from the configuration does not incur a strictly more beneficial valid strategy for any of them. In particular, the most beneficial strategy will include either $\{e_{[i_1+q]_n}, f_{[i_1+q]_n}\}$ or $\{e_{[i_2+q]_n}, f_{[i_2+q]_n}\}$. However S_{i_1} and S_{i_2} may intersect in at most one facility from each of these sets. Hence there is a facility with cost $C(S) - 1$ in both sets. Now we consider successive removal of players i_r , $r = 3, \dots, k$. Strategy S_{i_r} can intersect at best with the strategy of one of i_1, i_2 at $\{e_{[i_r+q]_n}\}$ and with the strategy of the other at $\{f_{[i_r+q]_n}\}$. Thus, removal of i_3 creates a beneficial strategy $\{e_{[i_3+q]_n}, f_{[i_3+q]_n}\}$ of cost $C(S) - 2$, for one player (out of i_1, i_2, i_3). At most q such strategies may be created by removal of q players whose strategies intersect “in their middle” with S_{i_1}, S_{i_2} appropriately. In any case, we may not create more than $k - 2$ beneficial strategies, by removal of k players. Thus S is k -strong.

By combining the derived lower bound $C(S)/C(S^*)$ with the lower bound from Theorem 4, we conclude the proof of the first statement of the result.

For the second statement, we present a network version of the example discussed above. A link is made for each facility e_j , which is followed by a link for facility f_j . Arc e_j leads by an additional link to arc $e_{[j+1]_n}$. Similarly for each arc f_j and $f_{[j+1]_n}$. Finally, we use $n - 1$ additional arcs for making all players emanate from a single source and $n - 1$ arcs to guide them to a single sink. The source is the tail vertex of some arc in E and the sink is the head vertex of some arc in F . In total we use $m = 6n - 2$ arcs, which yields a k -SPoA lower bound of $\lceil \frac{n}{k} \rceil = \lceil \frac{m+2}{6k} \rceil$. By combining this bound with the lower bound of Theorem 4 (which is a network example, as depicted in Figure 2), we conclude the proof of the second statement.

For an illustration of this construction, Figure 5 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 . \square

5 Summary and Open Problems

In this work, we derived inefficiency bounds for equilibria of linear Bottleneck Congestion games. These games generalize Load Balancing games, which have been investigated extensively in the literature [21, 11, 15]. In particular, we proved upper and lower bounds on the Price of Anarchy of pure Nash equilibria and strong equilibria of linear BCGs. We considered several special cases of these games, including identical players, identical facilities and symmetric strategy spaces. For most of our lower bounds we were able to provide asymptotically equivalent *network* worst-case examples, wherein the players’ strategies constitute paths over directed networks.




Fig. 5 A lower bounding example for the 2-SPoA of symmetric network linear BCGs with identical links (Theorem 8), involving 6 identical players. Player indices mark the links used by each player at 2-strong equilibrium. In the socially optimal configuration each of the players uses a distinct path from s to t , that is disjoint from the paths used by other players. 6 such paths are visibly existent in the figure.

There are still several problems that remain to be resolved, towards completing our understanding of the social inefficiency of equilibria in linear BCGs. In particular, we miss tight upper and lower bounds for symmetric linear BCGs (with arbitrarily weighted players). Additionally, it would be interesting to derive tight upper bounds for the k -strong Price of Anarchy for BCGs with identical facilities to obtain a complete picture of the model’s performance in dependence on the size of the coalitions that can deviate. Finally, it would be also very interesting to derive tight bounds for networks BCGs.

Acknowledgements. We thank two anonymous reviewers who helped us significantly in improving the presentation of this work.

Orestis Telelis acknowledges support by the research project “DDCOD” (PE6-213). The project is implemented within the framework of the Action “Supporting Postdoctoral Researchers” of the Operational Program “Education and Lifelong Learning” (Action’s Beneficiary: General Secretariat for Research and Technology), and is co-financed by the European Union (European Social Fund – ESF) and the Greek State.

References

1. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. In B. Durand and W. Thomas, editors, *STACS*, volume 3884 of *LNCS*, pages 218–229. Springer, 2006.
2. N. Andelman, M. Feldman, and Y. Mansour. Strong Price of Anarchy. In N. Bansal, K. Pruhs, and C. Stein, editors, *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 189–198. SIAM, 2007.
3. R. J. Aumann. Acceptable Points in Games of Perfect Information. *Pacific Journal of Mathematics*, 10:381–417, 1960.
4. B. Awerbuch, Y. Azar, and A. Epstein. Large the Price of Routing Unsplittable Flow. In H. N. Gabow and R. Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC’05)*, pages 57–66. ACM, 2005.
5. B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in Worst-Case Equilibria. *Theoretical Computer Science*, 361(2–3):200–209, 2006.

6. R. Banner and A. Orda. Bottleneck Routing Games in Communication Networks. *IEEE Journal on Selected Areas in Communication*, 25(6):1173–1179, 2007.
7. C. Busch and R. Kannan. Stretch in Bottleneck Games. In J. Gudmundsson, J. Mestre, and T. Viglas, editors, *COCOON*, volume 7434 of *LNCS*, pages 592–603. Springer, 2012.
8. C. Busch and M. Magdon-Ismail. Atomic Routing Games on Maximum Congestion. *Theoretical Computer Science*, 410(36):3337–3347, 2009.
9. I. Caragiannis, C. Galdi, and C. Kaklamanis. Network Load Games. In X. Deng and D.-Z. Du, editors, *ISAAC*, volume 3827 of *LNCS*, pages 809–818. Springer, 2005.
10. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In H. N. Gabow and R. Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73. ACM, 2005.
11. A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. *ACM Transactions on Algorithms*, 3(1), 2007.
12. B. de Keijzer, G. Schäfer, and O. Telelis. On the Inefficiency of Equilibria in Linear Bottleneck Congestion Games. In S. C. Kontogiannis, E. Koutsoupias, and P. G. Spirakis, editors, *SAGT*, volume 6386 of *LNCS*, pages 335–346. Springer, 2010.
13. A. Fiat, H. Kaplan, M. Levy, and S. Olonetsky. Strong Price of Anarchy for Machine Load Balancing. In L. Arge, C. Cachin, T. Jurdzinski, and A. Tarlecki, editors, *ICALP*, volume 4596 of *LNCS*, pages 583–594. Springer, 2007.
14. G. Finn and E. Horowitz. A Linear Time Approximation Algorithm for Multiprocessor Scheduling. *BIT*, 19:312–320, 1979.
15. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The Price of Anarchy for Restricted Parallel Links. *Parallel Processing Letters*, 16(1):117–132, 2006.
16. T. Harks, M. Hofer, M. Klimm, and A. Skopalik. Computing Pure Nash and Strong Equilibria in Bottleneck Congestion Games. In M. de Berg and U. Meyer, editors, *ESA (2)*, volume 6347 of *LNCS*, pages 29–38. Springer, 2010.
17. T. Harks, M. Klimm, and R.H. Möhring. Strong Nash Equilibria in Games with the Lexicographical Improvement Property. In S. Leonardi, editor, *WINE*, volume 5929 of *LNCS*, pages 463–470. Springer, 2009.
18. N. Immorlica, E. Li, V. S. Mirrokni, and A. S. Schulz. Coordination Mechanisms for Selfish Scheduling. *Theoretical Computer Science*, 410(17):1589–1598, 2009.
19. E. Koutsoupias, M. Mavronicolas, and P. G. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.
20. E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In C. Meinel and S. Tison, editors, *STACS*, volume 1563 of *LNCS*, pages 404–413. Springer, 1999.
21. E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. *Computer Science Review*, 3(2):65–69, 2009.
22. D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14:124–143, 1996.
23. C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.
24. R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
25. P. Schuurman and T. Vredeveld. Performance Guarantees of Local Search for Multiprocessor Scheduling. *INFORMS Journal on Computing*, 19(1):52–63, 2007.
26. Berthold Vöcking. Selfish Load Balancing. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.