On the Impact of Fair Best Response Dynamics*

Angelo Fanelli[1], Luca Moscardelli[2], and Alexander Skopalik[3]

[1] Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore
angelo.fanelli@ntu.edu.sg

[2] Department of Economic Studies, University of Chieti-Pescara, Italy
luca.moscardelli@unich.it

[3] Department of Computer Science, TU Dortmund, Germany
alexander.skopalik@tu-dortmund.de

Abstract. In this work we completely characterize how the frequency with which each player participates in the game dynamics affects the possibility of reaching efficient states, i.e., states with an approximation ratio within a constant factor from the price of anarchy, within a polynomially bounded number of best responses. We focus on the well known class of linear congestion games and we show that (i) if each player is allowed to play at least once and at most β times in T best responses, states with approximation ratio $O(\beta)$ times the price of anarchy are reached after $T\lceil \log\log n \rceil$ best responses, and that (ii) such a bound is essentially tight also after exponentially many ones. One important consequence of our result is that the fairness among players is a necessary and sufficient condition for guaranteeing a fast convergence to efficient states. This answers the important question of the maximum order of β needed to fast obtain efficient states, left open by [10,11] and [3], in which fast convergence for constant β and very slow convergence for $\beta = O(n)$ have been shown, respectively. Finally, we show that the structure of the game implicitly affects its performances. In particular, we prove that in the symmetric setting, in which all players share the same set of strategies, the game always converges to an efficient state after a polynomial number of best responses, regardless of the frequency each player moves with. All the results extend to weighted congestion games.

Keywords: Congestion Games, Speed of Convergence, Best Response Dynamics.

1 Introduction

Congestion games are used for modelling non-cooperative systems in which a set of resources are shared among a set of selfish players. In a congestion game we

* This research was partially supported by the grant NRF-RF2009-08 "Algorithmic aspects of coalitional games", by the PRIN 2008 research project COGENT (COmputational and GamE-theoretic aspects of uncoordinated NeTworks), funded by the Italian Ministry of University and Research, and by Deutsche Forschungsgemeinschaft (DFG) within the Collaborative Research Center SFB 876 "Providing Information by Resource-Constrained Analysis", project A2.

B. Rovan, V. Sassone, and P. Widmayer (Eds.): MFCS 2012, LNCS 7464, pp. 360–371, 2012.

have a set of m resources and a set of n players. Each player's strategy consists of a subset of resources. The delay of a particular resource e depends on its congestion, corresponding to the number of players choosing e, and the cost of each player i is the sum of the delays associated with the resources selected by i. In this work we focus on linear congestion games where the delays are linear functions. A congestion game is called symmetric if all players share the same strategy set. A state of the game is any combination of strategies for the players and its social cost, defined as the sum of the players' costs, denotes its quality from a global perspective. The social optimum denotes the minimum possible social cost among all the states of the game.

Related Work. Rosenthal [14] has shown, by a potential function argument, that the natural decentralized mechanism known as Nash dynamics consisting in a sequence of moves in which at each one some player switches its strategy to a better alternative, is guaranteed to converge to a pure Nash equilibrium [13].

In order to measure the degradation of social welfare due to the selfish behavior of the players, Koutsoupias and Papadimitriou [12] defined the price of anarchy as the worst-case ratio between the social cost in a Nash equilibrium and that of a social optimum. The price of anarchy for congestion games has been investigated by Awerbuch et al. [2] and Christodoulou and Koutsoupias [6]. They both proved that the price of anarchy for congestion games with linear delays is 5/2.

The existence of a potential function relates the class of congestion games to the class of polynomial local search problems (PLS) [8]. Fabrikant et al. [9] proved that, even for symmetric congestion games, the problem of computing Nash equilibria is PLS-complete [8]. One major consequence of the completeness result is the existence of congestion games with initial states such that any improvement sequence starting from these states needs an exponential number of steps in the number of players n in order to reach a Nash equilibrium. A recent result by Ackermann et al. [1] shows that the previous negative result holds even in the restricted case of congestion games with linear delay functions.

The negative results on computing equilibria in congestion games have lead to the development of the concept of ϵ-Nash equilibrium, in which no player can decrease its cost by a factor of more than ϵ. Unfortunately, as shown by Skopalik and Vöcking [15], also the problem of finding an ϵ-Nash equilibrium in congestion games is PLS-complete for any ϵ, though, under some restrictions on the delay functions, Chien and Sinclair [5] proved that in symmetric congestion games the convergence to ϵ-Nash equilibrium is polynomial in the description of the game and the minimal number of steps within which each player has a chance to move.

Since negative results tend to dominate the issues relative to the complexity of computing equilibria, another natural arising question is whether efficient states (with a social cost comparable to the one of any Nash equilibrium) can be reached by best response moves in a reasonable amount of time (e.g., [3,7,10,11]). We measure the efficiency of a state by the ratio among its cost and the optimal one, and we refer to it as the approximation ratio of the state. We generally say

that a state is efficient when its approximation ratio is within a constant factor from the price of anarchy. Since the price of anarchy of linear congestion games is known to be constant [2,6], efficient states approximate the social optimum by a constant factor. While Bilò et al. [4] considered such a problem restricted to the case in which the dynamics start from an empty state, proving that in such a setting an efficient state can be reached by allowing each player to move exactly once, we focus on the more general setting in which the dynamics start from a generic state. It is worth noticing that in the worst case, a generic Nash dynamics starting from an arbitrary state could never reach a state with an approximation ratio lower than the price of anarchy. Furthermore, by a potential function argument it is easy to show that in a linear congestion game, once a state S with a social cost $C(S)$ is reached, even if such a state is not a Nash equilibrium, we are guaranteed that for any subsequent state S' of the dynamics, $C(S') = O(C(S))$.
Awerbuch et al. [3] have proved that for linear congestion games, sequences of moves reducing the cost of each player by at least a factor of ϵ, converge to efficient states in a number of moves polynomial in $1/\epsilon$ and the number of players, under the minimal liveness condition that every player moves at least once every polynomial number of moves. Under the same liveness condition, they also proved that exact best response dynamics can guarantee the convergence to efficient states only after an exponential number of best responses [3]. Nevertheless, Fanelli et al. [10] have shown that, under more restrictive condition that each player plays exactly once every n best responses, any best response dynamics converges to an efficient state after $\Theta(n \log\log n)$ best responses. Subsequently, Fanelli and Moscardelli [11] extended the previous results to the more general case in which each player plays a constant number of times every $O(n)$ best responses.

Our Contribution. In this work we completely characterize how the frequency with which each player participates in the game dynamics affects the possibility of reaching efficient states. In particular, we close the most important open problem left open by [3] and [10,11] for linear congestion games. On the one hand, in [3] it is shown that, even after an exponential number of best responses, states with a very high approximation ratio, namely $\Omega\left(\frac{\sqrt{n}}{\log n}\right)$, can be reached. On the other hand, in [10,11] it is shown that, under the minimal liveness condition in which every player moves at least once every T steps, if players perform best responses such that each player is allowed to play at most $\beta = O(1)$ times any T steps (notice that $\beta = O(1)$ implies $T = O(n)$), after $T\lceil \log\log n \rceil$ best responses a state with a constant factor approximation ratio is reached.

The more β increases, the less the dynamics is fair with respect to the chance every player has of performing a best response: β measures the degree of unfairness of the dynamics. The important left open question was that of determining the maximum order of β needed to obtain fast convergence to efficient states: We answer this question by proving that, after $T\lceil \log\log n \rceil$ best responses, the dynamics reaches states with an approximation ratio of $O(\beta)$. Such a result is essentially tight since we are also able to show that, for any $\epsilon > 0$, there exist congestion games for which, even for an exponential number of best responses, states with an approximation ratio of $\Omega(\beta^{1-\epsilon})$ are obtained. Therefore, β constant as assumed in [10,11] is not only sufficient, but also necessary in order to reach efficient states after a polynomial number of best responses.

Finally, in the special case of symmetric congestion games, we show that the unfairness in best response dynamics does not affect the fast convergence to efficient states; namely, we prove that, for any β, after $T\lceil \log\log n \rceil$ best responses efficient states are always reached.

The paper is organized as follows: In the next section we provide the basic notation and definitions. Section 3 is devoted to the study of generic linear congestion games, while Section 4 analyzes the symmetric case. Finally, Section 5 provides some extensions of the results and gives some conclusive remarks.

2 Model and Definitions

A *congestion game* $\mathcal{G} = (N, E, (\Sigma_i)_{i\in N}, (f_e)_{e\in E}, (c_i)_{i\in N})$ is a non-cooperative strategic game characterized by the existence of a set E of resources to be shared by n players in $N = \{1, \ldots, n\}$.

Any strategy $s_i \in \Sigma_i$ of player $i \in N$ is a subset of resources, i.e., $\Sigma_i \subseteq 2^E$. A congestion game is called *symmetric* if all players share the same set of strategies, i.e., $\Sigma = \Sigma_i$ for every $i \in N$. Given a state or strategy profile $S = (s_1, \ldots, s_n)$ and a resource e, the number of players using e in S, called the congestion on e, is denoted by $n_e(S) = |\{i \in N \mid e \in s_i\}|$. A delay function $f_e : \mathbb{N} \mapsto \mathbb{Q}^+$ associates to resource e a delay depending on the congestion on e, so that the cost of player i for the pure strategy s_i is given by the sum of the delays associated with the resources in s_i, i.e., $c_i(S) = \sum_{e\in s_i} f_e(n_e(S))$.

In this paper we will focus on linear congestion games, that is having linear delay functions with nonnegative coefficients. More precisely, for every resource $e \in E$, $f_e(x) = a_e x + b_e$ for every resource $e \in E$, with $a_e, b_e \in \mathbb{Q}^+$.

Given the strategy profile $S = (s_1, \ldots, s_n)$, the social cost $C(S)$ of a given state S is defined as the sum of all the players' costs, i.e., $C(S) = \sum_{i\in N} c_i(S)$. An optimal strategy profile $S^* = (s_1^*, \ldots, s_n^*)$ is one having minimum social cost; we denote $C(S^*)$ by $\textsc{Opt}$. The *approximation ratio* of state S is given by the ratio between the social cost of S and the social optimum, i.e., $\frac{C(S)}{\textsc{Opt}}$. Moreover, given the strategy profile $S = (s_1, s_2, \ldots, s_n)$ and a strategy $s_i' \in \Sigma_i$, let $(S_{-i}, s_i') = (s_1, s_2, \ldots, s_{i-1}, s_i', s_{i+1}, \ldots, s_n)$, i.e., the strategy profile obtained from S if player i changes its strategy from s_i to s_i'.

The potential function is defined as $\Phi(S) = \sum_{e\in E} \sum_{j=1}^{n_e(S)} f_e(j)$. It is an *exact* potential function since it satisfies the property that for each player i and each strategy $s_i' \in \Sigma_i$ of i in S, it holds that $c_i(S_{-i}, s_i') - c_i(S) = \Phi(S_{-i}, s_i') - \Phi(S)$. It is worth noticing that in linear congestion games, for any state S, it holds $\Phi(S) \leq C(S) \leq 2\Phi(S)$.

Each player acts selfishly and aims at choosing the strategy decreasing its cost, given the strategy choices of other players. A *best response* of player i in S is a strategy $s_i^b \in \Sigma_i$ yielding the minimum possible cost, given the strategy choices of the other players, i.e., $c_i(S_{-i}, s_i^b) \leq c_i(S_{-i}, s_i')$ for any other strategy $s_i' \in \Sigma_i$. Moreover, if no $s_i' \in \Sigma_i$ is such that $c_i(S_{-i}, s_i') < c_i(S)$, the best response of i in S is s_i. We call a *best response dynamics* any sequence of best responses.

Given a best response dynamics starting from an arbitrary state, we are interested in the social cost of its final state. To this aim, we must consider dynamics in which each player performs a best response at least once in a given number T of best responses, otherwise one or more players could be "locked out" for arbitrarily long and we could not expect to bound the social cost of the state reached at the end of the dynamics. Therefore, we define a T-*covering* as a dynamics of T consecutive best responses in which each player moves at least once. More precisely, a T-covering $R = \left(S_R^0, \ldots, S_R^T\right)$ is composed of T best responses; S_R^0 is said to be the *initial* state of R and S_R^T is its *final* state. For every $1 \leq t \leq T$, let $\pi_R(t)$ be the player performing the t-th best response of R; π_R is such that every player performs at least a best response in R. In particular, for every $1 \leq t \leq T$, $S_R^t = \left((S_R^{t-1})_{-\pi_R(t)}, s'_{\pi_R(t)}\right)$ and $s'_{\pi_R(t)}$ is a best response of player $\pi_R(t)$ to S_R^{t-1}. For any $i = 1, \ldots, n$, the last best response performed by player i in R is the $\mathrm{last}_R(i)$-th best response of R, leading from state $S^{\mathrm{last}_R(i)-1}$ to state $S^{\mathrm{last}_R(i)}$. When clear from the context, we will drop the index R from the notation, writing S^i, π and $\mathrm{last}(i)$ instead of S_R^i, π_R and $\mathrm{last}_R(i)$, respectively.

Definition 1 (T-Minimum Liveness Condition). *Given any $T \geq n$, a best response dynamics satisfies the T-Minimum Liveness Condition if it can be decomposed into a sequence of consecutive T-coverings.*

In Section 3.2 we show that (for the general asymmetric case) under such a condition the quality of the reached state can be very bad even considering $T = O(n)$ (see Corollary 1): It is worth noticing that in the considered congestion game, only $\sqrt[4]{n}$ players perform a lot of best responses ($\sqrt{n}$ best responses) in each covering, while the remaining $n - \sqrt[4]{n}$ players perform only one best response every T-covering. The idea here is that there is a sort of unfairness in the dynamics, given by the fact that the players do not have the same chances of performing best responses.

In order to quantify the impact of fairness on best response dynamics, we need an additional parameter β and we define a β-bounded T-covering as a T-covering in which every player performs at most β best responses.

Definition 2 ((T, β)-Fairness Condition). *Given any positive integers β and T such that $n \leq T \leq \beta \cdot n$, a dynamics satisfies the (T, β)-Fairness Condition if it can be decomposed into a sequence of consecutive β-bounded T-coverings.*

Notice that β is a sort of (un)fairness index: If β is constant, it means that every player plays at most a constant number of times in each T-covering and therefore the dynamics can be considered *fair*.

In order to prove our upper bound results, we will focus our attention on particular congestion games to which any linear congestion game is best-response reducible. The following definition formally states such a notion of reduction.

Definition 3 (Best-Response Reduction). *A congestion game $\mathcal{G}$ is Best-Response reducible to a congestion game $\mathcal{G}'$ with the same set of players if there exists an injective function g mapping any strategy profile S of $\mathcal{G}$ to a strategy profile $g(S)$ of $\mathcal{G}'$ such that*

(i) for any $i = 1, \dots, n$ the cost of player i in S is equal to the one of player i in $g(S)$
(ii) for any $i = 1, \dots, n$, there exists, in the game $\mathcal{G}$, a best response of player i in S leading to state S' if and only if there exists, in the game $\mathcal{G}'$, a best response of player i in $g(S)$ leading to state $g(S')$.

3 Asymmetric Congestion Games

In this section we first (in Subsection 3.1) provide an upper bound to the approximation ratio of the states reached after a dynamics satisfying the (T, β)-Minimum Liveness Condition, starting from an arbitrary state and composed by a number of best responses polynomial in n. Finally (in Subsection 3.2), we provide an almost matching lower bound holding for dynamics satisfying the same conditions.

3.1 Upper Bound

All the results hold for linear congestion games having delay functions $f_e(x) = a_e x + b_e$ with $a_e, b_e \geq 0$ for every $e \in E$. Since our bounds are given as a function of the number of players, as shown in [10], the following proposition allows us to focus on congestion games with identical delay functions $f(x) = x$.

Proposition 1 ([10]). *Any linear congestion game is best-response reducible to a congestion game having the same set of players and identical delay functions $f(x) = x$.*

Since the dynamics satisfies the (T, β)-Fairness Condition, we can decompose it into k β-bounded T-coverings $R_1, \dots, R_k$.

Consider a generic β-bounded T-covering $R = \left(S^0, \dots, S^T\right)$. In the following we will often consider the *immediate* costs (or delays) of players during R, that is the cost $c_{\pi(t)}(S^t)$ right after the best response of player $\pi(t)$, for $t = 1, \dots, T$.

Given an optimal strategy profile S^*, since the t-th player $\pi(t)$ performing a best response, before doing it, can always select the strategy she would use in S^*, her immediate cost can be suitably upper bounded as $\sum_{e \in s^*_{\pi(t)}} \left(n_e(S^{t-1}) + 1\right)$.

By extending and strengthening the technique of [10,11], we are able to prove that the best response dynamics satisfying the (T, β)-Fairness Condition fast converges to states approximating the social optimum by a factor $O(\beta)$. It is worth noticing that, by exploiting the technique of [10,11], only a much worse bound of $O(\beta^2)$ could be proved. In order to obtain an $O(\beta)$ bound, we need to develop a different and more involved technique, in which also the functions ρ and H, introduced in [10,11], have to be redefined: roughly speaking, they now must take into account only the last move in R of each player, whereas in [10,11] they were accounting for all the moves in R.

We now introduce functions ρ and H, defined over the set of all the possible β-bounded T-coverings:

- Let $\rho(R) = \sum_{i=1}^{n} \sum_{e \in s_i^*} \left(n_e(S^{\text{last}_R(i)-1}) + 1\right)$;
- let $H(R) = \sum_{i=1}^{n} \sum_{e \in s_i^*} n_e(S^0)$.

Notice that $\rho(R)$ is an upper bound to the sum over all the players of the cost that she would experience on her optimal strategy s_i^* just before her last move in R, whereas $H(R)$ represents the sum over all the players of the delay on the moving player's optimal strategy s_i^* in the initial state S^0 of R. Moreover, since players perform best responses, $\sum_{i=1}^{n} c_i(S^{\text{last}_R(i)}) \leq \rho(R)$, i.e. $\rho(R)$ is an upper bound to the sum of the immediate costs over the last moves of every players.

The upper bound proof is structured as follows. Lemma 1 relates the social cost of the final state S^T of a β-bounded T-covering R with $\rho(R)$, by showing that $C(S^T) \leq 2\rho(R)$. Let $\overline{R}$ and R be two consecutive β-bounded T-coverings; by exploiting Lemmata 2 and 3, providing an upper (lower, respectively) bound to $H(R)$ in terms of $\rho(\overline{R})$ ($\rho(R)$, respectively), Lemma 4 proves that $\frac{\rho}{\text{OPT}}$ rapidly decreases between $\overline{R}$ and R, showing that $\frac{\rho(R)}{\text{OPT}} = O\left(\sqrt{\frac{\rho(\overline{R})}{\text{OPT}}}\right)$. In the proof of Theorem 1, after deriving a trivial upper bound equal to $O(n)$ for $\rho(R_1)$, Lemma 4 is applied to all the $k-1$ couples of consecutive β-bounded T-coverings of the considered dynamics satisfying the (T, β)-Fairness Condition.

The following lemmata show that the social cost at the end of any β-bounded T-covering R is at most $2\rho(R)$, and that $\frac{\rho(\cdot)}{\text{OPT}}$ fast decreases between two consecutive β-bounded T-coverings. They can be proved by adapting some proofs in [10,11] so that they still hold with the new definition of ρ.

Lemma 1. *For any $\beta \geq 1$, given a β-bounded T-covering R, $C(S^T) \leq 2\rho(R)$.*

Lemma 2. *For any $\beta \geq 1$, given a β-bounded T-covering R ending in S^T, $\frac{\sum_{e \in E} n_e(S^T) n_e(S^*)}{\text{OPT}} \leq \sqrt{2\frac{\rho(R)}{\text{OPT}}}$.*

In Lemma 3 we are able to relate $\rho(R)$ and $H(R)$ by much strengthening the technique exploited in [10,11].

Lemma 3. *For any $\beta \geq 1$, given a β-bounded T-covering R, $\frac{\rho(R)}{\text{OPT}} \leq 2\frac{H(R)}{\text{OPT}} + 4\beta + 1$.*

Proof. Let $\bar{N}$ be the set of players changing their strategies by performing best responses in R. First of all, notice that if the players in $\bar{N}$ never select strategies used by some player in S^*, i.e. if they select only resources e such that $n_e(S^*) = 0$, then, by recalling the definitions of $\rho(R)$ and $H(R)$, $\rho(R) \leq H(R) + \text{OPT}$ and the claim would easily follow for any $\beta \geq 1$.

In the following our aim is that of dealing with the generic case in which players moving in R can increase the congestions on resources e such that $n_e(S^*) > 0$.

For every resource $e \in E$, we focus on the congestion on such a resource above a "virtual" congestion frontier $g_e = 2\beta n_e(S^*)$.

We assume that at the beginning of covering R each resource $e \in E$ has a congestion equal to $\delta_{0,e} = \max\{n_e(S^0), g_e\}$, and we call $\delta_{0,e}$ the *congestion of level* 0 on resource e; moreover, $\Delta_0 = \sum_{e \in E} \delta_{0,e} \cdot n_e(S^*)$ is an upper bound to $H(R)$. We refer to Δ_0 as the total congestion of level 0.

The idea is that the total congestion of level 0 can induce on the resources a congestion (over the frontier g_e) being the total congestion of level 1, such a congestion a total congestion of level 2, and so on. More formally, for any $p \geq 1$ and any $e \in E$, we define $\delta_{p,e}$ as the congestion of level p on resource e above the frontier g_e; we say that a congestion $\delta_{p,e}$ of level p on resource e is *induced* by an amount $x_{p,e}$ of congestion of level $p-1$ if some players (say, players in $N_{p,e}$) moving on e can cause such a congestion of level p on e because they are experimenting a delay on the resources of their optimal strategies due to an amount $x_{p,e}$ of congestion of level p. Notice that, for each move of the players in $N_{p,e}$, such an amount $x_{p,e}$ of congestion of level $p-1$ can be used only once, i.e. it cannot be used in order to induce a congestion of level p for other resources in $E \setminus \{e\}$. In other words, $x_{p,e}$ is the overall congestion of level $p-1$ on the resources in the optimal strategies of players in $N_{p,e}$ used in order to induce the congestion $\delta_{p,e}$ of level p on resource e.

For any p, the total congestion of level p is defined as $\Delta_p = \sum_{e \in E} \delta_{p,e} \cdot n_e(S^*)$. Moreover, for any $p \geq 1$, we have that $\sum_{e \in E} x_{p,e} \leq \beta \Delta_{p-1}$ because each player can move at most β times in R and therefore the total congestion of level $p-1$ can be used at most β times in order to induce the total congestion of level p.

It is worth noticing that $\rho(R) \leq \sum_{p=0}^{\infty} \Delta_p + \textsc{Opt}$, because $\sum_{p=0}^{\infty} \delta_{p,e}$ is an upper bound on the congestion of resource e during the whole covering R:

$$\begin{aligned}
\rho(R) &= \sum_{i=1}^{n} \sum_{e \in s_i^*} \left(n_e(S^{\text{last}(i)-1}) + 1 \right) \\
&\leq \sum_{i=1}^{n} \sum_{e \in s_i^*} \left(\sum_{p=0}^{\infty} \delta_{p,e} + 1 \right) = \sum_{e \in E} \left(n_e(S^*) \left(\sum_{p=0}^{\infty} \delta_{p,e} + 1 \right) \right) \\
&= \sum_{e \in E} \sum_{p=0}^{\infty} \delta_{p,e} n_e(S^*) + \sum_{e \in E} n_e(S^*) = \sum_{p=0}^{\infty} \Delta_p + \textsc{Opt}
\end{aligned}$$

In the following, we bound $\sum_{p=0}^{\infty} \Delta_p$ from above.

$$\Delta_p = \sum_{e \in E} \delta_{p,e} \cdot n_e(S^*) \leq \sum_{e \in E} \frac{x_{p,e}}{g_e} \cdot n_e(S^*) \leq \sum_{e \in E} \frac{x_{p,e}}{2\beta n_e(S^*)} \cdot n_e(S^*) \leq \frac{\Delta_{p-1}}{2},$$

where the first inequality holds because $\delta_{p,e}$ is the portion of congestion on resource e above the frontier g_e due to some moving players having on the resources of their optimal strategy a delay equal to $x_{p,e}$, and the last inequality holds because each player can move at most β times in R and therefore the total congestion of level $p-1$ can be used at most β times in order to induce the total congestion of level p.

We thus obtain that, for any $p \geq 0$, $\Delta_p \leq \frac{\Delta_0}{2^p}$ and $\sum_{p=0}^{\infty} \Delta_p \leq 2\Delta_0$.

Since $\Delta_0 = \sum_{e \in E} \max\{n_e(S^0), 2\beta n_e(S^*)\} \cdot n_e(S^*) \leq H(R) + 2\beta\text{OPT}$ and $\rho(R) \leq \sum_{p=0}^{\infty} \Delta_p + \text{OPT} \leq 2\Delta_0 + \text{OPT}$, we finally obtain the claim. □

By combining Lemmata 2 and 3, it is possible to prove the following lemma showing that $\frac{\rho(R)}{\text{OPT}}$ fast decreases between two consecutive coverings.

Lemma 4. *For any $\beta \geq 1$, given two consecutive β-bounded T-coverings $\overline{R}$ and R,* $\frac{\rho(R)}{\text{OPT}} \leq 2\sqrt{2\frac{\rho(\overline{R})}{\text{OPT}}} + 4\beta + 1$.

By applying Lemma 4 to all the couples of consecutive β-bounded T-coverings, we are now able to prove the following theorem.

Theorem 1. *Given a linear congestion game, any best response dynamics satisfying the (T, β)-Fairness Condition converges from any initial state to a state S such that $\frac{C(S)}{\text{OPT}} = O(\beta)$ in at most $T\lceil \log \log n \rceil$ best responses.*

3.2 Lower Bound

Theorem 2. *For any $\epsilon > 0$, there exist a linear congestion game $\mathcal{G}$ and an initial state S^0 such that, for any $\beta = O(n^{-\frac{1}{\log_2 \epsilon}})$, there exists a best response dynamics starting from S^0 and satisfying the (T, β)-Fairness Condition such that for a number of best responses exponential in n the cost of the reached states is always $\Omega(\beta^{1-\epsilon} \cdot \text{OPT})$.*

By choosing $\beta = \sqrt{n}$ and considering a simplified version of the proof giving the above lower bound, it is possible to prove the following corollary. In particular, it shows that even in the case of best response dynamics verifying an $O(n)$-Minimum Liveness Condition, the speed of convergence to efficient states is very slow; such a fact implies that the T-Minimum Liveness condition cannot precisely characterize the speed of convergence to efficient states because it does not capture the notion of fairness in best response dynamics.

Corollary 1. *There exist a linear congestion game $\mathcal{G}$, an initial state S^0 and a best response dynamics starting from S^0 and satisfying the $O(n)$-Minimum Liveness Condition such that for a number of best responses exponential in n the cost of the reached states is always* $\Omega\left(\frac{\sqrt[4]{n}}{\log n} \cdot \text{OPT}\right)$.

4 Symmetric Congestion Games

In this section we show that in the symmetric case the unfairness in best response dynamics does not affect the speed of convergence to efficient states. In particular, we are able to show that, for any β, after $T\lceil \log \log n \rceil$ best responses an efficient state is always reached. To this aim, in the following we consider best response dynamics satisfying only the T-Minimum Liveness Condition, i.e. best response dynamics decomposable into k T-coverings $R_1, \ldots, R_k$.

All the results hold for linear congestion games having delay functions $f_e(x) = a_e x + b_e$ with $a_e, b_e \geq 0$ for every $e \in E$. Analogously to the asymmetric case, since our bounds are given as a function of the number of players, the following proposition allows us to focus on congestion games with identical delay functions $f(x) = x$.

Proposition 2. *Any symmetric linear congestion game is best-response reducible to a symmetric congestion game having the same set of players and identical delay functions $f(x) = x$.*

Consider a generic T-covering $R = \left(S^0, \ldots, S^T\right)$. Given an optimal strategy profile S^*, since player i, before performing her last best response, can always select any strategy s_j^*, for $j = 1 \ldots n$, of S^*, her immediate cost $c_i(S^{last(i)})$ can be upper bounded as $\frac{1}{n}\sum_{j=1}^{n}\sum_{e \in s_j^*}(n_e(S^{last(i)-1})+1) = \frac{1}{n}\sum_{e\in E} n_e(S^*)(n_e(S^{last(i)-1})+1)$. In order to prove our upper bound result, we introduce the following function:

- $\Gamma(R) = \frac{1}{n}\sum_{i=1}^{n}\sum_{e\in E} n_e(S^*)(n_e(S^{\mathrm{last}(i)-1})+1)$.

Notice that $\Gamma(R)$ is an upper bound to the sum of the immediate cost over the last moves of every players, i.e., $\Gamma(R) \geq \sum_{i=1}^{n} c_i(S^{\mathrm{last}_R(i)})$. Therefore, by exploiting the same arguments used in the proof of Lemma 1, it is possible to prove the following lemma relating the social cost $C(S^T)$ at the end of R with $\Gamma(R)$.

Lemma 5. *Given any T-covering R, $C(S^T) \leq 2\Gamma(R)$.*

Moreover, given any T-covering R, we can relate the social cost $C(S^T)$ of the final state of R with the cost $C(S^0)$ of its initial state.

Lemma 6. *Given any T-covering R, $\frac{C(S^T)}{\mathrm{OPT}} \leq (2+2\sqrt{2})\sqrt{\frac{C(S^0)}{\mathrm{OPT}}}$.*

Proof.

$$\frac{C(S^T)}{\mathrm{OPT}} \leq \frac{2\Gamma(R)}{\mathrm{OPT}} \tag{1}$$

$$= \frac{2}{n\mathrm{OPT}}\sum_{i=1}^{n}\sum_{e\in E} n_e(S^*)(n_e(S^{\mathrm{last}(i)-1})+1)$$

$$= \frac{2}{n\mathrm{OPT}}\left(\sum_{i=1}^{n}\sum_{e\in E} n_e(S^*)n_e(S^{\mathrm{last}(i)-1}) + \sum_{i=1}^{n}\sum_{e\in E} n_e(S^*)\right)$$

$$\leq 2 + \frac{2}{n\mathrm{OPT}}\sum_{i=1}^{n}\sum_{e\in E} n_e(S^*)n_e(S^{\mathrm{last}(i)-1})$$

$$\leq 2 + \frac{2}{n\mathrm{OPT}}\sum_{i=1}^{n}\sqrt{\sum_{e\in E} n_e^2(S^*)}\sqrt{\sum_{e\in E} n_e^2(S^{\mathrm{last}(i)-1})} \tag{2}$$

$$= 2 + \frac{2}{n\mathrm{OPT}}\sum_{i=1}^{n}\sqrt{\mathrm{OPT}}\sqrt{C(S^{\mathrm{last}(i)-1})}$$

$$\leq 2 + \frac{2}{n\sqrt{\textsc{Opt}}} \sum_{i=1}^{n} \sqrt{2\Phi(S^{\text{last}(i)-1})} \quad (3)$$

$$\leq 2 + \frac{2}{n\sqrt{\textsc{Opt}}} \sum_{i=1}^{n} \sqrt{2\Phi(S^0)} \quad (4)$$

$$\leq 2 + 2\sqrt{2}\sqrt{\frac{C(S^0)}{\textsc{Opt}}} \quad (5)$$

$$\leq (2 + 2\sqrt{2})\sqrt{\frac{C(S^0)}{\textsc{Opt}}},$$

where inequality (1) follows from Lemma 5, inequality (2) is due to the application of the Cauchy-Schwarz inequality, inequality (3) holds because $C(S^{\text{last}(i)-1}) \leq 2\Phi(S^{\text{last}(i)-1})$, inequality (4) holds because the potential function can only decrease at each best response and inequality (5) holds because $\Phi(S^0) \leq C(S^0)$. □

By applying Lemma 6 to all the pairs of consecutive T-coverings, we are now able to prove the following theorem.

Theorem 3. *Given a linear symmetric congestion game, any best response dynamics satisfying the T-Minimum Liveness Condition converges from any initial state to a state S such that $\frac{C(S)}{\textsc{Opt}} = O(1)$ in at most $T\lceil \log\log n \rceil$ best responses.*

5 Extensions and Final Remarks

All the results extend to the setting of weighted congestion games, in which any player $i \in N$ is associated a weight $w_i \geq 1$; notice that it is possible to assume without loss of generality that $w_i \geq 1$ for any $i \in N$ because it is always possible to suitably scale all the weights (and accordingly the coefficients of the latency functions) in order to obtain such a condition. Let $W = \sum_{i=1}^{n} w_i$. We denote by $l_e(S)$ the congestion on resource e in a state S, i.e. $l_e(S) = \sum_{i|e\in S_i} w_i$. The cost of player i in state S is $c_i(S) = w_i \sum_{e\in S_i} f_e(l_e(S))$. The social cost is given by the sum of the players costs: $C(S) = \sum_{i\in N} c_i(S) = \sum_{e\in E} l_e f_e(l_e(S))$. The following theorems hold.

Theorem 4. *Given a linear weighted congestion game, any best response dynamics satisfying the (T, β)-Fairness Condition converges from any initial state to a state S such that $\frac{C(S)}{\textsc{Opt}} = O(\beta)$ in at most $T\lceil \log\log W \rceil$ best responses.*

Theorem 5. *Given a linear weighted symmetric congestion game, any best response dynamics satisfying the T-Minimum Liveness Condition converges from any initial state to a state S such that $\frac{C(S)}{\textsc{Opt}} = O(1)$ in at most $T\lceil \log\log W \rceil$ best responses.*

As a final remark, it is worth to note that our techniques provide a much faster convergence to efficient states with respect to the previous result in the literature.

In particular, in the symmetric setting, Theorem 3 shows that best response dynamics leads to efficient states much faster than how ϵ-Nash dynamics (i.e., sequences of moves reducing the cost of a player by at least a factor of ϵ) leads to ϵ-Nash equilibria [5]. Furthermore, also in the more general asymmetric setting, Theorem 1 shows that the same holds for fair best response dynamics with respect to ϵ-Nash ones [3].

References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. J. ACM 55(6) (2008)
2. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: STOC, pp. 57–66. ACM (2005)
3. Awerbuch, B., Azar, Y., Epstein, A., Mirrokni, V.S., Skopalik, A.: Fast convergence to nearly optimal solutions in potential games. In: ACM Conference on Electronic Commerce, pp. 264–273. ACM (2008)
4. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: Performance of one-round walks in linear congestion games. Theory Comput. Syst. 49(1), 24–45 (2011)
5. Chien, S., Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. In: SODA, pp. 169–178. SIAM (2007)
6. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: STOC, pp. 67–73. ACM (2005)
7. Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and Approximation in Potential Games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 349–360. Springer, Heidelberg (2006)
8. Yannakakis, M., Johnson, D.S., Papadimitriou, C.H.: How easy is local search? Journal of Computer and System Sciences 37, 79–100 (1988)
9. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The complexity of pure Nash equilibria. In: STOC, pp. 604–612. ACM (2004)
10. Fanelli, A., Flammini, M., Moscardelli, L.: The Speed of Convergence in Congestion Games under Best-Response Dynamics. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part I. LNCS, vol. 5125, pp. 796–807. Springer, Heidelberg (2008)
11. Fanelli, A., Moscardelli, L.: On best response dynamics in weighted congestion games with polynomial delays. Distributed Computing 24, 245–254 (2011)
12. Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
13. Nash, J.F.: Equilibrium points in n-person games. Proceedings of the National Academy of Sciences, 36, 48–49 (1950)
14. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory 2, 65–67 (1973)
15. Skopalik, A., Vöcking, B.: Inapproximability of pure Nash equilibria. In: STOC, pp. 355–364. ACM (2008)