Convergence Issues in Congestion Games

Luca Moscardelli *

Abstract

Congestion games are a widely studied class of non-cooperative games. In fact, besides being able to model many practical settings, they constitute a framework with nice theoretical properties: Congestion games always converge to pure Nash Equilibria by means of improvement moves performed by the players, and many classes of congestion games guarantee a low price of anarchy, that is the ratio between the worst Nash Equilibrium and the social optimum. Unfortunately, the time of convergence to Nash Equilibria, even under best response moves of the players, can be very high, i.e., exponential in the number of players, and in many setting also computing a Nash equilibrium can require a high computational complexity.

Motivated by the above facts, in order to guarantee a fast convergence to Nash Equilibria, in the last decade many computer science researchers focused on special classes of congestion games (e.g., with linear or polynomial delay functions), on simplified structures of the strategy space (e.g., on symmetric games in which all players share the same set of strategies or on matroid congestion games in which the set of strategies constitutes a matroid) and on the relaxation of the notion of Nash Equilibria (e.g., exploiting the notion of ϵ-Nash Equilibria). We survey such attempts that, however, only in some very specific cases have led to satisfactory results on the speed of convergence to Nash Equilibria.

If we relax the constraint of reaching a Nash Equilibrium, and our goal becomes that of reaching states approximating the social optimum by a "low" factor, i.e., a factor being order of the price of anarchy, significantly better results on the speed of convergence under best response dynamics can be achieved.

Interestingly, in the more general asymmetric setting, fairness among players influences the speed of convergence. For instance, considering the fundamental class of linear congestion games, if each player is allowed to play at least once and at most β times every T best responses, states with approximation ratio $O(\beta)$ times the price of anarchy are reached after

*Department of Economic Studies, University of Chieti-Pescara, Viale Pindaro 42, 65127 Pescara, Italy. Email: luca.moscardelli@unich.it

$T\lceil \log \log n \rceil$ best responses, and such a bound is essentially tight also after exponentially many ones. It is worth noticing that the structure of the game implicitly affects its performances in terms of convergence speed: In particular, in the symmetric setting the game always converges to an efficient state after a polynomial number of best responses, regardless of the frequency each player moves with. Most of these results extend to polynomial and weighted congestion games.

1 Introduction

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

In the weighed version of the game, each player is assigned with a non-negative weight and the congestion of each resource is the sum of the weights of the player using that resource.

In this work we focus on a fundamental classes of congestion games, having particular delay functions for their resources. In particular, we consider polynomial delay function with maximum degree d. We are mainly interested in studying and analyzing the time of convergence to solutions being a good approximation of the optimum.

The paper is organized as follows: In the next section we provide the basic notation and definitions. Section 3 surveys the state of the art on convergence issues in congestion games. Section 4 is devoted to the study of the quality of the solution reached after a move performed by each player, while in Section 5 it is studied the speed of convergence to good solutions. Finally, Section 6 gives some conclusive remarks and lists some interesting open problems.

2 Model and Definitions

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

Each player i has a weighted demand $w_i \in \mathbb{R}^+$. Since it is always possible to suitably scale all the weights, without loss of generality, we assume that $w_i \geq 1$ for any $i \in N$; furthermore, we denote by W the sum of the weights of all players, i.e. $W = \sum_{i \in N} w_i$. Moreover, let us denote by w_{max} the maximum weight. Σ_i is the strategy space of player i, and any strategy $s_i \in \Sigma_i$ of player i is a subset of resources, i.e. $\Sigma_i \subseteq 2^E$. Given a *strategy profile* (also called *state*) $S = (s_1, \ldots, s_n)$ and a resource e, we define the congestion $\theta_e(S)$ on resource e by $\theta_e(S) = \sum_{i \in N | e \in s_i} w_i$. A delay function $f_e : \mathbb{R}^+ \mapsto \mathbb{R}^+$ associates to resource e a delay depending on its congestion, so that the cost of player i for the pure strategy s_i is given by the weighted sum of the delays associated with resources in s_i, i.e. $c_i(S) = \sum_{e \in s_i} w_i f_e(\theta_e(S))$.

In this paper we will focus on congestion games with *polynomial* delay functions with maximum degree d and non-negative coefficients, that is for every resource $e \in E$ the delay function is of the form $f_e(x) = \sum_{j=1}^{d} a_{e,j} x^j$ with $a_{e,j} \geq 0$ for all $j = 0, \ldots, d$. We call a congestion game *linear* if $d = 1$, that is $f_e(x) = a_e x + b_e$ for every resource $e \in E$, with $a_e, b_e \geq 0$.

We call a congestion game *unweighted* whenever, for each player $i \in N$, $w_i = 1$. In this case, the congestion $\theta_e(S)$ on resource e is also denoted by $n_e(S)$ and it is equal to $|\{i \in N | e \in s_i\}|$, i.e., $n_e(S) = \theta_e(S) = |\{i \in N | e \in s_i\}|$.

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., $C(S) = \sum_{i \in N} c_i(S) = \sum_{e \in E} \theta_e(S) f_e(\theta_e(S))$. 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'.

A notable special class of congestion games is that of *Network Congestion Games*, in which we are given a communication graph G whose edges are the resources of the congestion game; each player is associated to a source and a destination node of G and her strategies are given by all the resources corresponding to the edges of the simple paths connecting her source with her destination in G.

Finally, a *symmetric* game is a game in which all the players share the same set of strategies.

Each player acts selfishly and aims at choosing the strategy decreasing its cost, given the strategy choices of other players. For a strategy profile $S = (s_1, s_2, \ldots, s_n)$, an *improvement move* of player i is a strategy s_i' such that $c_i(S_{-i}, s_i') < c_i(S)$. 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.

A state S is a *Nash equilibrium* if for every player $i \in N$ and all strategy $s_i' \in \Sigma_i$, $c_i(S) \leq c_i(S_{-i}, s_i')$, i.e., no player in N can improve her individual cost by unilaterally changing her strategy. The *price of anarchy* (PoA) is the ratio $C(S)/\text{Opt}$, where S is the worst Nash equilibrium, i.e. it is the approximation ratio of the Nash equilibrium having maximum social cost. Notice that Nash equilibria correspond to sinks in the Nash dynamics graph.

By extending the Nash dynamics and the Nash equilibrium concepts, it is possible to define their approximated versions as follows: Given a strategy profile $S = (s_1, s_2, \ldots, s_n)$, an *$\epsilon$-improvement move* of player i is a strategy s_i' such that $c_i(S_{-i}, s_i') < (1-\epsilon)c_i(S)$. S is an ϵ-Nash equilibrium if no player has an ϵ-improvement move. An *ϵ-better response* of player i in S is either an ϵ-improvement move, if it is available, or the current strategy s_i, otherwise. An *ϵ-better response dynamics* any sequence of ϵ-better responses.

In Section 3 we will briefly survey the results achieved in the literature with respect to the speed of convergence to Nash equilibria and ϵ-Nash equilibria. Then, in the remainder of the paper, we will study the approximation ratio of states reached after sequences of best responses, that we call a *best response dynamics*.

The selfish behavior of players performing best responses can be modelled by the *(Best Response) Nash Dynamics Graph*. Formally the Nash Dynamics Graph associated to a congestion game $\mathcal{G}$ is a directed graph $\mathcal{B} = (V, A)$ where each vertex in V corresponds to a strategy profile and there is an edge $(S, S') \in A$ with label i, where $S' = (S_{-i}, s_i')$ and $s_i' \in \Sigma_i$, if and only if both the following conditions are met: (*i*) s_i' is a best response of i in S; (*ii*) if $S \neq S'$, s_i' is also an improvement move of i in S. Observe that $\mathcal{B}$ may contain loops, corresponding to best response in which a player maintains her current strategy. A *best response walk* is a directed walk in $\mathcal{B}$.

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:

Definition 1 (T-Minimum Liveness Condition)**.** *Given any $T \geq n$, a best response dynamics satisfies the T-Minimum Liveness Condition if each player performs at least a best response every T consecutive responses.*

We consider the following notions of best response walks, that are a refinement of the ones introduced in [10, 21]; notice that all of them satisfy the Minimum Liveness Condition above stated.

1**-round walk**: it's a best response walk $R = \left(\left(S_R^0, S_R^1\right), \left(S_R^1, S_R^2\right), \ldots, \left(S_R^i, S_R^{i+1}\right), \ldots, \left(S_R^{n-1}, S_R^n\right)\right)$ in $\mathcal{B}$ of length n, where the edge (S_R^i, S_R^{i+1})

has label $\pi_R(i)$ for every $0 \leq i \leq n-1$, i.e. $\pi_R(i)$ is the player performing the i-th best response of the walk. π_R is such that every player performs exactly a best response in R. S^0_R is said the *initial* state of R and S^n_R its *final* state. For simplicity we denote R by a sequence of states, i.e. $R = \left(S^0_R, \ldots, S^n_R\right)$.

(ℓ,β)-bounded covering walk: it's a best response walk $R = \left(\left(S^0_R, S^1_R\right), \left(S^1_R, S^2_R\right), \ \ldots, \left(S^i_R, S^{i+1}_R\right), \ldots, \left(S^{\ell-1}_R, S^\ell_R\right)\right)$ in $\mathcal{B}$ of length ℓ, where the edge (S^i_R, S^{i+1}_R) has label $\pi_R(i)$ for every $0 \leq i \leq \ell - 1$, i.e. $\pi_R(i)$ is the player performing the i-th best response of the walk. π_R is such that every player performs at least a best response and at most β best responses in R. S^0_R is said the *initial* state of R and S^ℓ_R its *final* state. For simplicity we denote R by a sequence of states, i.e. $R = \left(S^0_R, \ldots, S^\ell_R\right)$. 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 $S^{\mathrm{last}_R(i)}$.

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*.

Definition 2 ((T,β)-Fairness Condition)**.** *Given any $T \geq n$, a best response dynamics satisfies the (T,β)-Fairness Condition if it ca be decomposed in k (ℓ,β)-bounded covering walks such that $\ell \leq T$ for all such coverings.*

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^i_R, π_R and $last(i)_R$, respectively.

We will consider the following two scenarios:

1. all walks are assumed to start from an arbitrary initial state;

2. all walks are assumed to start from the “empty" state in which no receiver has still selected its communication path from the source; in order to include such a situation in the above framework we include an additional special empty path $\emptyset$ in the strategies sets P_i, assuming that the only possible best response moves from $\emptyset$ are the ones selecting a path of minimum payment actually connecting the source to the receiver; all the notation and definitions are trivially extended accordingly.

3 Related Work and Our Contribution

On the one hand, Rosenthal [23] has shown, by a potential function argument, that for unweighted congestion games the natural decentralized mechanism known as

Nash dynamics, in which at each step some player performs an improvement step switching her strategy to a better alternative, is guaranteed to converge to a pure Nash equilibrium [22], i.e. a fixed point in which no player can perform an improvement step (note that in a Nash dynamics players play their improvement steps sequentially, and not in parallel). On the other hand, weighted congestion games do not necessarily admit a Nash equilibrium [20] unless specific settings are considered (linear delay functions [18], singleton congestion games [17], matroid congestion games [2]). With a little abuse of notation, also in the context of weighted congestion games that do not admit a Nash equilibrium, we call Nash dynamics the dynamics in which each player aims at minimizing its current cost, given the strategic choices of all the other players. More formally, an *exact* real-valued potential function Φ defined on the set of states of the game satisfies the property that for each player i and each strategy $s'_i \in \Sigma_i$ of i in S, it holds that $c_i(S_{-i}, s'_i) - c_i(S) = \Phi(S_{-i}, s'_i) - \Phi(S)$. Unweighted congestion games admit (see [23]) the (exact) potential function

$$\Phi(S) = \sum_{e \in E} \sum_{j=1}^{n_e(S)} f_e(j).$$

It is worth noticing that in unweighted congestion games with polynomial delays with maximum degree d, for any state S, it holds $\Phi(S) \leq C(S) \leq (d+1)\Phi(S)$. Unfortunately, in general, weighted congestion games do not admit a potential function; nevertheless, it is possible to show that weighted congestion games with linear delays do (see [18]), and such a potential function is

$$\Phi(S) = \frac{1}{2}\left(\sum_{e \in E} \theta_e(S) f_e(\theta_e(S)) + \sum_{i=1}^{n} \sum_{e \in s_i} w_i f_e(w_i)\right).$$

It is worth noticing that in weighted linear congestion games, for any state S, it holds $\Phi(S) \leq C(S) \leq 2\Phi(S)$.

Even in the unweighted setting in which Nash equilibria are guaranteed to exist, Fabrikant et al. [12] show that such dynamics may require a number of steps exponential in the number of players n in order to reach such an equilibrium. Their analysis relates congestion games to local search problems by showing that it is PLS-complete [11] to compute a Nash equilibrium for general unweighted congestion games. Moreover from their completeness proof and from previous results about local search problems, it follows that there exist congestion games with initial states such that any improvement sequence starting from these states needs an exponential number of steps in order to reach a Nash equilibrium. More recently, Ackermann et al. [1] show that the previous negative result holds even in the restricted case of linear unweighted congestion games. On the positive side, Fabrikant et al. [12], by exploiting a reduction to the Min-cost flow problem, show that it is possible to Compute a Nash equilibrium in symmetric Network congestion games with non-decreasing delay function in polynomial time. Such a result can be also extended to network congestion games in which all the players share the same source (multicast games). Unfortunately, such a result does not imply a fast convergence to Nash equilibria, as Ackermann et al. [1] show the existence of network congestion games in which any sequence of improvement moves is exponentially long in n. It is possible to obtain a polynomial convergence by restricting the combinatorial structure of congestion games, in particular by imposing that the strategy set of each player is the basis of a matroid; note that load balancing games, in which every strategy is composed by only one resource, belong to such a class of congestion games.

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 [24], also the problem of finding an ϵ-Nash equilibrium in (unweighted) congestion games is PLS-complete for any ϵ, though, under some restrictions on the delay functions, Chien and Sinclair [8] proved, for some classes of delay functions including the polynomial ones, 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.

In order to measure the degradation of social welfare due to the selfish behavior of the players, Koutsoupias and Papadimitriou [19] 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. [4] and Christodoulou and Koutsoupias [9]. They both proved that the price of anarchy for congestion games with linear delays is $5/2$.

The performances of Nash equilibria in unweighted and weighted congestion games with polynomial delay functions have been investigated by Aland et al. [3], who have provided a tight price of anarchy both for the unweighted case and the weighted one. In particular they have improved the previous results due to Awerbuch et al. [4] (for the weighted case) and Christodoulou and Koutsoupias [9] (for the unweighted case).

Since negative results tend to dominate the issues relative to the complexity of computing equilibria, also considering that in the more general setting of weighted congestion games Nash equilibria may not be guaranteed to exist, another natural arising question is whether efficient states (with a social cost constant or at least comparable to the one of any Nash equilibrium) can be reached by best response moves in a reasonable amount of time (e.g., [5, 10, 13, 14]). We measure the efficiency of a state by the ratio among its cost and the optimal one, i.e., its approximation ratio. We generally say that a state is efficient when its approximation ratio is within a constant factor from the price of anarchy.

Awerbuch et al. [5] have proved that for some classes of delay functions (including the polynomial delay functions), 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 [5].

Nevertheless, we are able to show that best response dynamics in congestion games with polynomial delays can actually fast converge to good solutions, provided that the dynamics obeys some mild constraint on the order of the moving players.

Our Contribution. In this paper we first provide some preliminary result [6] on the approximation ratio of the solutions achieved after a one-round walk in linear congestion games. For one-round walks starting from the empty strategy profile, we close the existing gap between the upper bound of $2+\sqrt{5}\approx 4.24$ given in [10] and the lower bound of 4 derived in [7] by providing a family of instances yielding lower bounds approaching $2+\sqrt{5}$ as the number of players goes to infinity. The construction and the analysis of these instances require nontrivial arguments.

Then, we focus on the more general setting in which the dynamics start from a generic state.

We show [13, 14] that, for weighted congestion games with polynomial delays, the approximation ratio achieved after a sequence of k $O(1)$-bounded covering walks is $O\left(W^{d\left(\frac{d}{d+1}\right)^{k-1}}\right)$ and $\Omega\left(\frac{W^{d\left(\frac{d}{d+1}\right)^{k-1}}}{k}\right)$ (which is asymptotically matching for constant values of k), where W is the sum of the players' weights. As a consequence, we prove that, for any given d, $\Theta(\log\log W)$ $O(1)$-bounded covering walks are necessary and sufficient to achieve a constant factor approximate solution.

Finally, we completely characterize [15] 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 [5] and [13, 14] for linear congestion games. On the one hand, in [5] 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 [13, 14] 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 $T\lceil\log\log W\rceil$ 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 [13, 14] is not only sufficient, but also necessary in order to reach efficient states after a polynomial number of best responses. Furthermore, in the special case of symmetric congestion games with linear delays, 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 $T\lceil\log\log W\rceil$ best responses efficient states are always reached.

4 One-Round Analysis [6]

Christodoulou et al. [10] proved that for any linear congestion game and 1-round walk $(S^0, S^1, \ldots, S^n)$ with $S^0 = \emptyset$, it holds $C(S^n) \leq (2 + \sqrt{5})\textsc{Opt} \approx 4.24\textsc{Opt}$. Surprisingly enough, the best known lower bound is the one derived by Caragiannis at al. [7] for the restricted case of load balancing on identical servers, which poses $C(S^n) \geq 4\textsc{Opt} - \epsilon$, for any $\epsilon > 0$. We close this gap by proving that, for any $\epsilon > 0$, there always exist a linear congestion game and a 1-round walk $(S^0, S^1, \ldots, S^n)$, with $S^0 = \emptyset$, such that $C(S^n) \geq (2 + \sqrt{5} - \epsilon)\textsc{Opt}$.

Given three positive integers n, k and o, with $n \geq 2k + o - 1$ and $k \geq 2o$, we define the game $\mathcal{G}_{n,k,o}$ 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.

- $s_i = \{e_i\}$ and $s'_i = \{e_{i+1}\} \cup \bigcup_{j=k+1}^{k+i} \{e_j\}$, for any $i \in [k-1]$;
- $s_k = \{e_k\}$ and $s'_k = \bigcup_{j=k+1}^{2k} \{e_j\}$;
- $s_i = \bigcup_{j=k+1}^{i} \{e_j\}$ and $s'_i = \bigcup_{j=i+1}^{k+i} \{e_j\}$, for any $k + 1 \leq i \leq k + o$;

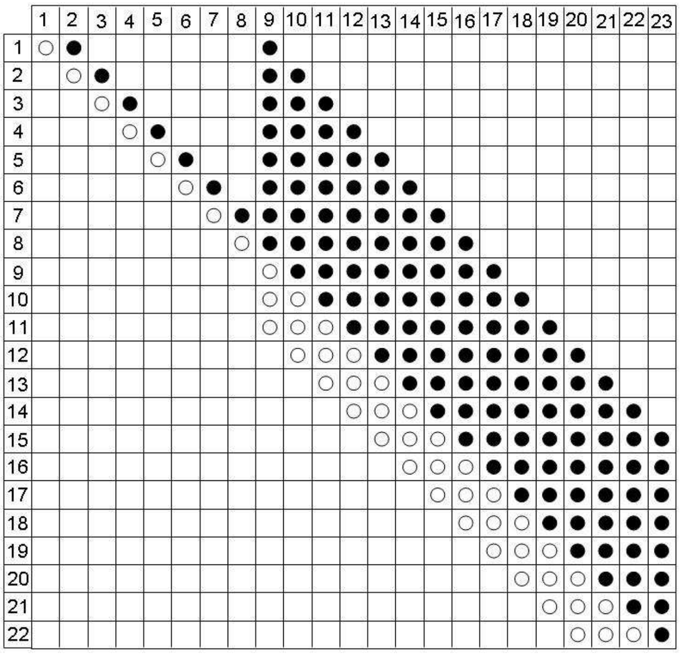

Figure 1: The set of strategies available to each player in the game $\mathcal{G}_{22,8,3}$. Rows are associated with players, while columns with resources. White and black circles represent the first and the second strategy, respectively.

- $s_i = \bigcup_{j=i-o+1}^{i} \{e_j\}$ and $s'_i = \bigcup_{j=i+1}^{\min\{k+i,m\}} \{e_j\}$, 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 $j \in [m]$ we associate the linear latency function $f_j(x) = a_j \cdot x$ with resource e_j, where each a_j 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:

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

	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 $\mathcal{G}_{22,8,3}$.

- $$2a_k - \sum_{j=k+1}^{2k} \Big((2k - j + 1)a_j\Big) = 0,$$

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

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

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.

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

Let $a = (a_1, \dots, a_m)^T$. 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 1. *The system of linear equations $Ba = 0$ admits a strictly positive solution.*

We claim that the strategy profile in which all players choose the second of their strategies is a possible outcome for a 1-round walk starting from the empty strategy profile.

Lemma 2. *For any game $\mathcal{G}_{n,k,o}$, there exists a* 1*-round walk* $(S^0, S^1, \ldots, S^n)$ *such that* $S^0 = \emptyset$ *and* $S^n = (s'_1, \ldots, s'_n)$.

Proof. The claim is a direct consequence of the definition of system A. □

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 3. *In any solution of system A it holds* $a_1 \leq 4 \sum_{j=k+1}^{2k} a_j$.

Lemma 4. *In any solution of system A it holds*

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

We can now prove our main result.

Theorem 1. *For any* $\epsilon > 0$*, there exist a linear congestion game* $\mathcal{G}_{n,k,o}$ *and a* 1*-round walk* $(S^0, S^1, \ldots, S^n)$*, with* $S^0 = \emptyset$*, such that* $C(S^n) \geq (2 + \sqrt{5} - \epsilon)\textsc{Opt}$.

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 big n, these values are consistent with the definition of $\mathcal{G}_{n,k,o}$ 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

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

Let $S^* = (s_1 \ldots, s_n)$ be the strategy profile in which all players choose the first of their strategies. Because of Lemma 2, we have that there exists a 1-round walk

$(S^0, S^1, \dots, S^n)$ such that $S^0 = \emptyset$ and $S^n = (s'_1, \dots, s'_n)$. By comparing the social costs of S^n and S^*, we obtain

$$\frac{C(S^n)}{\textsc{Opt}} \geq \frac{C(S^n)}{C(S^*)} \geq \frac{\displaystyle\sum_{i=2}^{k} a_i + k^2 \sum_{i=k+1}^{m} a_i}{\displaystyle\sum_{i=1}^{k} a_i + o^2 \sum_{i=k+1}^{m} a_i},$$

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

By using Equality 1, we get

$$\frac{C(S^n)}{\textsc{Opt}} \geq \frac{\displaystyle\sum_{i=2}^{k} a_i + k^2 \sum_{i=k+1}^{m} a_i}{\displaystyle\sum_{i=1}^{k} a_i + o^2 \sum_{i=k+1}^{m} a_i} =$$

$$\frac{\left((k+1)\left(\frac{k}{2} - o\right) + k^2\right) \displaystyle\sum_{i=k+1}^{m} a_i + (k+1) \sum_{i=m-o+1}^{m} ((i - m + o)a_i) - a_1}{\left((k+1)\left(\frac{k}{2} - o\right) + o^2\right) \displaystyle\sum_{i=k+1}^{m} a_i + (k+1) \sum_{i=m-o+1}^{m} ((i - m + o)a_i)} \geq$$

$$\frac{\left((k+1)\left(\frac{k}{2} - o\right) + k^2 + \frac{k^3}{n-2k-o+1} - 4\right) \displaystyle\sum_{i=k+1}^{m} a_i}{\left((k+1)\left(\frac{k}{2} - o\right) + o^2 + \frac{k^3}{n-2k-o+1}\right) \displaystyle\sum_{i=k+1}^{m} a_i},$$

where, in the last inequality, we have used Lemmas 3 and 4 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 $\lim_{k\to\infty} \frac{C(S^n)}{\textsc{Opt}} \geq \lim_{k\to\infty} \frac{k\left(\frac{k}{2}-o\right)+k^2}{k\left(\frac{k}{2}-o\right)+o^2} = \lim_{k\to\infty} \frac{\frac{\sqrt{5}-2}{2}k^2+k^2}{\frac{\sqrt{5}-2}{2}k^2+\frac{7-3\sqrt{5}}{2}k^2} = \lim_{k\to\infty} \frac{\frac{\sqrt{5}}{2}k^2}{\frac{5-2\sqrt{5}}{2}k^2} = \frac{\sqrt{5}}{5-2\sqrt{5}} = 2 + \sqrt{5}$, which implies the claim. □

Some numerical results, obtained on particular games, are shown in Figure 3. It is possible to appreciate there that the value of k may be chosen much higher than the bound $\lfloor \sqrt[4]{n} \rfloor$ fixed in the proof of Theorem 1. This is due to the fact that, for the sake of simplicity, the bound proved in Lemma 4 is really far from being tight.

n	k	o	$\frac{SUM(S^n)}{SUM(S^*)}$
70	8	3	4.001152
100	8	3	4.012482
500	80	30	4.185590
700	80	30	4.208719
1000	100	38	4.216734
1500	100	38	4.220854
2000	200	76	4.224342
3000	300	114	4.226854

Figure 3: Lower bounds on the approximation ratio of the solution achieved after a 1-round walk starting from the empty strategy profile in the games $\mathcal{G}_{n,k,o}$ for some particular values of n, k and o.

5 Convergence to Good Solutions

In this section we provide upper and lower bounds 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.

We first provide (in Subsection 5.1) an upper bound to the the social cost of the state achieved after a best response dynamics satisfying the (T,β)-Fairness condition with $\beta = O(1)$ and starting from an arbitrary state, and then (in Subsection 5.2) we deal with the case of general values of β.

All the results hold for congestion games having polynomial delay functions with non-negative coefficients and maximum degree d, i.e. for every $e \in E$, $f_e(x) = \sum_{j=0}^{d} a_{e,j}x^j$ with $a_{e,j} \geq 0$ for all $j = 0, \ldots, d$. Without loss of generality, we can assume that for every $e \in E$, $f_e(x) = x^{d_e}$ with $0 \leq d_e \leq d$.

In fact, given a congestion game $\mathcal{G}$ having polynomial delays with coefficients being non-negative integers and maximum degree d, it is possible to obtain an equivalent congestion game $\mathcal{G}'$ (i.e, a congestion game having a Nash Dynamics Graph isomorphic to the one of $\mathcal{G}$ and in which any strategy profile S' corresponding to the strategy profile S of $\mathcal{G}$ is such that for every player $i \in N$, $c_i(S) = c_i(S')$), having the same set of players and delay functions of the form $f_e(x) = x^{d_e}$ in the following way. For each resource e in $\mathcal{G}$, we include in $\mathcal{G}'$ a set $A_{e,j}$ of $a_{e,j}$ resources for $j = 0, \ldots, d$ with delay function $f_e(x) = x^j$; moreover, given any strategy set $s_i \in \Sigma_i$ in $\mathcal{G}$, $i = 1, \ldots, n$, we build a corresponding strategy set $s'_i \in \Sigma'_i$ (in $\mathcal{G}'$) by including in s'_i, for each $e \in s_i$, all the resources in the sets $A_e = \cup_{j=0}^{d} A_{e,j}$.

If the coefficients $a_{e,j}$ are rationals (not all being integers) we can perform a similar reduction by exploiting a simple scaling argument.

Finally, if the coefficients are real numbers (not all being rationals), we can obtain from a congestion game $\mathcal{G}$ a new game $\mathcal{G}'$ with coefficients being rationals by approximating the real numbers with a sufficiently high precision so that any best response of $\mathcal{G}$ corresponds to a "quasi"-best response of $\mathcal{G}'$, up to an additive ϵ depending on the precision of the considered approximation.

Moreover, we can also assume without loss of generality that all the player weights are at least 1, by suitably scaling all the weights if the considered game does not satisfy such a property.

Therefore, in the sequel of this section we will consider, for every $e \in E$, $f_e(x) = x^{d_e}$, with $d = \max_e d_e$ and $w_i \geq 1$ for every $i = 1, \ldots, n$.

5.1 $\beta = O(1)$ [14]

The following claim, whose proof can be found in [16] (Lemma 3.6 of [16], with $\gamma = \frac{1}{2}$) will be useful in the sequel.

Claim 1 ([16]). *For every pair of reals $x, y \geq 0$ and every integer $d \geq 1$, it holds $x^d \geq \frac{1}{2^{d-1}}(x+y)^d - y^d$.*

Let $R = \left(S^0, \ldots, S^\ell\right)$ be a (T,β)-bounded covering walk. Given the optimal strategy profile S^*, since the i-th moving player $\pi(i)$ before moving can always select the strategy she would use in S^*, $c_{\pi(i)}(S^i)$ (that is the player's cost immediately after her best response) can be suitably upper bounded by $\sum_{e \in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e(S^{i-1}) + w_{\pi(i)}\right)^{d_e}$. In order to state our results we define the following function

$$\rho(R) = \sum_{i=1}^{\ell} \sum_{e \in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e(S^{i-1}) + w_{\pi(i)}\right)^{d_e},$$

which, by the same argument explained earlier, clearly represents an upper bound to $\sum_{i=1}^{\ell} c_{\pi(i)}(S^i)$.

Lemmas 5 and 6 provide a lower and an upper bound to $\rho(R)$, respectively. From such Lemmas, we can easily derive the approximation achieved after a (T,β)-bounded covering walk.

Lemma 5. *For any $\beta \geq 1$, given a (T,β)-bounded covering walk R ending in S^ℓ, it holds $\rho(R) \geq \frac{C(S^\ell)}{(d+1)}$.*

Proof. Since the players perform best responses, inequality (2) below holds. In order to justify inequality 3, let us consider a resource e. Recall that the cost $c_{\pi(i)}(S^i)$ incurred by a player $\pi(i)$ on e is $w_{\pi(i)} f_e(\theta_e(S^i))$; since f_e is a non decreasing function and $\theta_e(S^i)$ is given by the sum of the players already on e at S^{i-1} plus $w_{\pi(i)}$, the sum of all the cost that players using e incur on e can be lower bounded by considering the resource used by many players each having infinitesimal weight; thus, the summation $\sum_{e\in E}\sum_{i\in\{1,\dots,\ell\}|e\in s_{\pi(i)}} w_{\pi(i)}\theta_e^{d_e}(S^i)$ can be replaced by the integral $\int_{x=0}^{\theta_e(S^\ell)} x^{d_e}dx$. Therefore, by recalling the definition of $\rho(R)$,

$$\begin{aligned}
\rho(R) &= \sum_{i=1}^{\ell}\sum_{e\in s^*_{\pi(i)}} w_{\pi(i)}\left(\theta_e(S^{i-1})+w_{\pi(i)}\right)^{d_e} \\
&\geq \sum_{i=1}^{\ell} c_{\pi(i)}(S^i) && (2)\\
&= \sum_{e\in E}\sum_{\substack{i\in\{1,\dots,\ell\}|\\ e\in s_{\pi(i)}}} w_{\pi(i)}\theta_e^{d_e}(S^i) \\
&\geq \sum_{e\in E}\int_{x=0}^{\theta_e(S^\ell)} x^{d_e}dx && (3)\\
&\geq \frac{1}{(d+1)}\sum_{e\in E}\theta_e^{d_e+1}(S^\ell) = \frac{C(S^\ell)}{(d+1)}.
\end{aligned}$$

□

Lemma 6. *For any $\beta \geq 1$, given a (T,β)-bounded covering walk R, it holds $\rho(R) \leq \beta(W + w_{max})^d$*Opt.

Proof. By the definition of $\rho(R)$, it holds

$$\begin{aligned}
\rho(R) &= \sum_{i=1}^{\ell}\sum_{e\in s^*_{\pi(i)}} w_{\pi(i)}\left(\theta_e(S^{i-1})+w_{\pi(i)}\right)^{d_e} \\
&\leq \sum_{i=1}^{\ell}\sum_{e\in s^*_{\pi(i)}} w_{\pi(i)}(W+w_{max})^{d_e} \\
&\leq (W+w_{max})^d\sum_{i=1}^{\ell}|s^*_{\pi(i)}|w_{\pi(i)} \\
&\leq \beta(W+w_{max})^d\textsc{Opt}, && (4)
\end{aligned}$$

where (4) holds by observing that $\textsc{Opt} \geq \frac{1}{\beta}\sum_{i=1}^{\ell}|s^*_{\pi(i)}|w_{\pi(i)}$.

□

As an immediate consequence of Lemmas 5 and 6, $\mathrm{Apx}_1^{\beta}(\mathcal{G}) \leq \beta(d+1)(W + w_{max})^d$. Lemmas 7 and 8 will be useful to extend such result to a (T,β)-bounded k-covering walk $P = \langle R_1,\dots,R_k\rangle$ by exploiting the relationship among two consecutive walks. To this aim, recalling that $R = \left(S^0,\dots,S^\ell\right)$, we define the following function

$$H(R) = \sum_{i=1}^{\ell} \sum_{e \in s^*_{\pi(i)}} w_{\pi(i)} \theta_e^{d_e}(S^0).$$

By exchanging the order of the summations, we obtain

$$H(R) = \sum_{e \in E^*} \theta_e^{d_e}(S^0) x_e,$$

where $x_e = \sum_{j \in X_e} w_{\pi(j)}$ and $X_e = \left\{ i \in \{1, \dots, \ell\} | e \in s^*_{\pi(i)} \right\}$. Informally speaking, $H(R)$ represents the sum over all the moves in the covering walk R of the weighted delay that the moving player $\pi(i)$ would experience in the initial state S^0 of R on her optimal strategy $s^*_{\pi(i)}$.

Clearly, as it can be noticed by their definition, $H(R)$ and $\rho(R)$ are correlated and Lemma 7 (resp. Lemma 8) will provide a lower bound (resp. an upper bound) to $\rho(R)$ in terms of $H(R')$ (resp. $H(R)$), where R and R' are two consecutive covering walks, that is the final state of R coincides with the initial one of R'.

Roughly speaking, on the one hand, Lemma 7 shows that the ratio between $H(R')$ and $\textsc{Opt}$ is significantly less than the one between $\Gamma(R)$ and $\textsc{Opt}$; in other words, recalling that $\Gamma(R)$ is an upper bound to $\sum_{i=1}^{\ell_R} c_{\pi_R(i)}(S^i_R)$ and that $H(R')$ is the sum over all the moves in R' of the delays that the moving players would experience, in the first state of R', on her optimal strategies, we are able to show that in the first state of the next covering R' the congestion on the resources used in the considered optimal solution is such that $\frac{H(R')}{\textsc{Opt}} \leq \beta(d+1)\left((d+1)\frac{\rho(R)}{\textsc{Opt}}\right)^{\frac{d}{d+1}}$. On the other hand, Lemma 8 shows that the ratio between $\Gamma(R')$ and $\textsc{Opt}$ is not much greater than the one between $H(R')$ and $\textsc{Opt}$. ; in other words, $H(R')$, even if referring to the delay in the first state of R', can be already considered as an "approximated" bound to $\Gamma(R')$, that is in turn an upper bound to $\sum_{i=1}^{\ell_{R'}} c_{\pi_{R'}(i)}(S^i_{R'})$. Therefore, the combination of Lemmata 7 and 8 gives that the ratio between $\Gamma(R')$ and $\textsc{Opt}$ is significantly less than the one between $\Gamma(R)$ and $\textsc{Opt}$.

In fact, by combining $k-1$ times the results of Lemmata 7 and 8, Theorem 2 finally derives an upper bound to $\mathrm{Apx}_k^{\beta}(\mathcal{G})$.

Lemma 7. *For any $\beta \geq 1$, given two consecutive (ℓ,β)-bounded covering walks R and R' such that the final state of R coincides with the initial one of R', it holds*

$$\frac{H(R')}{\textsc{Opt}} \leq \beta(d+1)\left((d+1)\frac{\rho(R)}{\textsc{Opt}}\right)^{\frac{d}{d+1}}.$$

Proof. Since the final state S^{ℓ} of R corresponds to the initial state of R', and recalling that each player can perform at most β best response within the same

(ℓ,β)-bounded covering walk, we obtain

$$
\begin{aligned}
H(R') &\leq \beta \sum_{e\in E^*} \theta_e^{d_e}(S^\ell)\theta_e(S^*) \\
&\leq \beta \sum_{j=0}^{d} \sum_{e\in E^*:d_e=j} \theta_e^{j}(S^\ell)\theta_e(S^*) \\
&\leq \beta \sum_{e\in E^*:d_e=0} \theta_e(S^*) + \\
&\quad +\beta \sum_{j=1}^{d} \left(\left(\sum_{e\in E^*:d_e=j} (\theta_e^{j}(S^\ell))^{\frac{j+1}{j}} \right)^{\frac{j}{j+1}} \cdot \right. \\
&\quad \left. \cdot \left(\sum_{e\in E^*:d_e=j} (\theta_e^{j+1}(S^*)) \right)^{\frac{1}{j+1}} \right) \qquad (5)
\end{aligned}
$$

$$
\begin{aligned}
&= \beta \sum_{e\in E^*:d_e=0} \theta_e(S^*) + \\
&\quad \beta \sum_{j=1}^{d} \left(\left(\sum_{e\in E^*:d_e=j} (\theta_e^{j+1}(S^\ell)) \right)^{\frac{j}{j+1}} \cdot \right. \\
&\quad \left. \cdot \left(\sum_{e\in E^*:d_e=j} (\theta_e^{j+1}(S^*)) \right)^{\frac{1}{j+1}} \right)
\end{aligned}
$$

$$
\begin{aligned}
&\leq \beta\textsc{Opt} + \beta \sum_{j=1}^{d} \left(\left(C(S^\ell) \right)^{\frac{j}{j+1}} \textsc{Opt}^{\frac{1}{j+1}} \right) \\
&= \beta \sum_{j=0}^{d} \left(\left(C(S^\ell) \right)^{\frac{j}{j+1}} \textsc{Opt}^{\frac{1}{j+1}} \right) \\
&\leq \beta \sum_{j=0}^{d} \left(\left((d+1)\frac{\rho(R)}{\textsc{Opt}}\textsc{Opt} \right)^{\frac{j}{j+1}} \textsc{Opt}^{\frac{1}{j+1}} \right) \qquad (6) \\
&= \beta \sum_{j=0}^{d} \left(\left((d+1)\frac{\rho(R)}{\textsc{Opt}} \right)^{\frac{j}{j+1}} \textsc{Opt} \right) \\
&\leq \beta(d+1) \left((d+1)\frac{\rho(R)}{\textsc{Opt}} \right)^{\frac{d}{d+1}} \textsc{Opt}.
\end{aligned}
$$

where (5) follows from Hölder's inequality, stating that, for r and s such that $\frac{1}{r} + \frac{1}{s} = 1$,

$$\sum_{j=1}^{q} a_j b_j \le \left(\sum_{j=1}^{q} a_j^r\right)^{1/r} \left(\sum_{j=1}^{q} b_j^s\right)^{1/s}$$

by replacing r with $\left(\frac{d+1}{d}\right)$ and s with $(d+1)$, and (6) follows from Lemma 5.

□

Lemma 8. *For any $\beta \ge 1$, given a (ℓ,β)-bounded covering walk R, it holds*

$$H(R) \ge \left(\frac{1}{2^{d-1}} - \frac{d}{\alpha}\right)\rho(R) - \beta\left((\alpha\beta)^d + 1\right)\textsc{Opt},$$

for any $\alpha > d\, 2^{d-1}$.

Proof. In order to lower bound $H(R)$ with respect to $\rho(R)$, we define the following suitable potential function $h_i(R) = \sum_{e\in E^*} g_e(S^i)x_e^{>i}$ for $i \in \{0,\dots,\ell\}$, where for a generic state S, $g_e(S) = \max\{0,\ f_e(\theta_e(S)) - f_e(\alpha\beta\theta_e(S^*))\}$ and $x_e^{>k} = \sum_{j\in X_e^{>k}} w_{\pi(j)}$ where $X_e^{>k} = \left\{i \in \{k+1,\dots,\ell\} | e \in s^*_{\pi(i)}\right\}$. Informally speaking, such a potential function takes into account the delay due to the congestion of the not yet moving players during walk R above a "virtual" congestion frontier given by all the values $\alpha\beta n_e(S^*)$. Let $\Delta_i(R) = h_{i-1}(R) - h_i(R)$ for $i \in \{1,\dots,\ell\}$. Notice that by the definition of the potential function $h_i(R)$, since $h_\ell(R) = 0$, $\sum_{i=1}^{\ell} \Delta_i(R) = h_0(R) \le H(R)$, that is a lower bound for $\sum_{i=1}^{\ell} \Delta_i(R)$ is also a lower bound for $H(R)$; therefore, in the following we focus on lower bounding $\sum_{i=1}^{\ell} \Delta_i(R)$.

Consider a generic step i in walk R, in which player $\pi(i)$ performs a best response by selecting resources in $s^i_{\pi(i)}$ and let us bound from below the value of $\Delta_i(R)$ by evaluating how much player $\pi(i)$ removes from $h_{i-1}(R)$ and how much she adds to $h_i(R)$. Player $\pi(i)$ in order to obtain $h_i(R)$ removes at least $\sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} g_e(S^{i-1})$ from $h_{i-1}(R)$, due to the decrease of the coefficients $x_e^{>(i-1)}$ to $x_e^{>i}$. Let us evaluate how much player $\pi(i)$ adds to $h_i(R)$. Player $\pi(i)$ increases the value of $h_i(R)$ only by resources whose congestion is above the virtual frontier after player $\pi(i)$ plays her best response. Thus for each resource $e \in s^i_{\pi(i)}$ such that $\theta_e(S^i) > \alpha\beta\theta_e(S^*)$, the increase of $h_i(R)$ is equal to $\left(g_e(S^i) - g_e(S^{i-1}))\right) x_e^{>i}$ which, by the definition of g_e, is equal to $\left(f_e(\theta_e(S^i)) - f_e(\theta_e(S^{i-1}))\right) x_e^{>i}$. Since f_e is convex, such quantity is at most $\left(\theta_e(S^i) - \theta_e(S^{i-1})\right) f_e'(\theta_e(S^i))x_e^{>i}$, that is equal to $w_{\pi(i)} f_e'(\theta_e(S^i))x_e^{>i}$, where f_e' is the derivative of f_e. Moreover since $x_e^{>i} \le x_e \le \beta\theta_e(S^*) \le \theta_e(S^i)/\alpha$, we obtain that the increase for each resource e is at most $w_{\pi(i)} f_e'(\theta_e(S^i))\theta_e(S^i)/\alpha$. Thus, considering the previous quantity as an upper bound of the increase for all the resources in $s^i_{\pi(i)}$, player $\pi(i)$ in order to obtain $h_i(R)$ adds at most $\frac{1}{\alpha}\sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} f'_e(\theta_e(S^i))\theta_e(S^i)$ to $h_{i-1}(R)$. Therefore,

$$\begin{aligned}\Delta_i(R) \quad \geq \quad & \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} g_e(S^{i-1}) \\ & -\frac{1}{\alpha}\sum_{e\in s^i_{\pi(i)}\cap E^*} w_{\pi(i)} f'_e(\theta_e(S^i))\theta_e(S^i).\end{aligned}$$

Finally, since $g_e(S^{i-1}) \geq f_e(\theta_e(S^{i-1})) - f_e(\alpha\beta\theta_e(S^*))$ for every $e \in E^*$, it follows that

$$\begin{aligned}\Delta_i(R) \quad \geq \quad & \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(f_e(\theta_e(S^{i-1})) - f_e(\alpha\beta\theta_e(S^*))\right) \\ & -\frac{1}{\alpha}\sum_{e\in s^i_{\pi(i)}\cap E^*} w_{\pi(i)} f'_e(\theta_e(S^i))\theta_e(S^i).\end{aligned} \tag{7}$$

Since $f_e(x) = x^{d_e}$ with $d_e \geq 0$, we obtain

$$\begin{aligned}\Delta_i(R) \quad \geq \quad & \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e^{d_e}(S^{i-1}) - (\alpha\beta)^{d_e}\theta_e^{d_e}(S^*)\right) \\ & -\frac{1}{\alpha}\sum_{e\in s^i_{\pi(i)}\cap E^*} d_e w_{\pi(i)}\theta_e^{d_e}(S^i) \\ \geq \quad & \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e^{d_e}(S^{i-1}) - (\alpha\beta)^{d_e}\theta_e^{d_e}(S^*)\right) \\ & -\frac{d}{\alpha} c_{\pi(i)}(S^i).\end{aligned}$$

Since players perform best responses,

$$c_{\pi(i)}(S^i) \leq \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e(S^{i-1}) + w_{\pi(i)}\right)^{d_e}; \tag{8}$$

moreover, $w_{\pi(i)} \leq \theta_e(S^*)$ for every $e \in s^*_{\pi(i)}$; thus, by using Claim 1 in (9), it follows that

$$\begin{aligned}\Delta_i(R) \quad \geq \quad & \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e^{d_e}(S^{i-1}) - (\alpha\beta)^{d_e}\theta_e^{d_e}(S^*)\right) \\ & -\frac{d}{\alpha}\sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e(S^{i-1}) + w_{\pi(i)}\right)^{d_e}\end{aligned}$$

$$\begin{aligned}
&\geq \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \Bigg(\frac{1}{2^{d_e-1}} \left(\theta_e(S^{i-1}) + w_{\pi(i)}\right)^d \\
&\quad -w^{d_e}_{\pi(i)} - (\alpha\beta)^{d_e}\theta^{d_e}_e(S^*)\Bigg) \\
&\quad -\frac{d}{\alpha} \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e(S^{i-1}) + w_{\pi(i)}\right)^{d_e} \qquad (9)\\
&\geq \left(\frac{1}{2^{d-1}} - \frac{d}{\alpha}\right) \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e(S^{i-1}) + w_{\pi(i)}\right)^{d_e} \\
&\quad -\left((\alpha\beta)^d + 1\right) \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)}\theta^{d_e}_e(S^*).
\end{aligned}$$

By summing up the values $\Delta_i(R)$, we obtain

$$\begin{aligned}
&\sum_{i=1}^{\ell} \Delta_i(R) \\
&\geq \left(\frac{1}{2^{d-1}} - \frac{d}{\alpha}\right) \sum_{i=1}^{\ell} \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)} \left(\theta_e(S^{i-1}) + w_{\pi(i)}\right)^{d_e} \\
&\quad -\left((\alpha\beta)^d + 1\right) \sum_{i=1}^{\ell} \sum_{e\in s^*_{\pi(i)}} w_{\pi(i)}\theta^{d_e}_e(S^*) \\
&\geq \left(\frac{1}{2^{d-1}} - \frac{d}{\alpha}\right) \rho(R) - \beta\left((\alpha\beta)^d + 1\right) \textsc{Opt},
\end{aligned}$$

and thus, since $H(R) \geq \sum_{i=1}^{\ell} \Delta_i(R)$, the claim follows.

□

Theorem 2. *For any $k \geq 1$ and any given $d \geq 1$, in every polynomial congestion game $\mathcal{G}$ with delay functions having maximum degree d, it holds that the approximation ratio of the state reached after k (ℓ,β)-bounded covering walks is $O\left(W^{d\left(\frac{d}{d+1}\right)^{k-1}}\right)$.*

Proof. Let $P = \langle R_1, \ldots, R_k \rangle$ be the sequence of k (ℓ,β)-bounded covering walks. From Lemma 8 we obtain that for each R_j with $j = 2, \ldots, k$ it holds that for any $\alpha > d\, 2^{d-1}$

$$H(R_j) \geq \left(\frac{1}{2^{d-1}} - \frac{d}{\alpha}\right) \rho(R_j) - \beta\left((\alpha\beta)^d + 1\right) \textsc{Opt}. \qquad (10)$$

Furthermore, by applying Lemma 7 we obtain that for each R_j with $j = 2, \ldots, k$ it holds

$$H(R_j) \leq \beta(d+1)\left((d+1)\frac{\rho(R_{j-1})}{\textsc{Opt}}\right)^{\frac{d}{d+1}} \textsc{Opt}. \tag{11}$$

By combining (10) and (11) we achieve a relation between $\frac{\rho(R_j)}{\textsc{Opt}}$ and $\frac{\rho(R_{j-1})}{\textsc{Opt}}$ for every $j = 2, \ldots, k$

$$\begin{aligned}\frac{\rho(R_j)}{\textsc{Opt}} &\leq \beta\left(\frac{\alpha 2^{d-1}}{\alpha - d\, 2^{d-1}}\right)\cdot \\ &\cdot\left((d+1)\left((d+1)\frac{\rho(R_{j-1})}{\textsc{Opt}}\right)^{\frac{d}{d+1}} + \right. \\ &\left. + \left((\alpha\beta)^d + 1\right)\right)\end{aligned} \tag{12}$$

for any $\alpha > d\, 2^{d-1}$.
From the previous inequalities (12) (for $j = 2, \ldots, k$), since $\beta = O(1)$, we obtain that for constant values of α and d it holds

$$\frac{\rho(R_k)}{\textsc{Opt}} = O\left(\left(\frac{\rho(R_1)}{\textsc{Opt}}\right)^{\left(\frac{d}{d+1}\right)^{k-1}}\right). \tag{13}$$

By applying Lemma 5 to $\rho(R_k)$ and Lemma 6 to $\rho(R_1)$ in (13), since $\beta = O(1)$, by constant value of d we obtain that the cost of the final state of walk P is

$$O\left(\left((W + w_{max})^d\right)^{\left(\frac{d}{d+1}\right)^{k-1}} \textsc{Opt}\right),$$

and thus the approximation is

$$\mathrm{Apx}_k^{O(1)}(\mathcal{G}) = O\left(W^{d\left(\frac{d}{d+1}\right)^{k-1}}\right).$$

□

The following corollary is an immediate consequence of Theorem 2.

Corollary 1. *For any polynomial congestion game $\mathcal{G}$ with $d = O(1)$, a best response dynamics satisfying the (T,β)-Fairness Condition converges from any initial state to a state having approximation ratio $O(1)$ in at most $\log\log W$ best responses.*

Now we provide an almost matching lower bound to the approximation ratio achieved after a (ℓ,β)-bounded k-covering walk.

Theorem 3. *For any $d \geq 1$ there exists a congestion game $\mathcal{G}$ with polynomial delay functions having maximum degree d such that the approximation ratio achieved after k (ℓ,β)-bounded covering walks is* $\Omega\left(\frac{W^{d\left(\frac{d}{d+1}\right)^{k-1}}}{k}\right)$.

Proof. Given any integers k and c, we consider a congestion game $\mathcal{G}$ defined on a set of $n = m(k+1)$ players having weight equal to 1 (i.e. $W = n$) and $m(k+2)$ resources, with $m = c^{(d+1)^k}$. The set of players can be partitioned into $k+1$ sets $P_0, \ldots, P_k$, each containing m players, i.e. $P_j = \{p_j^1, \ldots, p_j^m\}$ for $j = 0, \ldots, k$; the set of resources can be partitioned into $k+2$ sets $T_0, \ldots, T_{k+1}$ each containing m resources, i.e. $T_j = \{e_j^1, \ldots, e_j^m\}$ for $j = 0, \ldots, k+1$, and such that each resource has delay function $f(x) = x^d$.

The strategy set of player p_j^i (the i-th player of set P_j) consists of two strategies: the *left* strategy and the *right* strategy. The left strategy for player p_j^i of set P_j $(j = 0, \ldots, k)$ consists of the only resource e_j^i belonging to set T_j.

The right strategy of each player in P_0 is T_1; in order to define the right strategies of the remaining players, we need some additional definitions.

Let $b_{i,j} = \left\lfloor \frac{i}{m^{\left(\frac{d}{d+1}\right)^j}} \right\rfloor$; for each $j = 2, \ldots, k$, we partition the set T_j into $\frac{m}{m^{\left(\frac{d}{d+1}\right)^{j-1}}}$ (by the definition of m it is an integer) *blocks* $T_j^0, \ldots, T_j^{\frac{m}{m^{\left(\frac{d}{d+1}\right)^{j-1}}}-1}$. The right strategy for player p_j^i of set P_j $(j = 1, \ldots, k-1)$ is $T_{j+1}^{b_{i,j}}$. Notice that, by the definition of $b_{i,j}$, the number of players in P_j having as their right strategy block T_{j+1}^h for a fixed h is $m^{\left(\frac{d}{d+1}\right)^j}$, and for any resource in such a block, the number of players whose right strategy contains it is $m^{\left(\frac{d}{d+1}\right)^j}$.

Let us compute the social cost of the configuration obtained after k walks starting from the state in which all the players select their right strategy.

We assume that, during all the walks, the players of set P_j perform their best response before the ones of set P_{j-1} for $j = 1, \ldots, k$, and for every $i = 0, \ldots, k$ players $p_j^i \in P_j$ perform their best response in increasing order with respect to i.

Let us consider walk R_j, $j = 1, \ldots, k$. We now show by induction on the number of walks that the players belonging to P_x, $x = j, \ldots, k$, choose their right strategy, while the ones belonging to P_x, $x = 0, \ldots, j-1$, choose their left strategy.

Assume by induction that this is true until R_{j-1}; we now show that it holds also for R_j.

- The players belonging to P_x, $x = j, \ldots, k$, do not change their (right) strategy. In fact, on the one hand, by the inductive hypothesis, since each resource in their right strategies is used by $m^{\left(\frac{d}{d+1}\right)^x}$ players, their right strategy has a cost $m^{\left(\frac{d}{d+1}\right)^x} \cdot \left(m^{\left(\frac{d}{d+1}\right)^x}\right)^d = m^{\frac{d^x}{(d+1)^{x-1}}}$. On the other hand, always by the inductive hypothesis, their left strategies (composed by a unique resource) is used by $m^{\left(\frac{d}{d+1}\right)^{x-1}}$ players and therefore has a cost equal to $\left(m^{\left(\frac{d}{d+1}\right)^{x-1}}\right)^d = m^{\frac{d^x}{(d+1)^{x-1}}}$.

- The players belonging to P_{j-1} select their left strategies. In fact, on the one hand, since their right strategy is composed by $m^{\left(\frac{d}{d+1}\right)^{j-1}}$ resources, it has a cost at least $m^{\left(\frac{d}{d+1}\right)^{j-1}}$. On the other hand, by the inductive hypothesis, their left strategies (composed by a unique resource) is free, and therefore has a cost equal to 1.

- The players belonging to P_x, $x = 0, \ldots, j-2$, do not change their (left) strategy where they have a delay equal to 1 (i.e. the minimum possible one in such an instance).

Thus, in the final state of walk R_k, all the players in P_k are using their right strategy consisting of $m^{\left(\frac{d}{d+1}\right)^k}$ resources in T_{k+1}. Moreover, each of such resources is used by $m^{\left(\frac{d}{d+1}\right)^k}$ players, and it follows that the social cost of the final state of walk R_k is at least $m \cdot \left(m^{\left(\frac{d}{d+1}\right)^k}\right)^d \cdot m^{\left(\frac{d}{d+1}\right)^k} = m \cdot m^{d\left(\frac{d}{d+1}\right)^{k-1}}$. Since the configuration in which each player uses her left strategy costs n,

$$\mathrm{Apx}_k^1(\mathcal{G}) \geq \frac{m \cdot m^{d\left(\frac{d}{d+1}\right)^{k-1}}}{n} = \Omega\left(\frac{W^{d\left(\frac{d}{d+1}\right)^{k-1}}}{k}\right).$$

□

5.2 General β [15]

In this subsection, we focus on the special case of congestion games with linear delays, and completely characterize how the fairness of the dynamics affects the time of convergence to good solutions, approximating an optimal one by a constant factor.

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

Consider a generic (ℓ,β)-bounded covering $R = \left(S^0, \ldots, S^\ell\right)$. 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* $c_{\pi(t)}(S^t)$ can be suitably upper bounded as $\sum_{e \in s^*_{\pi(t)}} \left(\theta_e(S^{t-1}) + w_{\pi(t)}\right)$.

By extending and strengthening the technique of [13, 14], 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 [13, 14], only a 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 [13, 14], have to be redefined: roughly speaking, they now must take into account only the last move in R of each player, whereas in [13, 14] they were accounting for all the moves in R.

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

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

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, it holds that, for any $i = 1, \ldots, n$,

$$c_i(S^{\text{last}_R(i)}) \leq \sum_{e \in s_i^*} \left(\theta_e(S^{\text{last}_R(i)-1}) + w_i\right) \tag{14}$$

and, any summing over all players, we obtain that $\sum_{i=1}^{n} c_i(S^{\text{last}_R(i)}) \leq \rho(R)$, i.e. $\rho(R)$ is an upper bound to the sum of the immediate costs over the last moves of every players. Finally, it is worth noticing that, by inverting the order of the summations, $H(R) = \sum_{e \in E} \theta_e(S^0)\theta_e(S^*)$.

The upper bound proof is structured as follows. Lemma 9 relates the social cost of the final state S^ℓ of a (ℓ,β)-bounded covering R with $\rho(R)$, by showing that $C(S^\ell) \leq 2\rho(R)$. Let $\overline{R}$ and R be two consecutive (ℓ,β)-bounded coverings; by exploiting Lemmata 10 and 11, providing an upper (lower, respectively) bound to $H(R)$ in terms of $\rho(\overline{R})$ ($\rho(R)$, respectively), Lemma 12 proves that $\frac{\rho}{\textsc{Opt}}$ rapidly decreases between $\overline{R}$ and R, showing that $\frac{\rho(R)}{\textsc{Opt}} = O\left(\sqrt{\frac{\rho(\overline{R})}{\textsc{Opt}}} + \beta\right)$. In the proof of Theorem 1, after deriving a trivial upper bound equal to $O(W)$ for $\rho(R_1)$, Lemma 12 is applied to all the $k-1$ couples of consecutive (ℓ,β)-bounded coverings of the considered dynamics satisfying the (T,β)-Fairness Condition.

Similarly to Lemma 5 and Lemma 7, the following lemmata show that the social cost at the end of any (ℓ,β)-bounded covering R is at most $2\rho(R)$, and that $\frac{H(R')}{\textsc{Opt}}$ is significantly less than $\frac{\rho(R)}{\textsc{Opt}}$ for two consecutive coverings R and R'.

Lemma 9. *For any $\beta \geq 1$, given a (ℓ,β)-bounded covering R, $C(S^\ell) \leq 2\rho(R)$.*

Lemma 10. *For any $\beta \geq 1$, given two consecutive (ℓ,β)-bounded covering walks R and R' such that the final state of R coincides with the initial one of R', it holds,* $\frac{H(R')}{\textsc{Opt}} \leq \sqrt{2\frac{\rho(R)}{\textsc{Opt}}}$.

In Lemma 11 we are able to relate $\rho(R)$ and $H(R)$ by strengthening the technique exploited in [13, 14].

Lemma 11. *For any $\beta \geq 1$, given a (ℓ,β)-bounded covering R,* $\frac{\rho(R)}{\textsc{Opt}} \leq 2\frac{H(R)}{\textsc{Opt}}+4\beta+3$.

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 $\theta_e(S^*) = 0$, then, by recalling the definitions of $\rho(R)$ and $H(R)$, we would obtain

$$\begin{aligned}
\rho(R) &= \sum_{i=1}^{n} w_i \sum_{e\in s_i^*} \left(\theta_e(S^{\mathrm{last}_R(i)-1}) + w_i\right)\\
&= \sum_{i=1}^{n} w_i \sum_{e\in s_i^*} \left(\theta_e(S^0) + w_i\right)\\
&\leq \sum_{i=1}^{n} w_i \left(\sum_{e\in s_i^*} \theta_e(S^0) + \sum_{e\in s_i^*} w_i\right)\\
&= \sum_{e\in E} \theta_e(S^0)\theta_e(S^*) + \sum_{i=1}^{n}\sum_{e\in s_i^*} w_i^2\\
&\leq \sum_{e\in E} \theta_e(S^0)\theta_e(S^*) + \sum_{e\in E} \theta_e^2(S^*)\\
&= H(R) + \textsc{Opt},
\end{aligned}$$

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 congestion on resources e such that $\theta_e(S^*) > 0$.

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

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

$$\Delta_0 = \sum_{e\in E} \max\left\{\theta_e(S^0) + \theta_e(S^*), 2\beta\theta_e(S^*)\right\} \cdot \theta_e(S^*)$$

$$
\begin{aligned}
&\le \sum_{e\in E}\Big(\theta_e(S^0)+\theta_e(S^*)\Big)\cdot\theta_e(S^*)+2\beta\sum_{e\in E}\theta_e^2(S^*) \qquad (15)\\
&= H(R)+(2\beta+1)\textsc{Opt}.
\end{aligned}
$$

The idea is that the total delay of level 0 can induce on the resources a congestion (over the frontier g_e) contributing to the total delay of level 1, such a delay a congestion contributing (always over the frontier g_e) to the total delay 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 delay of level p on resource e; we say that a delay $\delta_{p,e}$ of level p on resource e is *induced* by an amount $x_{p-1,e}$ of delay of level $p-1$ if some players (say, players in $N_{p-1,e}$) moving on e can cause such a delay of level p on e because they are experimenting a delay of level $p-1$ on the resources of their optimal strategies equal to $x_{p-1,e}$. In other words, $x_{p-1,e}$ is the overall delay of level $p-1$ on the resources in the optimal strategies of players in $N_{p-1,e}$ used in order to induce the delay $\delta_{p,e}$ of level p on resource e.

For every $i = 1, \dots, n$, since players perform best responses, player i, in order to select a strategy s_i, must suffer a delay when selecting her optimal strategy s_i^* at least equal to the cost of the selected strategy s_i. In the following, we will assume that a player i can perform a best response selecting a strategy s_i if her cost for strategy s_i^* computed according to the delays $\delta_{p,e}$ ($e \in s_i^*$, $p \geq 0$) is at least her cost for strategy i. Such a delay is initially induced, for every $e \in s_i^*$, by the congestion $\theta_e(S^0)$; the additive term $\theta_e(S^*)$ in the definition of $\delta_{0,e}$ is due to the fact that, before performing her best response, player i weight might not belong to $\theta_e(S^0)$, while w_i has to be taken into account when computing the delay suffered by i when selecting her optimal strategy s_i^*.

We have to clarify how a total delay belonging to various levels is exploited in order to induce the delay of a given level, say level p, on a resource, say resource e. In fact, only at the beginning we can assume that all the delay is of level 0 and therefore it is entirely used in order to induce a delay of level 1, while at a generic move the delay on a resource can belong to different levels. Consider a best response s_i of a generic player i, with $e \in s_i$; if D is the total amount of delay needed in order to select resource e (delay D is belonging to some resources of s_i^*), we assume that the amount of the induced delay of a given level p on resource e is proportional to the amount of delay of level $p-1$ contained in D.

Assume that resource e has a congestion $\alpha \geq g_e$ and that player i is selecting strategy s_i with $e \in s_i$. The delay of e increases of w_i, and in order to perform such a best response, player i has to suffer a cost equal to $D = w_i(\alpha + w_i)$ on her optimal strategy. We have that $w_i(\alpha + w_i) \geq \alpha w_i$, i.e., the delay on s_i^* used in order to increase the delay on e is at least α times the amount of the increase.

If $\alpha < g_e$ and $\alpha + w_i > g_e$, the delay of e increases of ϵ, with $\epsilon = \alpha + w_i - g_e < w_i$

and in order to perform such a best response, player i has to suffer a cost equal to $D = w_i(g_e + \epsilon) > \epsilon(g_e + \epsilon)$ on her optimal strategy. Therefore, we again obtain that the delay on s_i^* used in order to increase the delay on e is at least α times the amount of the increase.

Since we are assuming that the amount of the induced delay of a given level p on resource e is proportional to the amount of delay of level $p-1$ contained in D, we obtain that

$$\delta_{p,e} \leq \frac{x_{p-1,e}}{\alpha} \leq \frac{x_{p-1,e}}{g_e}. \tag{16}$$

For any p, the total delay of level p is defined as $\Delta_p = \sum_{e \in E} \delta_{p,e} \cdot \theta_e(S^*)$. Moreover, for any $p \geq 1$, we have that

$$\sum_{e \in E} x_{p-1,e} \leq \beta \Delta_{p-1} \tag{17}$$

because each player can move at most β times in R and therefore the total delay of level $p-1$ can be used at most β times in order to induce the total delay of level p.

We also have 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 delay of resource e during the whole covering R.

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

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

where the first inequality holds by inequality (16), and the last inequality holds by inequality (17).

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

$$\rho(R) \leq \sum_{p=0}^{\infty} \Delta_p + \textsc{Opt} \leq 2\Delta_0 + \textsc{Opt}.$$

Finally, the claim follows by combining this upper bounds to $\rho(R)$ with the upper bound to Δ_0 derived in (15).

□

By combining Lemmata 10 and 11, the following lemma, showing that $\frac{\rho(\cdot)}{\textsc{Opt}}$ fast decreases between two consecutive coverings, holds.

Lemma 12. *For any $\beta \geq 1$, given two consecutive (ℓ, β)-bounded coverings $\overline{R}$ and R,* $\frac{\rho(R)}{\textsc{Opt}} \leq 2\sqrt{2\frac{\rho(\overline{R})}{\textsc{Opt}}} + 4\beta + 3$.

By applying Lemma 12 to all the couples of consecutive (ℓ,β)-bounded coverings, we are now able to prove the following theorem.

Theorem 4. *Given a weighted congestion game with linear delay functions, 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.*

Proof. Given a best response dynamics satisfying the (T,β)-Fairness Condition, let $R_1,\ldots,R_k$ be the k (ℓ,β)-bounded coverings in which it can be decomposed. By applying Lemma 12 to all the pairs of consecutive (ℓ,β)-bounded coverings R_j and R_{j+1}, for any $j = 1,\ldots,k-1$ we obtain

$$\frac{\rho(R_{j+1})}{\textsc{Opt}} \leq 2\sqrt{2\frac{\rho(R_j)}{\textsc{Opt}}} + 4\beta + 3.$$

By combining all the above inequalities for $j = 1,\ldots,k-1$ and by performing some basic algebraic manipulations, for any constant value of d we obtain that $\frac{\rho(R_k)}{\textsc{Opt}} = O\left(\sqrt[2^{k-1}]{\frac{\rho(R_1)}{\textsc{Opt}}} + \beta\right)$. Thus, by Lemma 9, the cost of the final state S of walk R_k is such that

$$\frac{C(S)}{\textsc{Opt}} = O\left(\sqrt[2^{k-1}]{\frac{\rho(R_1)}{\textsc{Opt}}} + \beta\right).$$

By the definition of $\rho(R)$, since $\sum_{e\in E}\theta_e(S^*) \leq \sum_{e\in E}\theta_e^2(S^*) = \textsc{Opt}$, for any possible (ℓ,β)-bounded covering R it holds that

$$\begin{aligned}
\rho(R) &= \sum_{i=1}^{n} w_i \sum_{e\in s_i^*}\left(\theta_e(S^{\text{last}(i)-1}) + w_i\right) \leq \sum_{i=1}^{n} w_i \sum_{e\in s_i^*}(W+W) \\
&= 2W\sum_{i=1}^{n} w_i|s_i^*| \leq 2W\sum_{e\in E}\theta_e(S^*) \leq 2W\textsc{Opt}.
\end{aligned}$$

Therefore, $\frac{\rho(R_1)}{\textsc{Opt}} \leq 2W$ and we obtain $\frac{C(S)}{\textsc{Opt}} = O\left(\sqrt[2^{k-1}]{W} + \beta\right)$.

It is worth noticing that $\log\log W$ (ℓ,β)-bounded coverings are sufficient in order to obtain $\frac{C(S)}{\textsc{Opt}} = O(\beta)$. Since every (ℓ,β)-bounded covering, by its definition, contains at most T best responses, the claim follows. □

It is also possible to prove the following lower bounds.

Theorem 5. *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 $\Omega(\beta^{1-\epsilon}\cdot\textsc{Opt})$.*

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 2. *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 $\Omega\left(\frac{\sqrt[4]{n}}{\log n} \cdot \textsc{Opt}\right)$.*

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 $T\lceil \log\log W \rceil$ best responses an efficient state is always reached.

Theorem 6. *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.*

6 Conclusions and Future Work

In this work we have surveyed the state of the art about convergence issues in congestion games, with a special focus on the results concerning the speed of convergence of best response dynamics. In particular, we have shown that in congestion games with polynomial delays *fair* dynamics, in which each player is allowed to player at least once and at most a constant number of times every T best responses, fast converges to solutions approximating the optimum by a factor proportion to the price of anarchy.

Moreover, we have completely characterized how, in weighted congestion games with linear delays, the frequency with which each player participates in the game dynamics affects the possibility of reaching states with an approximation ratio within a constant factor from the price of anarchy, within a polynomially bounded number of best responses. We have shown that, while in the asymmetric setting the fairness among players is a necessary and sufficient condition for guaranteeing a fast convergence to efficient states, in the symmetric one the game always converges to an efficient state after a polynomial number of best responses, regardless of the frequency each player moves with. We conjecture that such results can be extended to broader classes of congestion games, such as congestion games with polynomial delay functions.

It is worth to note that our techniques provide a much faster convergence to efficient states with respect to previous results in the literature. In particular, in the symmetric setting, Theorem 6 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 [8]. Furthermore, also in the more general asymmetric setting, Theorem 4 shows that the same holds for fair best response dynamics with respect to ϵ-Nash ones [5].

An interesting open question is that of studying the time of convergence to solutions approximating by a low factor the optimum with respect to other social function, such as the maximum cost among the players. Finally, considering other kinds of dynamics (such as coalitional responses in which two or more players coordinate in order to decrease their costs) is a left open problem that deserves further research effort.

References

[1] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *J. ACM*, 55(6), 2008.

[2] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. *Theor. Comput. Sci.*, 410(17):1552–1563, 2009.

[3] Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. Exact price of anarchy for polynomial congestion games. In *STACS*, volume 3884, pages 218–229. Springer, 2006.

[4] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *STOC*, pages 57–66. ACM, 2005.

[5] Baruch Awerbuch, Yossi Azar, Amir Epstein, Vahab S. Mirrokni, and Alexander Skopalik. Fast convergence to nearly optimal solutions in potential games. In *ACM Conference on Electronic Commerce*, pages 264–273. ACM, 2008.

[6] Vittorio Bilò, Angelo Fanelli, Michele Flammini, and Luca Moscardelli. Performance of one-round walks in linear congestion games. *Theory Comput. Syst.*, 49(1):24–45, 2011.

[7] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. In *ICALP (1)*, volume 4051, pages 311–322. Springer, 2006.

[8] S. Chien and A. Sinclair. Convergence to approximate Nash equilibria in congestion games. In *SODA*, pages 169–178. SIAM, 2007.

[9] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC*, pages 67–73. ACM, 2005.

[10] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *STACS*, volume 3884, pages 349–360. Springer, 2006.

[11] M. Yannakakis D. S. Johnson, C. H. Papadimitriou. How easy is local search? *Journal of Computer and System Sciences*, 37:79–100, 1988.

[12] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *STOC*, pages 604–612. ACM, 2004.

[13] Angelo Fanelli, Michele Flammini, and Luca Moscardelli. The speed of convergence in congestion games under best-response dynamics. In *ICALP (1)*, volume 5125, pages 796–807. Springer, 2008.

[14] Angelo Fanelli and Luca Moscardelli. On best response dynamics in weighted congestion games with polynomial delays. *Distributed Computing*, 24:245–254, 2011.

[15] Angelo Fanelli, Luca Moscardelli, and Alexander Skopalik. On the impact of fair best response dynamics. In *MFCS*, volume 7464 of *LNCS*, pages 360–371. Springer, 2012.

[16] Babak Farzad, Neil Olver, and Adrian Vetta. A priority-based model of routing. *Chicago J. Theor. Comput. Sci.*, 2008, 2008.

[17] Dimitris Fotakis, Spyros C. Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul G. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *ICALP*, volume 2380, pages 123–134. Springer, 2002.

[18] Dimitris Fotakis, Spyros C. Kontogiannis, and Paul G. Spirakis. Atomic congestion games among coalitions. *ACM Transactions on Algorithms*, 4(4), 2008.

[19] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *STACS*, volume 1563, pages 404–413. Springer, 1999.

[20] Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111 – 124, 1996.

[21] Vahab S. Mirrokni and Adrian Vetta. Convergence issues in competitive games. In *APPROX-RANDOM*, volume 3122 of *LNCS*, pages 183–194. Springer, 2004.

[22] J. F. Nash. Equilibrium points in n-person games. In *Proceedings of the National Academy of Sciences*, volume 36, pages 48–49, 1950.

[23] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[24] Alexander Skopalik and Berthold Vöcking. Inapproximability of pure Nash equilibria. In *STOC*, pages 355–364. ACM, 2008.