One-Round Walks in Linear Congestion Games

Vittorio Bilò[†]

[†] Dipartimento di Matematica "Ennio De Giorgi", Università del Salento

Provinciale Lecce-Arnesano P.O. Box 193, 73100 Lecce, Italy.

vittorio.bilo@unisalento.it

Abstract

Which is the approximation ratio of the solutions achieved after a one-round walk starting from the empty strategy profile in linear congestion games for the social function defined as the sum of the players' costs? The exact answer has been given only for the cases of unweighted players with unrestricted sets of strategies and weighted players with singleton strategies on identical resources. In this note, we survey results, techniques and open problems from this research topic.

1 Introduction

After the publication of the seminal papers by Koutsoupias and Papadimitriou [15] and by Anshelevich et al. [2], a tremendous quantity of results analyzing

Bulletin of the EATCS no 101, pp. 45–63, June 2010

the prices of anarchy and stability of non-cooperative systems has been produced. Just to get an idea of this phenomenon, a rough search for the number of citations obtained by these two papers, carried out on GOOGLE SCHOLAR, returns the values of 867 and 276, respectively.

Hence, inefficiencies of equilibria solutions have been nicely characterized in a variety of applications and, not rarely, the discovered bounds have been proved to be surprisingly low. For example, Roughgarden and Tardos [17] showed that, in non-atomic selfish routing with linear latencies, the price of anarchy coincides with that of stability and it is exactly $4/3$ for pure Nash equilibria. As to its discrete counterpart, represented by congestion games with linear latencies (from now on, linear congestion games), Christodoulou and Koutsoupias [9, 10] set the price of anarchy of pure and mixed Nash equilibria to 2.5, while Awerbuch et al. [3] derived an exact 2.618 bound for their generalization to weighted players. Christodoulou and Koutsoupias [10] also contributed to the determination of the price of stability of pure Nash equilibria in linear congestion games, which was shown to be $1 + \sqrt{3}/3$ in conjunction with the result of Caragiannis et al. [7].

Nevertheless, computing pure Nash equilibria is being proved to be PLS-complete for more and more games represented in succinct form. At the same time, a recent escalation of results [13, 14, 8] has definitively shown that computing mixed Nash equilibria is PPAD-complete even for games represented in strategic form and for any number of players. Because of these limitations, the notions of price of anarchy and stability of either pure and mixed Nash equilibria can only be considered as theoretical measures of the inefficiency of solutions to which games may tend, but which are unlikely to be achieved in practical settings. Thus, studying the quality of polynomially computable outcomes, representing reasonable relaxations of Nash equilibria, is of great practical interest. To this aim, one can consider either approximate Nash equilibria (see [11] for an application to congestion games) or solutions achieved after a (polynomially) bounded number of best responses.

According to the latter approach, Mirrokni and Vetta [16] defined the notions of covering walks and one-round walks. Covering walks are sequences of best responses in which each player is allowed for at least one of them. One-round walks are covering walks in which each player is allowed for at most one best response, that is, they are sequences of best responses in which each player is allowed for exactly one of them.

In this note, we focus on the solutions achieved after a one-round walk starting from the empty strategy profile in linear congestion games when the social function measuring the quality of a given solution is defined as the sum of the players' costs. This approach is justified by the result of Ackermann et al. [1] who showed that computing pure Nash equilibria in linear congestion games is PLS-complete. In particular, we host results originally given in [4, 6, 7, 12].

2 Definitions

Strategic Games. A *strategic game* is defined by a triple $SG = (P, \Sigma_{i\in P}, U_{i\in P})$, where $P = \{1, \ldots, n\} := [n]$ denotes the set of *n players*, Σ_i the set of *strategies* for player i and $U_i : \times_{i\in[n]}\Sigma_i \mapsto \mathbb{R}_{\geq 0}$ defines the *utility* that player i gets in any possible strategy profile.

Strategy Profiles. Let $\Sigma = \times_{i\in[n]}\Sigma_i$ be the set of *strategy profiles* of the game and $\mathbf{s} = (s_1, \ldots, s_n) \in \Sigma$ be a generic strategy profile (or *solution*) in which each player $i \in [n]$ chooses strategy $s_i \in \Sigma_i$. We call *empty strategy profile* the profile $\mathbf{s}^{\emptyset}$ in which $s_i^{\emptyset} = \emptyset$ for any $i \in [n]$, that is, the solution in which no player has already made any strategic choice.

Improving Deviations and Best Responses. Given a strategy profile $\mathbf{s}$ and a strategy $t_i \in \Sigma_i$, let $(\mathbf{s}_{-i} \diamond t_i) = (s_1, \ldots, s_{i-1}, t_i, s_{i+1}, \ldots, s_n)$ be the strategy profile obtained from $\mathbf{s}$ when player i changes her strategy from s_i to t_i. From now on, we assume that the utility functions $U_{i\in[n]}$ model a cost that each player wants to minimize. The strategy $t_i \in \Sigma_i$ is called an *improving deviation* for player i in profile $\mathbf{s}$, if $U_i(\mathbf{s}_{-i} \diamond t_i) < U_i(\mathbf{s})$. Furthermore, a *best response* for player i in $\mathbf{s}$ is a strategy $t_i^* \in \Sigma_i$ giving player i the minimum cost once fixed the strategies of the other players, i.e., such that $U_i(\mathbf{s}_{-i} \diamond t_i^*) \leq U_i(\mathbf{s}_{-i} \diamond t_i)$ for any strategy $t_i \in \Sigma_i$. Note that, for any profile $\mathbf{s}$, there always exists a best response for player i in $\mathbf{s}$, while there may exist a profile $\mathbf{s}$ (for example a pure Nash equilibrium) in which a player does not possess any improving deviation. Thus, best responses are also improving deviations only in those profiles in which the player can lower her cost by changing her strategy.

One-Round Walks. Assume that the players have been arranged according to a particular ordering. For the ease of presentation, we denote as i the ith player in such an ordering. A *one-round walk* $W = (\mathbf{s}^0, \ldots, \mathbf{s}^n)$ is an $(n+1)$-tuple of strategy profiles such that, for any $i \in [n]$, it holds $\mathbf{s}^i = (\mathbf{s}^{i-1}_{-i} \diamond t_i^*)$, where t_i^* is a best response for player i in $\mathbf{s}^{i-1}$. The profiles $\mathbf{s}^0$ and $\mathbf{s}^n$ are called the initial and the final profile of the walk, respectively. Note that, if player i does not possess any improving deviation in $\mathbf{s}^{i-1}$, it may be $\mathbf{s}^i = \mathbf{s}^{i-1}$. However, when $\mathbf{s}^0 = \mathbf{s}^{\emptyset}$ this can never happen, i.e., each player always has an improving deviation. Throughout this note, we will always consider one-round walks starting from the empty strategy profiles, that is, such that $\mathbf{s}^0 = \mathbf{s}^{\emptyset}$.

Congestion Games. A *congestion game* $CG = (P, R, \Sigma_{i\in P}, \ell_{r\in R}, U_{i\in P})$ is a succinctly represented game in which there is a set R of m *resources* to be shared among the n players in P. A strategy $s_i \in \Sigma_i$ for player i is a subset of resources, i.e., $\Sigma_i \subseteq 2^R$. Given a strategy profile $\mathbf{s}$ and a resource $r \in R$, the number of players using r in $\mathbf{s}$, called the *congestion* of r, is denoted by $c_r(\mathbf{s}) = |\{i \in [n] \mid r \in s_i\}|$. A *latency function* $\ell_r : \mathbb{N} \mapsto \mathbb{R}_{\geq 0}$ associates resource r a cost (or latency) depending on the number of players using r in a given profile. The *cost* of player i in $\mathbf{s}$ is defined as $U_i(\mathbf{s}) = \sum_{r \in s_i} \ell_r(c_r(\mathbf{s}))$. We will consider *linear congestion games*, that is, the subclass of congestion games $\mathcal{G}_{lin}$ in which $\ell_r(x) = a_r x + b_r$, with $a_r, b_r \in \mathbb{R}_{\geq 0}$. Interesting special cases are also the class $\mathcal{G}^1_{lin}$ of *singleton linear congestion games* in which all of the players' strategies are restricted to contain a single resource and the class $\mathcal{G}^1_{ide}$ of *singleton load balancing games* which are singleton congestion games with identical resources, i.e., all having the same latency function. For all such games, we also consider the generalization with weighted players. In such a case, each player i has an associated weight $w_i \in \mathbb{R}_{\geq 0}$, the congestion of a resource r in $\mathbf{s}$ becomes $c_r(\mathbf{s}) = \sum_{i: r \in s_i} w_i$, that is, the sum of the weights of all players using r in $\mathbf{s}$, and the domain of the latency functions is extended to the set $\mathbb{R}_{\geq 0}$.

Quality of the Solutions. We measure the quality of a given profile by means of the *social function* $\textsc{Sum}(\mathbf{s}) = \sum_{i \in [n]} U_i(\mathbf{s})$, defined as the sum of the players' costs (another standard measure is the maximum cost of a player, i.e., $\textsc{Max}(\mathbf{s}) = max_{i \in [n]} U_i(\mathbf{s})$). For the generalizations to weighted players, the definition of $\textsc{Sum}$ is extended as follows, $\textsc{Sum}(\mathbf{s}) = \sum_{i \in [n]} w_i U_i(\mathbf{s})$.

We call an *instance*, any triple $(CG, \mathbf{s}^*, W)$ such that CG is a congestion game, $\mathbf{s}^* = (s^*_1, \ldots, s^*_n)$ is an *optimal strategy profile* in CG with respect to $\textsc{Sum}$ and $W = (\mathbf{s}^\emptyset, \ldots, \mathbf{s}^n)$ is a one-round walk for CG. We assume that the players are numbered in the order in which they are allowed to choose best responses in the walk. The approximation ratio of the solution achieved in the instance $(CG, \mathbf{s}^*, W)$ is then defined as $Apx^1_\emptyset(CG, \mathbf{s}^*, W) = \frac{\textsc{Sum}(\mathbf{s}^n)}{\textsc{Sum}(\mathbf{s}^*)}$. Finally, for a given class of games $\mathcal{G}$, we denote as $Apx^1_\emptyset(\mathcal{G}) = \sup_{(CG \in \mathcal{G}, \mathbf{s}^*, W)} Apx^1_\emptyset(CG, \mathbf{s}^*, W)$ the worst-case approximation ratio of the solutions achieved after a one-round walk starting from the empty strategy profile in $\mathcal{G}$.

3 Singleton Load Balancing Games

In this section, we consider the most basic case $\mathcal{G}^1_{ide}$. For the ease of notation, given an instance $(CG \in \mathcal{G}^1_{ide}, \mathbf{s}^*, W)$, for any $j \in [m]$, we set $o_j = c_{r_j}(\mathbf{s}^*)$ and $n_j = c_{r_j}(\mathbf{s}^n)$. Note that, in such a setting, results are independent of the particular latency function, thus we can assume $\ell_r(x) = x$, which yields $\textsc{Sum}(\mathbf{s}) = \sum_{r \in R} c_r(\mathbf{s})^2$ and, consequently, $\textsc{Sum}(\mathbf{s}^n) = \sum_{j \in [m]} n_j^2$ and $\textsc{Sum}(\mathbf{s}^*) = \sum_{j \in [m]} o_j^2$.

In a first paper on this subject, Suri et al. [18] introduced the term "greedy inequality" to denote the inequality characterizing the definition of best responses and showed the following result.

Lemma 1 (Greedy inequality, [18]). *For any instance* $(CG \in \mathcal{G}^1_{ide}, \mathbf{s}^*, W)$*, it holds* $\sum_{j\in[m]} n_j^2 \leq \sum_{j\in[m]} \left(2n_j o_j + o_j\right)$.

Suri et al. [18] used the greedy inequality to show that $3,083 \leq Apx^1_\emptyset(\mathcal{G}^1_{ide}) \leq 4.236$. Such a gap was tremendously narrowed by Caragiannis et al. [7] as shown in what follows.

In order to upper bound the value of $Apx^1_\emptyset(\mathcal{G}^1_{ide})$, we need the following two lemmas.

Lemma 2 ([7]). *Let* $g(x,y) = 2xy + (1+\xi)y - \xi x$ *and* $h(x,y) = \frac{x^2}{\psi^2} + \psi y^2$*, where* $\xi = \frac{7\sqrt{21}-3}{30}$ *and* $\psi = \frac{9+\sqrt{21}}{6}$*. For any non-negative integers* x, y *such that either* $x \neq 1$ *or* $y \neq 1$*, it holds* $g(x,y) \leq h(x,y)$*. Moreover,* $g(0,1) + g(1,1) = h(0,1) + h(1,1)$.

We say that a resource r_j is of type x/y if $n_j = x$ and $o_j = y$.

Lemma 3 ([7]). *Among the instances yielding the highest value for* $Apx^1_\emptyset(\mathcal{G}^1_{ide})$*, there is one in which the number of resources of type* $0/1$ *is not smaller than the number of resources of type* $1/1$.

Theorem 1 ([7]). $Apx^1_\emptyset(\mathcal{G}^1_{ide}) \leq \frac{2}{3}\sqrt{21} + 1 \approx 4.055$.

Proof. Consider an instance $(SG \in \mathcal{G}^1_{ide}, \mathbf{s}^*, W)$ yielding the highest value for $Apx^1_\emptyset(\mathcal{G}^1_{ide})$. Because of Lemma 3, we can assume that $\mathbf{s}^*$ and $\mathbf{s}^n$ are such that the number of resources of type $0/1$ is not smaller than the number of resources of type $1/1$. Let S be a function associating to each resource of type $1/1$ a resource of type $0/1$. Denote by F the set of resources of type $1/1$ and by $S = \bigcup_{r\in F} S(r)$ the set of resources of type $1/1$ which are associated with a resource in F. By the greedy inequality, the fact that $\sum_{j\in[m]} n_j = \sum_{j\in[m]} o_j$, the definitions of functions g and h, and by Lemma 2, we obtain

$$
\begin{aligned}
\sum_{j\in[m]} n_j^2 &\leq \sum_{j\in[m]} \left(2n_j o_j + o_j\right) \\
&= \sum_{j\in[m]} \left(2n_j o_j + \frac{27+7\sqrt{21}}{30} o_j - \frac{7\sqrt{21}-3}{30} n_j\right) \\
&= \sum_{j\notin S\cup F} g(n_j, o_j) + \sum_{j\in F} \left(g(n_{S(j)}, o_{S(j)}) + g(n_j, o_j)\right) \\
&\leq \sum_{j\notin S\cup F} h(n_j, o_j) + \sum_{j\in F} \left(h(n_{S(j)}, o_{S(j)}) + h(n_j, o_j)\right) \\
&= \frac{9-\sqrt{21}}{10} \sum_{j\in[m]} n_j^2 + \frac{9+\sqrt{21}}{6} \sum_{j\in[m]} o_j^2.
\end{aligned}
$$

Hence, it follows $\sum_{j\in[m]} n_j^2 \leq \left(\frac{2}{3}\sqrt{21}+1\right)\sum_{j\in[m]} o_j^2$. □

An almost matching lower bound is provided by the following result.

Theorem 2 **([7]).** *For any $\epsilon > 0$, it holds $Apx_{\emptyset}^1(\mathcal{G}_{ide}^1) \geq 4 - \epsilon$.*

Proof. For any fixed integer $k > 0$, let m be an integer such that m is a multiple of j^2, for any $j \in [k]$. We define a singleton load balancing game with m resources and $n = m\sum_{j=1}^{k}\frac{1}{j^2}$ players partitioned in k groups $g_1, \ldots, g_k$ such that $|g_j| = \frac{m}{j^2}$ for each $j \in [k]$. Denoted as p_i^j the ith player in the jth group, the set of strategies are such that p_i^j can only choose one of the first i resources.

Consider the solution $\mathbf{s}$ in which each player p_i^j chooses resource r_i. Since, for any $\frac{m}{(j+1)^2} + 1 \leq i \leq \frac{m}{j^2}$, the number of players using r_i is equal to $|g_j| - |g_{j+1}|$, we obtain

$$\begin{aligned}
\textsc{Sum}(\mathbf{s}) &= \sum_{j=1}^{k-1}\left(j^2(|g_j| - |g_{j+1}|)\right) + k^2|g_k| \\
&= m + m\sum_{j=1}^{k-1}\left(j^2\left(\frac{1}{j^2} - \frac{1}{(j+1)^2}\right)\right) \\
&= m\left(1 + 2\sum_{j=1}^{k-1}\frac{1}{i+1} - \sum_{j=1}^{k-1}\frac{1}{(i+1)^2}\right) \\
&= m\left(2\sum_{j=1}^{k}\frac{1}{i} - \sum_{j=1}^{k}\frac{1}{i^2}\right) \\
&< m\left(2H_k - \frac{\pi^2}{6}\right).
\end{aligned}$$

Hence, we have $\textsc{Sum}(\mathbf{s}^*) \leq m\left(2H_k - \frac{\pi^2}{6}\right)$.

Now consider the one-round walk W in which players p_i^j are arranged in non-increasing order according to the index i, and such that they always break ties in favor of the resource with smallest index when they have more than a choice of minimum cost. Once fixed the particular ordering of the players in W, we say that a player i belongs to row_j if $c_{t_i^*}(\mathbf{s}^i) = j$. By exploiting the fact that the number of resources with congestion i in $\mathbf{s}^n$ is equal to $|row_i| - |row_{i+1}|$, it can be verified that $|row_{2i}| = \frac{m}{(i+1)^2}$ and $|row_{2i-1}| = \frac{m}{i(i+1)}$ for any $i \in [k-1]$. Thus, even by simply considering the resources with congestion of at most $2k-4$, we obtain

$$\begin{aligned}
\textsc{Sum}(\mathbf{s}^n) &\geq \sum_{i=1}^{k-2}\Big((2i-1)^2(|row_{2i-1}|-|row_{2i}|)+(2i)^2(|row_{2i}|-|row_{2i+1}|)\Big)\\
&= m\sum_{i=1}^{k-2}\left((2i-1)^2\left(\frac{1}{i(i+1)}-\frac{1}{(i+1)^2}\right)+(2i)^2\left(\frac{1}{(i+1)^2}-\frac{1}{(i+1)(i+2)}\right)\right)\\
&\geq m\sum_{i=1}^{k-2}\left(\frac{8}{i+1}-\frac{20}{(i+1)^2}\right)\\
&\geq m\left(8H_k-\frac{10\pi^2}{3}\right).
\end{aligned}$$

Hence, for any $\epsilon > 0$, it is possible to choose k and m sufficiently large so as to have $Apx^1_{\emptyset}(\mathcal{G}^1_{ide}) \geq 4-\epsilon$. □

Thus, there is a small gap between the upper and lower bounds for $Apx^1_{\emptyset}(\mathcal{G}^1_{ide})$ which still needs to be closed. The conjecture is that the upper bound is not tight.

3.1 Weighted Players

For the case of weighted players, instead, we have an exact characterization obtained by combining the following two results.

Theorem 3 ([4]). $Apx^1_{\emptyset}(\mathcal{G}^1_{ide}) \leq 3+2\sqrt{2}$.

Proof. By the greedy inequality, for any $i \in [n]$, we have

$$\begin{aligned}
c_{t_i^*}(\mathbf{s}^i)^2-(c_{t_i^*}(\mathbf{s}^{i-1}))^2 &= (c_{t_i^*}(\mathbf{s}^{i-1})+w_i)^2-c_{t_i^*}(\mathbf{s}^{i-1})^2\\
&\leq (c_{s_i^*}(\mathbf{s}^{i-1})+w_i)^2-c_{s_i^*}(\mathbf{s}^{i-1})^2\\
&= 2w_ic_{s_i^*}(\mathbf{s}^i)+w_i^2 \leq 2w_in_{s_i^*}+w_i^2.
\end{aligned}$$

Let P_j and P_j^* be the set of players choosing resource r_j in the walk and in the optimal profile, respectively. We obtain

$$\sum_{j\in[m]}\sum_{i\in P_j}\left(c_j(\mathbf{s}^i)^2-c_j(\mathbf{s}^{i-1})^2\right) \leq \sum_{j\in[m]}\sum_{i\in P_j^*}\left(2w_in_j+w_i^2\right).$$

The right hand side of the above inequality telescopes for each j. Moreover, since $\sum_{i\in P_j^*} w_i = o_j$ implies $\sum_{i\in P_j^*} w_i^2 \leq o_j^2$, we have

$$\sum_{j\in[m]} n_j^2 \leq \sum_{j\in[m]}\left(n_jo_j+o_j^2\right) \leq 2\sqrt{\sum_{j\in[m]} n_j^2\sum_{j\in[m]} o_j^2}+\sum_{j\in[m]} o_j^2,$$

where the last inequality follows from the Cauchy-Schwartz inequality.

Let $x = \frac{\sum_{j\in[m]} n_j^2}{\sum_{j\in[m]} o_j^2}$ denote the approximation ratio achieved by $\mathbf{s}^n$. Dividing the last inequality by $\sum_{j\in[m]} o_j^2$, we obtain $x \leq 2\sqrt{x}+1$, which gives $x \leq 3+2\sqrt{2}$. □

Theorem 4 ([7]). *For any $\epsilon > 0$, it holds $Apx_\emptyset^1(\mathcal{G}_{ide}^1) \geq 3 + 2\sqrt{2} - \epsilon$.*

Proof. The lower bounding instance is constructed recursively as follows. Let $SG_0 \in \mathcal{G}_{ide}^1$ be a game with two resources and one player of weight 1 who can choose any of the two resources and let W_0 be the one-round walk in which the player chooses the second resource. For any $k \geq 1$, we define SG_k as the game obtained by merging two identical games SG_{k-1}^1 and SG_{k-1}^2 and by adding a player p_k of weight $2^{k/2}$ who can choose any of the two last resources in SG_{k-1}^1 and SG_{k-1}^2. The one-round walk W_k is obtained by merging the two one-round walks W_{k-1}^1 and W_{k-1}^2 and then by letting p_k be the last player to perform a best response in W_k. We assume that p_k chooses the last resource in SG_{k-1}^2. Clearly, this is a best response since, because of the fact that SG_{k-1}^1 and SG_{k-1}^2 are identical, both of the available resources for p_k have the same congestion.

Denote as $opt(k)$ and $sol(k)$ the social value of the optimal strategy profile in SG_k and the social value of the final profile of W_k, respectively. By the recursive construction and the fact that in an optimal profile for G_{k-1} the last resource has congestion zero, it follows that $opt(k) = 2opt(k-1) + 2^k$. On the other hand, the choice of p_k increases the congestion of the last resource in G_{k-1}^2 from $\sum_{i=0}^{k-1} 2^{i/2}$ to $\sum_{i=0}^{k} 2^{i/2}$. Thus, it follows that $sol(k) = 2sol(k-1) + 2^k + 2^{1+k/2}\sum_{i=0}^{k-1} 2^{i/2}$. Using these relationships, it can be shown by induction that $opt(k) = 2^k(k+1)$ and $sol(k) = 2^k\left((3+2\sqrt{2})k - 5 - 4\sqrt{2}\right) + 2^{1+k/2}(3+2\sqrt{2})$. Hence, for any $\epsilon > 0$, there exists a suitably large k such that $Apx_\emptyset^1(\mathcal{G}_{ide}^1) \geq 3 + 2\sqrt{2} - \epsilon$. □

4 Linear Congestion Games

For the class of games $\mathcal{G}_{lin}$, $Apx_\emptyset^1(\mathcal{G}_{lin})$ has been exactly characterized by the works of Christodoulou at al. [12] and Bilò et al. [6].

In order to show the upper bound, we need the following numerical lemma.

Lemma 4 ([12]). *For any pair of non-negative integers α and β, it holds $2\alpha\beta + 2\beta - \alpha \leq \frac{\alpha^2}{1+\phi} + (1+\phi)\beta^2$, where $\phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio.*

Theorem 5 ([12]). $Apx_\emptyset^1(\mathcal{G}_{lin}) \leq 2 + \sqrt{5}$.

Proof. For any $i \in [n]$, the greedy inequality in this setting is

$$\sum_{r\in t_i^*} a_r c_r(\mathbf{s}^{i-1}) \leq \sum_{r\in s_i^*}\left(a_r c_r(\mathbf{s}^{i-1}) + a_r + b_r\right) - \sum_{r\in t_i^*}(a_r + b_r).$$

By exploiting the greedy inequality, the social value of each intermediate profile produced during the walk can be bounded as follows.

$$\begin{aligned}
\textsc{Sum}(\mathbf{s}^i) &= \sum_{r\in R} c_r(\mathbf{s}^i)\ell_r(c_r(\mathbf{s}^i))\\
&= \sum_{r\in R\setminus t_i^*} c_r(\mathbf{s}^{i-1})\ell_r(c_r(\mathbf{s}^{i-1})) + \sum_{r\in t_i^*}(c_r(\mathbf{s}^{i-1})+1)\ell_r(c_r(\mathbf{s}^{i-1})+1)\\
&= \textsc{Sum}(\mathbf{s}^{i-1}) + \sum_{r\in t_i^*}(2a_rc_r(\mathbf{s}^{i-1})+a_r+b_r)\\
&\leq \textsc{Sum}(\mathbf{s}^{i-1}) + 2\sum_{r\in s_i^*}(a_rc_r(\mathbf{s}^{i-1})+a_r+b_r) - \sum_{r\in t_i^*}(a_r+b_r).
\end{aligned}$$

Summing up all these inequalities for any $i \in [n]$ and using Lemma 4, we obtain

$$\begin{aligned}
\textsc{Sum}(\mathbf{s}^n) &\leq \textsc{Sum}(\mathbf{s}^\emptyset) + 2\sum_{i\in[n]}\sum_{r\in s_i^*}(a_rc_r(\mathbf{s}^{i-1})+a_r+b_r) - \sum_{i\in[n]}\sum_{r\in t_i^*}(a_r+b_r)\\
&\leq 2\sum_{i\in[n]}\sum_{r\in s_i^*}(a_rc_r(\mathbf{s}^n)+a_r+b_r) - \sum_{i\in[n]}\sum_{r\in t_i^*}(a_r+b_r)\\
&= 2\sum_{r\in R} c_r(\mathbf{s}^*)(a_rc_r(\mathbf{s}^n)+a_r+b_r) - \sum_{r\in R} c_r(\mathbf{s}^n)(a_r+b_r)\\
&= \sum_{r\in R} a_r(2c_r(\mathbf{s}^*)c_r(\mathbf{s}^n)+2c_r(\mathbf{s}^*)-c_r(\mathbf{s}^n)) + \sum_{r\in R} b_r(2c_r(\mathbf{s}^*)-c_r(\mathbf{s}^n))\\
&\leq \sum_{r\in R} a_r\left(\frac{1}{1+\phi}c_r(\mathbf{s}^n)^2+(1+\phi)c_r(\mathbf{s}^*)^2\right) + \sum_{r\in R} b_r(2c_r(\mathbf{s}^*)-c_r(\mathbf{s}^n))\\
&\leq \frac{1}{1+\phi}\textsc{Sum}(\mathbf{s}^n) + (1+\phi)\textsc{Sum}(\mathbf{s}^*),
\end{aligned}$$

which gives $Apx_\emptyset^1(\mathcal{G}_{lin}) \leq 2+\sqrt{5}$. □

A matching lower bound can be achieved through the following construction.

Given three positive integers n, k and o, with $n \geq 2k+o-1$ and $k \geq 2o$, we define the game $CG_{n,k,o}$ in which there are n players, $m = n+1$ resources and each player $i \in [n]$ possesses exactly two strategies s_i and s_i' defined according to the following scheme.

- $s_i = \{r_i\}$ and $s_i' = \{r_{i+1}\} \cup \bigcup_{j=k+1}^{k+i}\{r_j\}$, for any $i \in [k]$;
- $s_i = \bigcup_{j=k+1}^{i}\{r_j\}$ and $s_i' = \bigcup_{j=i+1}^{k+i}\{r_j\}$, for any $k+1 \leq i \leq k+o$;

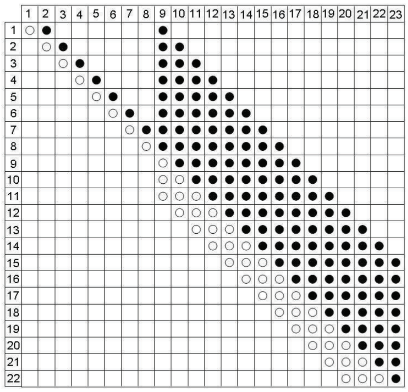

Figure 1: The set of strategies available to each player in the game $CG_{22,8,3}$. Rows are associated with players, while columns with resources. White and black circles represent the first and the second strategy, respectively.

- $s_i = \bigcup_{j=i-o+1}^{i} \{r_j\}$ and $s'_i = \bigcup_{j=i+1}^{\min\{k+i,m\}} \{r_j\}$, for any $k + o + 1 \leq i \leq n$.

A small example in which $n = 22$, $k = 8$ and $o = 3$ is shown in Figure 1.

For any $r \in R$, we set $b_r = 0$ and a_r is obtained as a solution of the following system of linear equations.

$$A = \begin{cases} eq_1 \\ eq_2 \\ \dots \\ eq_n \end{cases}$$

where each eq_i is defined as follows:

- $a_1 - a_2 - a_{k+1} = 0$,
- $2a_i - a_{i+1} - \sum_{j=k+1}^{k+i} \Big((k + i - j + 1)a_j\Big) = 0 \;\; \forall\, i = 2, \dots, k - 1$,
- $2a_k - \sum_{j=k+1}^{2k} \Big((2k - j + 1)a_j\Big) = 0$,
- $(k + 1) \sum_{j=k+1}^{i} a_j - \sum_{j=i+1}^{m} \Big((k + i - j + 1)a_j\Big) = 0 \;\; \forall\, i \in \{k + 1, \dots, k + o\}$,

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
1	1	-1							-1														
2		2	-1						-2	-1													
3			2	-1					-3	-2	-1												
4				2	-1				-4	-3	-2	-1											
5					2	-1			-5	-4	-3	-2	-1										
6						2	-1		-6	-5	-4	-3	-2	-1									
7							2	-1	-7	-6	-5	-4	-3	-2	-1								
8								2	-8	-7	-6	-5	-4	-3	-2	-1							
9									9	-8	-7	-6	-5	-4	-3	-2	-1						
10									9	9	-8	-7	-6	-5	-4	-3	-2	-1					
11									9	9	9	-8	-7	-6	-5	-4	-3	-2	-1				
12										9	9	9	-8	-7	-6	-5	-4	-3	-2	-1			
13											9	9	9	-8	-7	-6	-5	-4	-3	-2	-1		
14												9	9	9	-8	-7	-6	-5	-4	-3	-2	-1	
15													9	9	9	-8	-7	-6	-5	-4	-3	-2	-1
16														9	9	9	-8	-7	-6	-5	-4	-3	-2
17															9	9	9	-8	-7	-6	-5	-4	-3
18																9	9	9	-8	-7	-6	-5	-4
19																	9	9	9	-8	-7	-6	-5
20																		9	9	9	-8	-7	-6
21																			9	9	9	-8	-7
22																				9	9	9	-8

Figure 2: The coefficient matrix B generated by the game $CG_{22,8,3}$.

- $$(k+1)\sum_{j=i-o+1}^{i} a_j - \sum_{j=i+1}^{\min\{k+i,m\}} \Big((k+i-j+1)a_j\Big) = 0 \;\; \forall\, i \in \{k+o+1,\ldots,n\}.$$

Note that the definition of each equality is such that, for any $i \in [n]$, both strategies are equivalent for player i, provided all players $j < i$ have chosen s'_j and all players $j > i$ have not entered the game yet. Thus, we have that the strategy profile $\mathbf{s}^n$, in which all players choose the second of their strategies, is a possible outcome for a one-round walk starting from the empty strategy profile.

Let B be the $n{\times}m$ coefficient matrix defining system A. The matrix B generated by the game $CG_{22,8,3}$ is shown in Figure 2.

Let $a = (a_1,\ldots,a_m)^T$. In order for our instance to be well defined, we need to prove that there exists at least a strictly positive solution to the homogeneous system $Ba = 0$.

Lemma 5 **([6]).** *The system of linear equations $Ba = 0$ admits a strictly positive solution.*

For our purposes, we do not have to explicitly solve system A, but only need to prove some properties characterizing its set of solutions. We do this in the next two lemmas.

Lemma 6 **([6]).** *In any solution of system A it holds $a_1 \leq 4\sum_{j=k+1}^{2k} a_j$.*

Proof. It suffices to prove that for any $i \in \{2,\ldots,k+1\}$ it holds

$$a_i = 2^{i-2}a_1 - \sum_{j=k+1}^{k+i-1}\Big((2^{k+i+1-j} - k - i + j - 2)a_j\Big).$$

We use and inductive argument on i. For $i = 2$ we get $a_2 = a_1 - a_{k+1}$ which coincides with eq_1. Now, let us use the inductive hypothesis on a_{i-1}, that is, assume that

$$a_{i-1} = 2^{i-3}a_1 - \sum_{j=k+1}^{k+i-2} \left((2^{k+i-j} - k - i + j - 1)a_j\right).$$

By combining this equation with eq_{i-1} which is equal to

$$2a_{i-1} = a_i + \sum_{j=k+1}^{k+i-1} \left((k + i - j)a_j\right),$$

we obtain

$$\begin{aligned} a_i &= 2^{i-2}a_1 - \sum_{j=k+1}^{k+i-2} \left((2^{k+i-j+1} + 2(-k - i + j - 1))a_j\right) - \sum_{j=k+1}^{k+i-1} \left((k + i - j)a_j\right) \\ &= 2^{i-2}a_1 - \sum_{j=k+1}^{k+i-2} \left((2^{k+i-j+1} - k - i + j - 2)a_j\right) - a_{k+i-1} \\ &= 2^{i-2}a_1 - \sum_{j=k+1}^{k+i-1} \left((2^{k+i+1-j} - k - i + j - 2)a_j\right) \end{aligned}.$$

which completes the induction.
Setting $i = k + 1$, we have

$$a_{k+1} = 2^{k-1}a_1 - \sum_{j=k+1}^{2k} \left((2^{2k+2-j} - 2k - j - 3)a_j\right)$$

which gives

$$a_1 = \frac{\sum_{j=k+1}^{2k} \left((2^{2k+2-j} - 2k - j - 3)a_j\right) + a_{k+1}}{2^{k-1}}.$$

The claim follows directly from this last equation. □

Lemma 7 ([6]). *In any solution of system A it holds*

$$(k + 1) \sum_{i=m-o+1}^{m} ((i - m + o)a_i) \leq \frac{k^3}{n - 2k - o + 1} \sum_{i=k+1}^{m-o} a_i.$$

Proof. For any integer $0 \leq i \leq \lfloor \frac{n-2k-o+1}{k+o} \rfloor$, consider the sum of the last $i(k+o)+k+o-1$ equations of system A. Note that the last $k+o-1$ equations are the all and only ones involving a_j for any $j \in \{m-o+1,\ldots,m\}$. Since for each possible value of i we are always summing the last $k+o-1$ equations, we have that all a_j, with $j \in \{m-o+1,\ldots,m\}$, appear in the sum once with coefficient $-p$ for any $p \in [k]$ and $m-j$ times with coefficient $k+1$. Thus, for every $j \in \{m-o+1,\ldots,m\}$, a_j appears in the sum with coefficient $(k+1)(m-j-\frac{k}{2})$.

For $i=0$, the lowest j such that a_j is involved in the $k-o+1$ equations is $j=m-k-2o+2$. Clearly, each a_j appears in the sum with a coefficient not greater than $(k+1)o$. Thus, we get

$$(k+1)\sum_{j=m-o+1}^{m}\left(\left(j-m+\frac{k}{2}\right)a_j\right) \leq (k+1)o\sum_{j=m-k-2o+2}^{m-o} a_j.$$

For a generic $i>0$, the lowest j such that a_j is involved in the $i(k+o)+k-o+1$ equations is $j=m-(i+1)(k+o)-o+2$. Clearly, each a_j, with $j \in \{m-(i+1)(k+o)-o+2,\ldots,m-i(k+o)-o+1\}$, appears in the sum with a coefficient not greater than $(k+1)o$. All the a_j with $j \in \{m-i(k+o)-o+2,\ldots,m-o\}$ appear in the sum once with coefficient $-p$ for any $p \in [k]$ and o times with coefficient $k+1$, that is, with an overall negative coefficient since $\frac{k}{2} \geq o$. Thus, we get

$$(k+1)\sum_{j=m-o+1}^{m}\left(\left(j-m+\frac{k}{2}\right)a_j\right) \leq (k+1)o\sum_{j=m-(i+1)(k+o)-o+2}^{m-i(k+o)-o+1} a_j.$$

By summing for all possible indexes i, we obtain

$$\left(\left\lfloor\frac{n-2k-o+1}{k+o}\right\rfloor+1\right)(k+1)\sum_{j=m-o+1}^{m}\left(\left(j-m+\frac{k}{2}\right)a_j\right) \leq (k+1)o\sum_{j=k+1}^{m-o} a_j,$$

which gives

$$(k+1)\sum_{j=m-o+1}^{m}\left(\left(j-m+\frac{k}{2}\right)a_j\right) \leq \frac{(k+1)(k+o)o}{n-2k-o+1}\sum_{j=k+1}^{m-o} a_j,$$

which yields the claim since $\frac{k}{2} \geq o$ and $(k+1)(k+o)o \leq k^3$. □

The following lower bound, hence, can be achieved.

Theorem 6 **([6]).** *For any $\epsilon > 0$, it holds $Apx^1_{\emptyset}(\mathcal{G}_{lin}) \geq 2+\sqrt{5}-\epsilon$.*

Proof. For a fixed integer $n \gg 0$, set $k = \lfloor \sqrt[4]{n} \rfloor$ and $o = \lfloor \frac{3-\sqrt{5}}{2} k \rfloor$. Note that, for a sufficiently large n, these values are consistent with the definition of $CG_{n,k,o}$ since $n \geq 2k + o - 1$ and $k \geq 2o$.

Consider the sum of all the equations defining system A together with the dummy one $a_1 = a_1$. We obtain the equation

$$\sum_{i=1}^{k} 2a_i + (k+1)o \sum_{i=k+1}^{m} a_i - (k+1) \sum_{i=m-o+1}^{m} ((i-m+o)a_i) = \sum_{i=1}^{k} a_i + \frac{k(k+1)}{2} \sum_{i=k+1}^{m} a_i$$

which yields

$$\sum_{i=1}^{k} a_i = (k+1)\left(\frac{k}{2} - o\right) \sum_{i=k+1}^{m} a_i + (k+1) \sum_{i=m-o+1}^{m} ((i-m+o)a_i). \tag{1}$$

Let $\mathbf{s}^*$ be the strategy profile in which all players choose the first of their strategies. By comparing the social costs of $\mathbf{s}^n$ and $\mathbf{s}^*$, we obtain

$$\frac{\textsc{Sum}(\mathbf{s}^n)}{\textsc{Sum}(\mathbf{s}^*)} \geq \frac{\sum_{i=2}^{k} a_i + k^2 \sum_{i=k+1}^{m} a_i}{\sum_{i=1}^{k} a_i + o^2 \sum_{i=k+1}^{m} a_i},$$

where we have exploited the fact that $\textsc{Sum}(\mathbf{s}^*) \leq \sum_{i=1}^{k} a_i + o^2 \sum_{i=k+1}^{m} a_i$.

By using Equality 1, we get

$$\begin{aligned}
\frac{\textsc{Sum}(\mathbf{s}^n)}{\textsc{Sum}(\mathbf{s}^*)} &\geq \frac{\sum_{i=2}^{k} a_i + k^2 \sum_{i=k+1}^{m} a_i}{\sum_{i=1}^{k} a_i + o^2 \sum_{i=k+1}^{m} a_i} \\
&= \frac{\left((k+1)\left(\frac{k}{2} - o\right) + k^2\right) \sum_{i=k+1}^{m} a_i + (k+1) \sum_{i=m-o+1}^{m} ((i-m+o)a_i) - a_1}{\left((k+1)\left(\frac{k}{2} - o\right) + o^2\right) \sum_{i=k+1}^{m} a_i + (k+1) \sum_{i=m-o+1}^{m} ((i-m+o)a_i)} \\
&\geq \frac{\left((k+1)\left(\frac{k}{2} - o\right) + k^2 + \frac{k^3}{n-2k-o+1} - 4\right) \sum_{i=k+1}^{m} a_i}{\left((k+1)\left(\frac{k}{2} - o\right) + o^2 + \frac{k^3}{n-2k-o+1}\right) \sum_{i=k+1}^{m} a_i},
\end{aligned}$$

where, in the last inequality, we have used Lemmas 6 and 7 together with the fact that for any four positive numbers α, β, γ and δ such that $\alpha \geq \beta$ and $\gamma \geq \delta$, it holds $\frac{\alpha+\delta}{\beta+\delta} \geq \frac{\alpha+\gamma}{\beta+\gamma}$.

For n going to infinity, by considering only the dominant terms, we obtain

$$\begin{aligned}
\lim_{k\to\infty} \frac{\textsc{Sum}(\mathbf{s}^n)}{\textsc{Sum}(\mathbf{s}^*)} &\geq \lim_{k\to\infty} \frac{k\left(\frac{k}{2} - o\right) + k^2}{k\left(\frac{k}{2} - o\right) + o^2} \\
&= \lim_{k\to\infty} \frac{\frac{\sqrt{5}-2}{2}k^2 + k^2}{\frac{\sqrt{5}-2}{2}k^2 + \frac{7-3\sqrt{5}}{2}k^2} \\
&= \lim_{k\to\infty} \frac{\frac{\sqrt{5}}{2}k^2}{\frac{5-2\sqrt{5}}{2}k^2} \\
&= \frac{\sqrt{5}}{5-2\sqrt{5}} \\
&= 2 + \sqrt{5},
\end{aligned}$$

which implies the claim. □

4.1 Weighted Players

As to the extension to weighted players, Christodoulou at al. [12] provided the following upper bound.

Theorem 7 ([12]). *$Apx_{\emptyset}^{1}(\mathcal{G}_{lin}) \leq 4 + 2\sqrt{3}$.*

Proof. For any $i \in [n]$, the greedy inequality in this setting is

$$\sum_{r\in t_i^*} \left(a_r c_r(\mathbf{s}^{i-1}) + a_r w_i + b_r\right) \leq \sum_{r\in s_i^*} \left(a_r c_r(\mathbf{s}^{i-1}) + a_r w_i + b_r\right),$$

which, by multiplying both members by w_i, gives

$$\sum_{r\in t_i^*} \left(a_r c_r(\mathbf{s}^{i-1}) + a_r w_i + b_r\right) w_i \leq \sum_{r\in s_i^*} \left(a_r c_r(\mathbf{s}^{i-1}) + a_r w_i + b_r\right) w_i.$$

By exploiting the greedy inequality, the social value of each intermediate profile produced during the walk can be bounded as follows.

$$\begin{aligned}
\textsc{Sum}(\mathbf{s}^i) &= \sum_{r\in R} c_r(\mathbf{s}^i)\ell_r(c_r(\mathbf{s}^i)) \\
&= \sum_{r\in R\setminus t_i^*} c_r(\mathbf{s}^{i-1})\ell_r(c_r(\mathbf{s}^{i-1})) + \sum_{r\in t_i^*}(c_r(\mathbf{s}^{i-1})+w_i)\ell_r(c_r(\mathbf{s}^{i-1})+w_i) \\
&= \textsc{Sum}(\mathbf{s}^{i-1}) + \sum_{r\in t_i^*}(2a_rc_r(\mathbf{s}^{i-1})w_i + a_rw_i^2 + b_rw_i) \\
&\le \textsc{Sum}(\mathbf{s}^{i-1}) + 2\sum_{r\in s_i^*}(a_rc_r(\mathbf{s}^{i-1}) + a_rw_i + b_r)w_i.
\end{aligned}$$

Summing up all these inequalities for any $i \in [n]$ and using the Cauchy-Schwartz inequality, we obtain

$$\begin{aligned}
\textsc{Sum}(\mathbf{s}^n) &\le \textsc{Sum}(\mathbf{s}^\emptyset) + 2\sum_{i\in[n]}\sum_{r\in s_i^*}(a_rc_r(\mathbf{s}^{i-1}) + a_rw_i + b_r)w_i \\
&\le 2\sum_{i\in[n]}\sum_{r\in s_i^*}(a_rc_r(\mathbf{s}^n) + a_rw_i + b_r)w_i \\
&\le 2\sum_{r\in R} a_rc_r(\mathbf{s}^*)c_r(\mathbf{s}^n) + 2\sum_{r\in R}(a_rc_r(\mathbf{s}^*)^2 + b_rc_r(\mathbf{s}^*)) \\
&\le 2\sqrt{\sum_{r\in R} a_rc_r(\mathbf{s}^n)^2 \sum_{r\in R} a_rc_r(\mathbf{s}^*)^2} + 2\textsc{Sum}(\mathbf{s}^*) \\
&\le 2\sqrt{\textsc{Sum}(\mathbf{s}^n)\textsc{Sum}(\mathbf{s}^*)} + 2\textsc{Sum}(\mathbf{s}^*).
\end{aligned}$$

Dividing by $\textsc{Sum}(\mathbf{s}^*)$ and setting $x = \frac{\textsc{Sum}(\mathbf{s}^n)}{\textsc{Sum}(\mathbf{s}^*)}$, we obtain $x \le 2\sqrt{x} + 2$, which yields $x \le (1+\sqrt{3})^2$ and thus $Apx_\emptyset^1(\mathcal{G}_{lin}) \le 4 + 2\sqrt{3}$. □

For this case, no particular lower bounds are known.

5 Conclusions and Open Problems

Studying the performances of one-round walks has become a subject of great importance because of the intractability of the problems of computing either pure and mixed Nash equilibria in several games of interest. Moreover, the solution achieved after a one-round walk starting from the empty strategy profile can also be interpreted as the one produced by a competitive algorithm for an online problem in which requests arrive one at time and the algorithm has to perform an irrevocable decision on how to satisfy them. More precisely, the set of requests and the sets of the possible choices satisfying each request become, respectively, the set players and the sets of strategies in a game and the online algorithm will take its decisions by mimicking the best response choices of the players. However, it must be noted that best responses are based on the individual cost function experienced by a player in a given profile, while the quality of the produced solution is measured by means of the social function (which in this case will represent the objective function of the online problem) which is usually different from the individual cost functions of the players. This means that, in general, online algorithms defined in this way may significantly differ from online greedy algorithms in which, at each step, the choice yielding the smallest increase in the value of the objective function is performed. Anyway, there are cases in which the two algorithms coincide as for the case of singleton load balancing games. The result of Awerbuch et al. [4] presented in this note, in fact, comes from an early work on online load balancing published in 1995, that is, prior of the birth of Algorithmic Game Theory.

As to open problems and future research, besides closing the gaps in the existing results, there are lots of cases that still need to be exactly characterized. For example, the are no purpose-derived bounds for the class $\mathcal{G}^1_{lin}$ for either unweighted and weighted players. In fact, both lower bounds for $Apx^1_\emptyset(\mathcal{G}^1_{lin})$ come from the class $\mathcal{G}^1_{ide}$, while both its upper bounds come from $\mathcal{G}_{lin}$. Thus, it is not yet understood whether instances in this class are "easy" as those in $\mathcal{G}^1_{ide}$ or "hard" as those in $\mathcal{G}_{lin}$ or if they lie somehow in the between. Moreover, one can study the restricted case in which players are symmetric, that is, they all possess the same strategy set. Also a different social function may be analyzed: for the social function Max, some results are given in [5, 6]. Finally, does it make sense to study performances of one-round walks in which each player randomizes on her set of available best responses instead of deterministically choosing one of them?

References

[1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE Computer Society, pp. 613-622, 2006.

[2] E. Anshelevich, A. Dasgupta, J. M. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4), pp. 1602-1623, 2008.

[3] B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, pp. 57-66, 2005.

[4] B. Awerbuch, Y. Azar, E. F. Grove, M.-Y. Kao, P. Krishnan, and J. S. Vitter. Load balancing in the *Lp* norm. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE Computer Society, pp. 383-391, 1995.

[5] Y. Azar, J. Naor, and R. Rom. The Competitiveness of On-Line Assignments. In *Proceedings of 3rd Annual ACM-SIAM Symposium on Discrete Algorithms(SODA)*, SIAM, pp. 203-210, 1992.

[6] V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. Performances of One-Round Walks in Linear Congestion Games. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT)*, LNCS 5814, Springer, pp. 311-322, 2009.

[7] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight Bounds for Selfish and Greedy Load Balancing. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS 4051, Springer, pp. 311-322, 2006.

[8] X. Chen and X. Deng. Settling the Complexity of 2-Player Nash Equilibrium. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE Computer Society, pp. 261-252, 2006.

[9] G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, pp. 67-73, 2005.

[10] G. Christodoulou and E. Koutsoupias. On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, LNCS 3669, Springer, pp. 59-70, 2005.

[11] G. Christodoulou, E. Koutsoupias, and P. G. Spirakis. On the Performance of Approximate Equilibria in Congestion Games. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, LNCS 5757, Springer, pp. 251-262, 2009.

[12] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS)*, LNCS 3884, Springer, pp. 349-360, 2006.

[13] K. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, pp. 71-78, 2006.

[14] K. Daskalakis and C. H. Papadimitriou. Three-Player Games Are Hard. *Electronic Colloquium on Computational Complexity (ECCC)*, (139), 2005.

[15] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS)*, LNCS 1563, Springer, pp. 404-413, 1999.

[16] V. S. Mirrokni and A. Vetta. Convergence Issues in Competitive Games. In *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, LNCS 3122, Springer, pp. 183-194, 2004.

[17] T. Roughgarden and E. Tardos. How Bad is Selfish Routing? *Journal of ACM*, 49:236-259, 2002.

[18] S. Suri, C. Tóth and Y. Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1): 79-96, 2007.