On the Impact of Singleton Strategies in Congestion Games

Vittorio Bilò[1] and Cosimo Vinci[2]

1 **Department of Mathematics and Physics, University of Salento, Lecce, Italy**
`vittorio.bilo@unisalento.it`

2 **Gran Sasso Science Institute, L'Aquila, Italy**
`cosimo.vinci@gssi.it`

Abstract

To what extent does the structure of the players' strategy space influence the efficiency of decentralized solutions in congestion games? In this work, we investigate whether better performance is possible when restricting to load balancing games in which players can only choose among single resources. We consider three different solutions concepts, namely, approximate pure Nash equilibria, approximate one-round walks generated by selfish players aiming at minimizing their personal cost and approximate one-round walks generated by cooperative players aiming at minimizing the marginal increase in the sum of the players' personal costs. The last two concepts can also be interpreted as solutions of simple greedy online algorithms for the related resource selection problem. Under fairly general latency functions on the resources, we show that, for all three types of solutions, better bounds cannot be achieved if players are either weighted or asymmetric. On the positive side, we prove that, under mild assumptions on the latency functions, improvements on the performance of approximate pure Nash equilibria are possible for load balancing games with weighted and symmetric players in the case of identical resources. We also design lower bounds on the performance of one-round walks in load balancing games with unweighted players and identical resources (in this case, solutions generated by selfish and cooperative players coincide).

1998 ACM Subject Classification F.7.2 Algorithmic Game Theory and Mechanism Design, F.7.2.6 Quality of Equilibria, F.5.5 Online Algorithms

Keywords and phrases Congestion games, Nash equilibrium, price of anarchy, online load balancing, greedy algorithms

Digital Object Identifier 10.4230/LIPIcs.ESA.2017.17

1 Introduction

Congestion games [25] are non-cooperative games in which there is a set of selfish players competing for a set of resources, and each resource incurs a certain latency, expressed by a congestion-dependent function, to the players using it. Each player has a certain weight and an available set of strategies, where each strategy is a non-empty subset of resources, and aims at choosing a strategy minimizing her personal cost which is defined as the sum of the latencies experienced on all the selected resources. We speak of *weighted* games/players when players have arbitrary non-negative weights and of *unweighted* games/players when all players have unitary weight.

Stable outcomes in this setting are the pure Nash equilibria [24]: strategy profiles in which no player can lower her cost by unilaterally deviating to another strategy. However, they are demanding solution concepts, as they might not always exist in weighted games

25th Annual European Symposium on Algorithms (ESA 2017).
Editors: Kirk Pruhs and Christian Sohler; Article No. 17; pp. 17:1–17:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

[18] and, even when their existence is guaranteed, as, for instance, in unweighted games [25] and in weighted games with affine latency functions [18, 21], their computation might be an intractable problem [1, 16]. For such a reason, more relaxed solution concepts are usually considered in the literature, as ϵ-approximate pure Nash equilibria or ϵ-approximate one-round walks. An ϵ-approximate pure Nash equilibrium is the relaxation of the concept of pure Nash equilibrium in which no player can lower her cost of a factor more than $1 + \epsilon$ by unilaterally deviating to another strategy, while an ϵ-approximate one-round walk is defined as a myopic process in which players arrive in an arbitrary order and, upon arrival, each of them has to make an irrevocably strategic choice aiming at approximatively minimizing a certain cost function. In this work, we shall consider two variants of this process: in the first, players choose a strategy approximatively minimizing, up to a factor of $1 + \epsilon$, their personal cost (*selfish players*), while, in the second, players choose the strategy approximatively minimizing, up to a factor of $1 + \epsilon$, the marginal increase in the social cost (*cooperative players*) which is defined as the sum of the players' personal costs (for the case of $\epsilon = 0$, we use the term *exact one-round walk*). In particular, approximate one-round walks can be interpreted as simple greedy online algorithms for the equivalent resource selection problem associated with a given congestion game, and, in most of the cases, these algorithms are optimal in the context of online optimization of load balancing problems [9]. The worst-case efficiency of these solution concepts with respect to the optimal social cost is termed as the *ϵ-approximate price of anarchy* (for the case of pure Nash equilibria, the term *price of anarchy* [22] is adopted) and as the *competitive ratio* of ϵ-approximate one-round walks, respectively. Interesting special cases of congestion games are obtained by restricting the combinatorics of the players' strategy space. In *symmetric congestion games*, all players share the same set of strategies; in *network congestion games* the players' strategies are defined as paths in a given network; in *matroid congestion games* [1, 2], the strategy set of every player is given by the set of bases of a matroid defined over the set of available resources; in *k-uniform matroid congestion games* [15], each player can select any subset of cardinality k from a prescribed player-specific set of resources; finally, in *load balancing games*, players can only choose single resources.

To what extent does the structure of the players' strategy space influence the efficiency of decentralized solutions in congestion games? In this work, we investigate whether better performance is possible when restricting to load balancing games. Previous work established that the price of anarchy does not improve when restricting to unweighted load balancing games with polynomial latency functions [10, 20], while better bounds are possible in unweighted symmetric load balancing games with fairly general latency functions [17]. Under the assumption of identical resources with affine latency functions, improvements are also possible when restricting to both unweighted load balancing games [10, 27] and weighted symmetric load balancing games [23]. Finally, [6] proves that the price of anarchy does not improve when restricting to weighted symmetric load balancing games under polynomial latency functions. For the competitive ratio of exact one-round walks generated by cooperative players, no improvements are possible in unweighted load balancing games with affine latency functions [10, 27], while improved performance can be obtained under the additional assumption of identical resources [10] (we observe that, in this case, solutions generated by both types of players coincide); however, for weighted players, no improvements are possible even under the assumption of identical resources [9, 10]. For one-round walks generated by selfish players, instead, no specialized limitations are currently known.

Our Contribution. We obtain an almost precise picture of the cases in which improved performance can be obtained in load balancing congestion games. This is done by either solving open problems or extending previously known results to both approximate solution concepts and more general latency functions. Specifically, we provide the following characterizations.

Let $\mathcal{C}$ be a class of non-negative and non-decreasing functions such that, for each $f \in \mathcal{C}$ and $\alpha \in \mathbb{R}_{\geq 0}$, the function g such that $g(x) = \alpha f(x)$ belongs to $\mathcal{C}$ and let $\mathcal{C}' \subset \mathcal{C}$ be the subclass of $\mathcal{C}$ such that, for each $f \in \mathcal{C}'$ and $\alpha \in \mathbb{R}_{\geq 0}$, the function h such that $h(x) = f(\alpha x)$ belongs to $\mathcal{C}'$. A function f is *semi-convex* if $xf(x)$ is convex, it is *unbounded* if $\lim_{x\to\infty} f(x) = \infty$. We prove that:

- **for weighted players:** under unbounded latency functions drawn from $\mathcal{C}'$, the approximate price of anarchy does not improve when restricting to symmetric load balancing games (this solves an open problem raised in [6], where a similar limitation was shown only with respect to pure Nash equilibria and polynomial latency functions). Under latency functions drawn from $\mathcal{C}'$, the competitive ratio of approximate one-round walks generated by selfish players does not improve when restricting to load balancing games (this solves an open problem raised in [8]). If all functions in $\mathcal{C}'$ are semi-convex, then the same limitation applies to the competitive ratio of approximate one-round walks generated by cooperative players (this generalizes results in [5, 9, 10] which hold only with respect to exact one-round walks for games with polynomial latency functions). We also provide a parametric formula for the relative bounds which we use to obtain the exact values for polynomial latency functions;
- **for unweighted players:** under latency functions drawn from $\mathcal{C}$, either the approximate price of anarchy and the competitive ratio of approximate one-round walks generated by both selfish and cooperative players do not improve when restricting to load balancing games (these generalize a result in [10, 20] which holds only with respect to pure Nash equilibria and polynomial latency functions, a result in [10, 27] which holds only with respect to exact one-round walks generated by cooperative players in games with affine latency functions, and solve an open problem raised in [8] for one-round walks generated by selfish players). Also in this case we provide a parametric formula for the relative bounds which we use to obtain the exact bounds for polynomial latency functions.

These negative results, together with the positive ones achieved by [10, 17], imply that better bounds on the approximate price of anarchy are possible only when dealing with unweighted symmetric load balancing games. However, under the additional hypothesis of identical resources, better performance is still possible. Let f be an increasing, continuous and semi-convex function. We prove that the approximate price of anarchy of weighted symmetric load balancing games with identical resources whose latency functions coincide with f is equal to $\sup_{x\in\mathbb{R}_{>0}} \sup_{\lambda\in(0,1)} \left\{ \frac{\lambda x f(x)+(1-\lambda)inv(x)f(inv(x))}{opt(x)f(opt(x))} \right\}$, where $inv(x) := \inf\{t \geq 0 : f(x) \leq (1+\epsilon)f(x/2+t)\}$ and $opt(x) := \lambda x + (1-\lambda)inv(x)$. This generalizes a result by [23] which holds only with respect to the price of anarchy under affine latency functions. Furthermore, by using the previous formula, we compute the exact price of anarchy of weighted symmetric load balancing games with identical resources and polynomial latency functions.

Finally, still for the case of identical resources, we design lower bounds on the performance of exact one-round walks in load balancing games with unweighted players (this improves and generalizes a result in [10] which holds only for affine latency functions).

Related Work. The price of anarchy in congestion games was first considered in [4] and [11] where it was independently shown that the price of anarchy is $5/2$ and $(3+\sqrt{5})/2$ for, respectively, unweighted and weighted congestion games with affine latency functions. In [11], it is also proved that no improved bounds are possible both in symmetric unweighted games and in unweighted network games; these results were improved by [14] which shows that the price of anarchy stays the same even in symmetric unweighted network games. In [10], it is shown that the previous bounds are tight also for load balancing games. For the special case of load balancing games on identical resources, the works of [27] and [10] show that the price of anarchy is 2.012067 for unweighted games and at least 5/2 for weighted ones. In [23], it is proved that, for symmetric load balancing games, the price of anarchy drops to 4/3 if the games are unweighted, and to 9/8 if the games are weighted with identical resources. For symmetric unweighted k-uniform matroid congestion games with affine latency functions, [15] proves that the price of anarchy is at most 28/13 and at least 1.343 for a sufficiently large value of k (for $k = 5$, it is roughly 1.3428). Tight bounds on the price of anarchy of either weighted and unweighted congestion games with polynomial latency functions have been given by [3]. Under fairly general latency functions, [17] shows that the price of anarchy of unweighted symmetric load balancing games coincides with that of non-atomic congestion games (thus generalizing a first result by [19] which proves an upper bound of $\sum_{i\in[d]} \mathcal{B}_i$, where $\mathcal{B}_i$ is the ith Bell number for the case of polynomial latency functions of maximum degree equal to d), while [6] proves that assuming symmetric strategies does not lead to improved bounds in unweighted games and gives exact bounds for the case of weighted players. It also shows that, for the case of weighted players, no improvements are possible even in symmetric load balancing games with polynomial latency functions. Finally, [12] and [7] characterize the approximate price of anarchy, respectively, in unweighted and weighted games under affine latency functions.

The competitive ratio of exact one-round walks generated by cooperative players in load balancing games with polynomial latency functions has been first considered in [5], where, for the special case of affine functions, an upper bound of $3 + 2\sqrt{2}$ is provided for weighted players. For unweighted players, this result has been improved to 17/3 in [27], where it is also shown that, for identical resources, the upper bound drops to $2 + \sqrt{5}$ in spite of a lower bound of 3.0833. Finally, [10] shows matching lower bounds of $3 + 2\sqrt{2}$ and 17/3 for, respectively, weighted and unweighted players. For weighted games with polynomial latency functions, tight bounds have been given in [9]; the lower bounds, in particular, hold even for identical resources, thus improving previous results from [5]. In [10] it is also shown that, for unweighted players and identical resources, the competitive ratio lies between 4 and $\frac{2}{3}\sqrt{21} + 1$. For the case of selfish players and still under affine latency functions, [8, 13] show that the competitive ratio is $2 + \sqrt{5}$ for unweighted congestion games, while, for weighted players, [13] gives an upper bound of $4 + 2\sqrt{3}$. In this setting, no specialized results are known for restrictions to load balancing games.

2 Definitions and Notation

For two integers $0 \le k_1 \le k_2$, let $[k_1, k_2] := \{k_1, k_1 + 1, \ldots, k_2 - 1, k_2\}$ and $[k_1] := [1, k_1]$.

A *congestion game* is a tuple $\mathsf{CG} = (N, E, (\ell_e)_{e\in E}, (w_i)_{i\in N}, (\Sigma_i)_{i\in N})$, where N is a set of $n \ge 2$ players, E is a set of resources, $\ell_e : \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$ is the latency function of resource $e \in R$, and, for each $i \in N$, $w_i \ge 0$ is the weight of player i and $\Sigma_i \subseteq 2^R \setminus \emptyset$ is her set of strategies. We speak of *weighted* games/players when players have arbitrary weights and of *unweighted* games/players when $w_i = 1$ for each $i \in N$. A congestion game is *symmetric* if $\Sigma_i = \Sigma$ for each $i \in N$, i.e., if all players share the same strategy space. A *load balancing* game is a congestion game in which for each $i \in N$ and $S \in \Sigma_i$, $|S| = 1$, that is, all players' strategies are singleton sets. Given a class $\mathcal{C}$ of latency functions, let $\mathsf{W}(\mathcal{C})$ be the class of weighted congestion games, $\mathsf{U}(\mathcal{C})$ be the class of unweighted congestion games, $\mathsf{ULB}(\mathcal{C})$ be the class of unweighted load balancing games, $\mathsf{WLB}(\mathcal{C})$ be the class of weighted load balancing games, and $\mathsf{WSLB}(\mathcal{C})$ be the class of weighted symmetric load balancing games, all having latency functions in the class $\mathcal{C}$.

A *strategy profile* is an n-tuple of strategies $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$, that is, a state of the game in which each player $i \in N$ is adopting strategy $\sigma_i \in \Sigma_i$, so that $\Sigma := \times_{i\in N}\Sigma_i$ denotes the set of strategy profiles which can be realized in CG. For a strategy profile $\boldsymbol{\sigma}$, the *congestion* of resource $e \in E$ in $\boldsymbol{\sigma}$, denoted as $k_e(\boldsymbol{\sigma}) := \sum_{i\in N: e\in\sigma_i} w_i$, is the total weight of the players using resource e in $\boldsymbol{\sigma}$, (observe that, in unweighted games, $k_e(\boldsymbol{\sigma})$ coincides with the number of users of resource e in $\boldsymbol{\sigma}$). The personal cost of player i in $\boldsymbol{\sigma}$ is defined as $cost_i(\boldsymbol{\sigma}) = \sum_{e\in\sigma_i} \ell_e(k_e(\boldsymbol{\sigma}))$ and each player aims at minimizing it. For the sake of conciseness, when the strategy profile $\boldsymbol{\sigma}$ is clear from the context, we write k_e in place of $k_e(\boldsymbol{\sigma})$. Fix a strategy profile $\boldsymbol{\sigma}$ and a player $i \in N$. We denote with $\boldsymbol{\sigma}_{-i}$ the restriction of $\boldsymbol{\sigma}$ to all the players other than i; moreover, for a strategy $S \in \Sigma_i$, we denote with $(\boldsymbol{\sigma}_{-i}, S)$ the strategy profile obtained from $\boldsymbol{\sigma}$ when player i changes her strategy from σ_i to S, while the strategies of all the other players are kept fixed. The quality of a strategy profile in congestion games is measured by using the *social function* $\mathsf{SUM}(\boldsymbol{\sigma}) = \sum_{i\in N} w_i cost_i(\boldsymbol{\sigma}) = \sum_{e\in E} k_e(\boldsymbol{\sigma})\ell_e(k_e(\boldsymbol{\sigma}))$, that is, the sum of the players' personal costs. A *social optimum* is a strategy profile $\boldsymbol{\sigma}^*$ minimizing SUM. For the sake of conciseness, once a particular social optimum has been fixed, we write o_e to denote the value $k_e(\boldsymbol{\sigma}^*)$.

For any $\epsilon \geq 0$, an *ϵ-approximate pure Nash equilibrium* is a strategy profile $\boldsymbol{\sigma}$ such that, for any player $i \in N$ and strategy $S \in \Sigma_i$, $cost_i(\boldsymbol{\sigma}) \leq (1+\epsilon)cost_i(\boldsymbol{\sigma}_{-i}, S)$. We denote by $\mathsf{NE}_\epsilon(\mathsf{CG})$ the set of ϵ-approximate pure Nash equilibria of a congestion game CG. For any $\epsilon \geq 0$, an *ϵ-approximate one-round walk* is an online process in which players appear sequentially according to an arbitrary order and, upon arrival, each player irrevocably chooses a strategy approximatively minimizing a certain cost function. Let $\boldsymbol{\sigma}^i$ denote the strategy profile obtained when the first i players have performed their strategic choice, while the remaining ones have not entered the game yet (so, it may be assumed that each of them is playing the empty strategy). The i-th *selfish player* aims at minimizing her personal cost, so that $cost_i(\boldsymbol{\sigma}^i) \leq (1+\epsilon)\min_{S\in\Sigma_i} cost_i(\boldsymbol{\sigma}^{i-1}, S)$; the i-th *cooperative player* aims at minimizing the marginal increase in the social function SUM, so that $\mathsf{SUM}(\boldsymbol{\sigma}^i) - \mathsf{SUM}(\boldsymbol{\sigma}^{i-1}) \leq (1+\epsilon)\min_{S\in\Sigma_i}(\mathsf{SUM}(\boldsymbol{\sigma}^{i-1}, S) - \mathsf{SUM}(\boldsymbol{\sigma}^{i-1}))$. For $\epsilon = 0$, we speak of an *exact one-round walk*. We denote by $\mathsf{ORW}^s_\epsilon(\mathsf{CG})$ (resp. $\mathsf{ORW}^c_\epsilon(\mathsf{CG})$) the set of strategy profiles which can be constructed by an ϵ-approximate one-round walk involving selfish (resp. cooperative) players in a congestion game CG.

The *ϵ-approximate price of anarchy* of a congestion game CG is defined as $\mathsf{PoA}_\epsilon(\mathsf{CG}) = \max_{\boldsymbol{\sigma}\in\mathsf{NE}_\epsilon(\mathsf{CG})}\{\mathsf{SUM}(\boldsymbol{\sigma})/\mathsf{SUM}(\boldsymbol{\sigma}^*)\}$, where $\boldsymbol{\sigma}^*$ is a social optimum for CG. Similarly, the *competitive ratio* of ϵ-approximate one-round walks generated by selfish (resp. cooperative) players, is defined as $\mathsf{CR}^s_\epsilon(\mathsf{CG}) = \max_{\boldsymbol{\sigma}\in\mathsf{ORW}^s_\epsilon(\mathsf{CG})}\{\mathsf{SUM}(\boldsymbol{\sigma})/\mathsf{SUM}(\boldsymbol{\sigma}^*)\}$ (resp. $\mathsf{CR}^c_\epsilon(\mathsf{CG}) = \max_{\boldsymbol{\sigma}\in\mathsf{ORW}^c_\epsilon(\mathsf{CG})}\{\mathsf{SUM}(\boldsymbol{\sigma})/\mathsf{SUM}(\boldsymbol{\sigma}^*)\}$). Given a class of congestion games $\mathcal{G}$, the ϵ-approximate price of anarchy of $\mathcal{G}$ is defined as $\mathsf{PoA}_\epsilon(\mathcal{G}) = \sup_{\mathsf{CG}\in\mathcal{G}}\mathsf{PoA}_\epsilon(\mathsf{CG})$. For the case of $\epsilon = 0$, we refer to this metric simply as to the *price of anarchy*. The competitive ratio of ϵ-approximate one-round walks of $\mathcal{G}$ generated by both selfish and cooperative players is defined accordingly. Throughout the paper, we shall assume that, in any considered class of latency functions, there always exists a non-constant function, otherwise the inefficiency of all the ϵ-approximate solution concepts we consider is always equal to $1+\epsilon$.

3 Weighted Load Balancing Games

In this section, we first show that the approximate price of anarchy of weighted congestion games cannot improve even when restricting the players' strategy space to the simplest possible combinatorial structure, i.e., to the case of symmetric load balancing games.

▶ **Theorem 1.** *Let $\mathcal{C} = \{f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}\}$ be a class of non-decreasing latency functions whose members, except for the constant ones, are unbounded and such that, for any $f \in \mathcal{C}$ and $\alpha \geq 0$, the functions $g, h : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that $g(x) = \alpha f(x)$ and $h(x) = g(\alpha x)$ for each $x \in \mathbb{R}_{\geq 0}$ belong to $\mathcal{C}$. Then, $\mathsf{PoA}_\epsilon(\mathsf{W}(\mathcal{C})) = \mathsf{PoA}_\epsilon(\mathsf{WSLB}(\mathcal{C}))$.*

Proof Sketch. We make use of a multi-graph representation of a pair of strategy profiles for a (symmetric) load balancing game, denoted as *load balancing graph*, defined as follows: nodes are all the resources in E, and each player is associated to a weighted edge (e_1, e_2, w), where $\{e_1\}$ is denoted as her first strategy, $\{e_2\}$ is her second strategy, and w is her weight.

Let $k_1 > 0$ and $k_2 \geq 0$ be two real numbers, n be a positive integer, and f_1, f_2 be two non-constant (and so, unbounded) functions belonging to $\mathcal{C}$. Consider a load balancing graph $\mathsf{LB}(k_1, k_2)$ yielded by a directed n-ary tree, organized in $2s$ levels, numbered from 1 to $2s$, and whose edges are oriented from the root to the leaves, with the addition of n self-loops on the nodes of level $2s$. The weight of a player associated to an edge outgoing from a node at level $i \in [s]$ (resp. $i \in [s+1, 2s]$) is equal to $(k_1/n)^i$ (resp. $(k_1/n)^s(k_2/n)^{i-s}$). For $i, j \in [2]$, define $\theta_{i,j} = \frac{f_i(k_i)}{(1+\epsilon)f_j(k_j+1)}$ and $\theta_i = \theta_{i,i}$. Each resource at level i has latency $g_i(x) = \theta_1^{i-1} f_1\left(\left(\frac{n}{k_1}\right)^{i-1} x\right)$ if $i \in [1, s]$ and $g_i(x) = \theta_1^{s-1}\theta_{1,2}\theta_2^{i-s-1} f_2\left(\left(\frac{n}{k_1}\right)^{s}\left(\frac{n}{k_2}\right)^{i-s-1} x\right)$, otherwise.

For a sufficiently large n, the strategy profile $\boldsymbol{\sigma}$ in which all players select their first strategy is an ϵ-approximate pure Nash equilibrium. Towards this end, consider a player whose first strategy is a resource from level i. Since the game is symmetric, we have to consider the following cases: (1) if $i \in [1, 2s-1]$ and the player deviates to a resource from level $i+1$, her cost decreases exactly of a factor of $1+\epsilon$; (2) if $i \in [2, 2s]$ and the player deviates to a resource from level $j \leq i$, her cost does not decrease; if $i \in [1, 2s-2]$ and the player deviates to a resource from level $j > i+1$, for a sufficiently large n, her cost does not decrease. Let $\boldsymbol{\sigma}^*$ be the strategy profile in which each player plays her second strategy. We can show that, for each $M < \mathsf{PoA}_\epsilon(\mathsf{W}(\mathcal{C}))$, there exist $k_1 > 0$ and $k_2 \geq 0$ such that $\lim_{s\to\infty} \frac{\mathsf{SUM}(\boldsymbol{\sigma})}{\mathsf{SUM}(\boldsymbol{\sigma}^*)} > M$, thus proving the thesis. This technical claim, together with the full proof of the theorem, resembles a similar result used in [6, 26]. ◀

Then, we prove that no improvements are possible for approximate one-round walks when restricting to load balancing games.

▶ **Theorem 2.** *Let $\mathcal{C} = \{f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}\}$ be a class of non-decreasing latency functions such that, for any $f \in \mathcal{C}$ and $\alpha \geq 0$, the functions $g, h : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that $g(x) = \alpha f(x)$ and $h(x) = g(\alpha x)$ for each $x \in \mathbb{R}_{\geq 0}$ belong to $\mathcal{C}$. Then $\mathsf{CR}^s_\epsilon(\mathsf{W}(\mathcal{C})) = \mathsf{CR}^s_\epsilon(\mathsf{WLB}(\mathcal{C}))$. If all functions in $\mathcal{C}$ are semi-convex, we have that $\mathsf{CR}^c_\epsilon(\mathsf{W}(\mathcal{C})) = \mathsf{CR}^c_\epsilon(\mathsf{WLB}(\mathcal{C}))$.*

Proof Sketch. Let us start with the case of selfish players. We extend the load balancing graph $\mathsf{LB}(k_1, k_2)$ used in the proof of Theorem 1 as follows. Denote as $i(v)$ the level of resource v. For each node u in the load balancing graph, consider an arbitrary enumeration of all the n outgoing edges of u. Since each node has a unique incoming edge, we denote by $h(v) \in [n]$ the position associated to the unique edge entering v in the given ordering.

Consider the ϵ-approximate one-round walk in which players enter the game in non-increasing order of level (with respect to their first strategy) and, within the same level, players are processed in non-decreasing order of position.

For $i, j \in [2]$ and $h \in [n]$, define $\theta_{i,j}(h) = \frac{f_i\left(\frac{hk_i}{n}\right)}{(1+\epsilon)f_j(k_j+1)}$ and $\theta_i(h) = \theta_{i,i}(h)$. Resource v has latency function $g_v(x) = f_1(x)$ if $i(v) = 1$, $g_v(x) = \underbrace{\theta_1(h(v))A_u}_{A_v} f_1\left(\left(\frac{n}{k_1}\right)^{i(v)-1} x\right)$ if

d	Selfish Players	Coordinated Players	d	Selfish Players	Coordinated Players
1	7.464	5.828	6	27,089,557	7,553,550
2	90.3	56.94	7	974,588,649	222,082,591
3	1,521	780.2	8	39,729,739,895	7,400,694,480
4	32,896	13,755	9	1,809,913,575,767	275,651,917,450
5	868,567	296,476	∞	$(\Theta(d))^{d+1}$	$(\Theta(d))^{d+1}$

Figure 1 The competitive ratio of exact one-round walks generated by either selfish or cooperative players in weighted load balancing games with polynomial latency functions of maximum degree d.

$i(v) \in [2, s]$, $g_v(x) = \underbrace{\theta_{1,2}(h(v))A_u}_{A_v} f_2\left(\left(\frac{n}{k_1}\right)^s x\right)$ if $i(v) = s+1$, while, in all the other cases, $g_v(x) = \underbrace{\theta_2(h(v))A_u}_{A_v} f_2\left(\left(\frac{n}{k_1}\right)^s \left(\frac{n}{k_2}\right)^{i(v)-s-1} x\right)$, where (u, v) denotes the unique incoming edge of v and A_v is recursively defined on the basis of A_u by setting $A_v = 1$ for $i(v) = 1$, i.e., for v being the root of the tree.

The strategy profile $\boldsymbol{\sigma}$ in which all players select their first strategy is a possible outcome of an ϵ-approximate one-round walk generated by selfish players. Let $\boldsymbol{\sigma}^*$ be the strategy profile in which all players select their second strategy. As the game is not symmetric, we can assume that all players can choose among these two strategies only. We can show that, for each $M < \mathsf{CR}^s_\epsilon(\mathsf{W}(\mathcal{C}))$, there exist $k_1 > 0$ and $k_2 \geq 0$ such that $\lim_{s\to\infty} \lim_{n\to\infty} \frac{\mathsf{SUM}(\boldsymbol{\sigma})}{\mathsf{SUM}(\boldsymbol{\sigma}^*)} > M$, thus proving the thesis. Again, this technical claim, together with the full proof of the theorem, resembles a similar result used in [6, 26]. For the case of cooperative players, it suffices considering the same load balancing graph, with $n = 1$ and $\theta_{i,j}(1) = \frac{k_i f_i(k_i)}{(1+\epsilon)((k_j+1)f_i(k_j+1)-k_j f_j(k_j))}$. ◀

3.1 Polynomial Latency Functions

Consider the class $\mathcal{P}(d)$ of polynomials with non-negative coefficients and maximum degree d. Observe that this class of latency functions satisfies the hypothesis required by Theorems 1 and 2. By applying similar arguments to those used in [3], we get $\mathsf{CR}^s_\epsilon(\mathcal{P}(d)) = (\varphi_{\epsilon,d+1})^{d+1}$ and $\mathsf{CR}^c_\epsilon(\mathcal{P}(d)) = \left(\varphi'_{\epsilon,d+1}\right)^{d+1}$, where $\varphi_{\epsilon,d+1}$ and $\varphi'_{\epsilon,d+1}$ are the unique solutions of the equations $\frac{x^{d+1}}{d+1} - (1+\epsilon)(x+1)^d = 0$ and $(2+\epsilon)x^{d+1} - (1+\epsilon)(x+1)^{d+1} = 0$, respectively. Observe that $\varphi'_{\epsilon,d+1} = \frac{1}{\sqrt[d+1]{\frac{2+\epsilon}{1+\epsilon}}-1}$ which generalizes the bounds given in [9] for the case $\epsilon = 0$. Some values for the case of $\epsilon = 0$ are reported in Figure 1.

4 Unweighted Load Balancing Games

In this section, we first show that the ϵ-approximate price of anarchy of unweighted congestion games cannot improve when restricting to load balancing games.

▶ **Theorem 3.** *Let $\mathcal{C}$ be a class of non-decreasing latency functions such that $f \in \mathcal{C}, \alpha \geq 0 \Rightarrow \alpha f \in \mathcal{C}$. Then $\mathsf{PoA}_\epsilon(\mathsf{ULB}(\mathcal{C})) = \mathsf{PoA}_\epsilon(\mathsf{U}(\mathcal{C}))$.*

Proof Sketch. Let $k_1, o_1, o_2 > 0$ and $k_2 \geq 0$ be non-negative integers. Consider a load balancing game defined by a multi-partite directed graph $\mathsf{LB}(k_1, k_2, o_1, o_2)$ organized in $2s$ levels, numbered from 1 to $2s$, and defined as follows. For each $i \in [s]$ (resp. $i \in [s+1, 2s]$) there are $o_1^{s-i} k_1^{i-1} o_2^s$ (resp. $o_2^{2s-i} k_2^{i-s-1} k_1^s$) nodes/resources. Edges can only connect nodes of consecutive levels, except for nodes at level $2s$, each of which has k_2 self-loops. The out-degree of each node at level $i \in [s]$ (resp. $i \in [s+1, 2s]$) is k_1 (resp. k_2), and the in-degree of each node at level $i \in [2, s]$ (resp. $i \in [s+1, 2s]$ without considering self-loops) is o_1 (resp. o_2); observe that this configuration can be realized since the total number of nodes at level $i \in [s-1]$ (resp. $i = s$, resp. $i \in [s+1, 2s-1]$) multiplied by k_1 (resp. k_1, resp. k_2) is equal to the number of nodes at level $i+1$ multiplied by o_1 (resp. o_2, resp. o_2). For $i, j \in [2]$, define $\theta_{i,j} = \frac{f_i(k_i)}{(1+\epsilon)f_j(k_j+1)}$ and $\theta_i = \theta_{i,i}$. Each resource at level i has latency function $g_i(x) = \theta_1^{i-1} f_1(x)$ if $i \in [s]$, and $g_i(x) = \theta_1^{s-1}\theta_{1,2}\theta_2^{i-s-1} f_2(x)$ otherwise.

Let $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$ be the strategy profiles in which all players select their first and second strategy, respectively. As the game is not symmetric, we can assume that all players can choose among these two strategies only. Analogously to Theorem 1, it is possible to show that, for any $M < \mathsf{PoA}_\epsilon(\mathsf{U}(\mathcal{C}))$, there exist suitable non-negative integers k_1, k_2, o_1, o_2 such that $\boldsymbol{\sigma}$ is an ϵ-approximate pure Nash equilibrium and $\lim_{s\to\infty} \frac{\mathsf{SUM}(\boldsymbol{\sigma})}{\mathsf{SUM}(\boldsymbol{\sigma}^*)} > M$. ◀

Then, we prove a similar limitation for approximate one-round walks.

▶ **Theorem 4.** *Let $\mathcal{C} = \{f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}\}$ be a class of non-decreasing latency functions such that $f \in \mathcal{C}, \alpha \geq 0 \Rightarrow \alpha f \in \mathcal{C}$. Then $\mathsf{CR}^s_\epsilon(\mathsf{U}(\mathcal{C})) = \mathsf{CR}^s_\epsilon(\mathsf{WLB}(\mathcal{C}))$. If functions of $\mathcal{C}$ are semi-convex, we have that $\mathsf{CR}^c_\epsilon(\mathsf{U}(\mathcal{C})) = \mathsf{CR}^c_\epsilon(\mathsf{WLB}(\mathcal{C}))$.*

Proof Sketch. Let us start with the case of selfish players. Define $j(i) = 1$ if $i \in [s]$ and $j(i) = 2$ otherwise. We extend the load balancing graph $\mathsf{LB}(k_1, k_2, o_1, o_2)$ used in the proof of Theorem 3 according to the following recursive procedure.

- **Base Case:** partition the resources of the first level (resp. second level) in $o_{j(2)}$ (resp. $k_{j(1)}$) groups of equal size, and add edges from the first level to the second one in such a way that each resource in the first level has exactly $k_{j(1)}$ outgoing edges, each ending in a different group of the second level, and each resource in the second level has exactly $o_{j(2)}$ incoming edges, each coming from a different group of the first level; number the groups of the second level from 1 to $k_{j(1)}$ and label each resource with the number associated to the group it belongs to, for an illustrating example see figure 2 where resources belonging to different groups at level 1 are represented with different colors, resources belonging to different groups at level 2 belong to different squares and they are labeled with the number of the square they belong to.
- **Inductive Case:** as inductive hypothesis, suppose that resources at level $i \in [2s-1]$ have been partitioned into $m(i)$ groups of equal size and labeled with values from 1 to $k_{j(i-1)}$, where each label is assigned to $m(i)/k_{j(i-1)}$ distinct groups, and that all the edges from level $i-1$ to level i have been added. Partition resources at level $i+1$ in a temporary partition of $m(i)$ groups of equal size, and consider a bijective correspondence between groups at level i and groups at level $i+1$ (in Figure 2, groups at levels 2 and 3 which are in bijective correspondence, have been depicted in the same dashed square). Partition each group at level i into $o_{j(i+1)}$ subgroups of equal size, and the corresponding group at level $i+1$ into $k_{j(i)}$ subgroups of equal size (this defines the final partitioning of nodes at level $i+1$ into $m(i)k_{j(i)}$ groups), and add edges from the first group to the second one in the same way as described in the basic case, i.e. each resource in the first group has exactly $k_{j(i)}$ outgoing edges, each ending in a different subgroup of the second group, and each resource in the second group has exactly $o_{j(i)}$ incoming edges, each coming from a different subgroup of the first group. For each group at level $i+1$, number its subgroups with values from 1 to $k_{j(i)}$ and label each resource with the number associated to subgroup it belongs to. For instance, in Figure 2, consider an arbitrary dashed square including two groups at levels 2 and 3 which are in bijective correspondence. Analogously to the base case, resources belonging to different subgroups of the first (resp. second) group are represented with different colors (resp. belong to different squares and are labeled with the number of the square they belong to).

Let $h(v)$ be the label of resource v. Consider the ϵ-approximate one-round walk in which players enter the game in non-increasing order of level (with respect to their first strategy) and, within the same level, players are processed in non-decreasing order of position defined by labeling function h.

For $i, j \in [2]$ and $h \in k_i$, define $\theta_{i,j}(h) = \frac{f_i(h)}{(1+\epsilon)f_j(k_j+1)}$ and $\theta_i(h) = \theta_{i,i}(h)$. Resource v has latency function $g_v(x) = f_1(x)$ if $i(v) = 1$, $g_v(x) = \underbrace{\theta_1(h(v))A_u}_{A_v} f_1(x)$ if $i(v) \in [2, s]$, $g_v(x) = \underbrace{\theta_{1,2}(h(v))A_u}_{A_v} f_2(x)$ if $i(v) = s+1$, and $g_v(x) = \underbrace{\theta_2(h(v))A_u}_{A_v} f_2(x)$ otherwise, where (u, v) is an arbitrary incoming edge of v and A_v is recursively defined on the basis of A_u by setting $A_v = 1$ for $i(v) = 1$. By using the recursive structure of the load balancing graph, one can prove, by induction on the level of each resource v, that $A_u = A_{u'}$ if (u, v) and (u', v) are both edges of the load balancing graph, so that the definition of g_v is independent of the particular incoming edge of v.

The strategy profile $\boldsymbol{\sigma}$ all players select their first strategy is a possible outcome of an ϵ-approximate one-round walk generated by selfish players. Let $\boldsymbol{\sigma}^*$ be the strategy profile in which all players select their second strategy. We can show that, for each $M < \mathsf{CR}^s_\epsilon(\mathsf{U}(\mathcal{C}))$, there exist suitable non-negative integers k_1, k_2, o_1, o_2 such that $\lim_{s\to\infty} \lim_{n\to\infty} \frac{\mathsf{SUM}(\boldsymbol{\sigma})}{\mathsf{SUM}(\boldsymbol{\sigma}^*)} > M$, thus proving the claim. For the case of cooperative players, it suffices considering the same load balancing graph with $\theta_{i,j}(h) = \frac{h_i f_i(h_i) - (h_i-1) f_i(h_i-1)}{(1+\epsilon)((k_j+1) f_i(k_j+1) - k_j f_j(k_j))}$. ◀

4.1 Polynomial Latency Functions

Consider the class $\mathcal{P}(d)$ of polynomials with non-negative coefficients and maximum degree d. For ϵ-approximate one-round walks generated by cooperative players, by using similar arguments to those exploited in [3], one can prove that $\mathsf{CR}^c_\epsilon(\mathsf{ULB}(\mathcal{P}(d)))$ is equal to $\mathsf{CR}^c_\epsilon(\mathsf{WLB}(\mathcal{P}(d)))$ if $\varphi'_{d,\epsilon}$ is an integer (see Subsection 3.1), otherwise we get $\mathsf{CR}^c_\epsilon(\mathsf{ULB}(\mathcal{P}(d))) = \gamma_{d,\epsilon}\left(\left\lfloor \varphi'_{d,\epsilon} \right\rfloor\right)$, where $\gamma_{d,\epsilon}(k) := k^{d+1} + x_{d,\epsilon}(k)\left(-k^{d+1} + (1+\epsilon)\cdot\left((k+1)^{d+1} - k^{d+1}\right)\right)$, and $x_{d,\epsilon}(k)$ is such that $\gamma_{d,\epsilon}(k) = \gamma_{d,\epsilon}(k+1)$. Some values for the case of $\epsilon = 0$ are reported in Figure 3. For the case of selfish players, by using the approach in [7], we get that $\mathsf{CR}^s_0(\mathsf{ULB}(\mathcal{P}(1))) = 2 + \sqrt{5}$, $\mathsf{CR}^s_0(\mathsf{ULB}(\mathcal{P}(2))) = \frac{3383}{90}$ and $\mathsf{CR}^s_0(\mathsf{ULB}(\mathcal{P}(3))) = \frac{17929}{34}$.

5 The Case of Identical Resources

In this section, we characterize the approximate price of anarchy of weighted symmetric load balancing games with identical resources having semi-convex latency functions. We start by showing the upper bound.

▶ **Theorem 5** (Upper bound). *Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a non-decreasing and semi-convex latency function. Let $\mathsf{WSILG}(f)$ be the class of weighted symmetric load balancing games with identical resources having latency function f. For any $\epsilon \geq 0$, let*

$$inv(x) := \inf\{t \geq 0 : f(x) \leq (1+\epsilon) f(x/2 + t)\},$$
$$opt(x, \lambda) := \lambda x + (1-\lambda) inv(x),$$
$$upp(x, \lambda) := \frac{\lambda x f(x) + (1-\lambda) inv(x) f(inv(x))}{opt(x, \lambda) f(opt(x, \lambda))}.$$

If $inv(x) \neq 0$ for each $x \in \mathbb{R}_{>0}$, then:

$$\mathsf{PoA}_\epsilon(\mathsf{WSILG}(f)) \leq \sup_{x \in \mathbb{R}_{>0}} \max_{\lambda \in (0,1)} upp(x, \lambda). \tag{1}$$

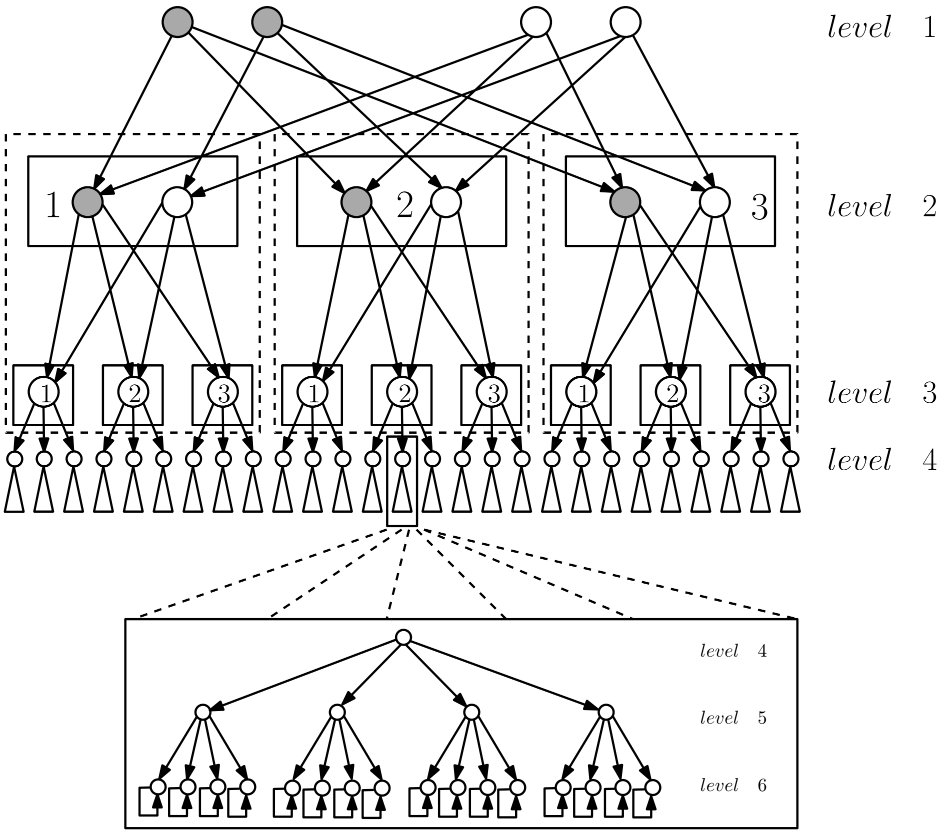

Figure 2 The load balancing graph described in the proof of Theorems 3 and 4, with $s = 3$, $k_1 = 3$, $o_1 = 2$, $k_2 = 4$ and $o_2 = 1$. We also describe the partitioning and labeling structures used in the proof of Theorem 4.

d	Competitive Ratio	d	Competitive Ratio	d	Competitive Ratio
1	5.66	4	13,170	7	220,349,064
2	55.46	5	289,648	8	7,022,463,077
3	755.2	6	7,174,495	∞	$(\Theta(d))^{d+1}$

Figure 3 The competitive ratio of exact one-round walks generated by cooperative players in unweighted load balancing games with polynomial latency functions of maximum degree d.

Proof Sketch. Let $\mathsf{WSILG}(f, W, m) \subseteq \mathsf{WSILG}(f)$ be the subclass of load balancing games having m resources and such that $\sum_{i \in N} w_i = W$. First, we prove that the optimal social cost of games in $\mathsf{WSILG}(f, W, m)$ is lower bounded by the cost of a strategy profile $\boldsymbol{\sigma}^*(W, m)$ in which all resources have the same congestion, so that $\mathsf{SUM}(\boldsymbol{\sigma}^*(W, m)) = Wf\left(\frac{W}{m}\right)$.

Furthermore, we prove that the supremum of the social cost over all ϵ-approximate pure Nash equilibria of games in $\mathsf{WSILG}(f, W, m)$ is upper bounded by the supremum of the social cost over all strategy profiles $\boldsymbol{\sigma}(m, x, h)$ in which all the resources can have three possible congestions, namely x, y, z, such that $z = inv(x) \leq y \leq x$, one resource has congestion y and $h \in [0, m-1]$ resources have congestion equal to x, so that $\mathsf{SUM}(\boldsymbol{\sigma}(m, x, h)) = hxf(x) + yf(y) + (m-h-1)inv(x)f(inv(x))$. Observe that it must be $W = hx + y + (m-h-1)inv(x)$.

One can show that that:

$$\begin{aligned}
&\mathsf{PoA}_\epsilon(\mathsf{WSILG}(f)) && (2)\\
&= \sup_{W\geq 0, m\in\mathbb{N}} \mathsf{PoA}_\epsilon(\mathsf{WSILG}(f,W,m))\\
&\leq \sup_{m\in\mathbb{N}, h\in[0,m-1], x\geq 0, y: inv(x)\leq y\leq x} \frac{\mathsf{SUM}(\boldsymbol{\sigma}(m,x,h))/m}{\mathsf{SUM}(\boldsymbol{\sigma}(hx+y+(m-h-1)inv(x)))/m}\\
&= \sup_{m\in\mathbb{N}, h\in[0,m-1], x\geq 0, y: inv(x)\leq y\leq x} \frac{\frac{hxf(x)+yf(y)+(m-h-1)inv(x)f(inv(x))}{m}}{\left(\frac{hx+y+(m-h-1)inv(x)}{m}\right) f\left(\frac{hx+y+(m-h-1)inv(x)}{m}\right)}\\
&= \lim_{m\to\infty} \sup_{h\in[0,m-1], x\geq 0, y: inv(x)\leq y\leq x} \frac{\frac{hxf(x)+yf(y)+(m-h-1)inv(x)f(inv(x))}{m}}{\left(\frac{hx+y+(m-h-1)inv(x)}{n}\right) f\left(\frac{hx+y+(m-h-1)inv(x)}{m}\right)} && (3)\\
&= \sup_{x\in\mathbb{R}_{>0}} \max_{\lambda\in(0,1)} \frac{\lambda xf(x)+(1-\lambda)inv(x)f(inv(x))}{opt(x,\lambda)f(opt(x,\lambda))}\\
&= \sup_{x\in\mathbb{R}_{>0}} \max_{\lambda\in(0,1)} upp(x,\lambda) && (4)
\end{aligned}$$

thus proving the claim (in (3) we have replaced h/m with λ, $(m-h-1)/m$ with $1-\lambda$ and y/m with 0). ◀

We show that, under mild assumptions, a tight lower bound can be obtained.

▶ **Theorem 6** (Lower Bound). *For any $\epsilon \geq 0$, let $\lambda^*(x) := \arg\max_{\lambda\in(0,1)} upp(x,\lambda)$ for any $x\in\mathbb{R}_{\geq 0}$. If $\lambda^*(x)\leq\frac{1}{2}$ and $opt(x,\lambda^*(x))-x/2\geq 0$, then*

$$\mathsf{PoA}_\epsilon(\mathsf{WSILG}(f)) = \sup_{x\in\mathbb{R}_{>0}} \max_{\lambda\in(0,1)} upp(x,\lambda). \tag{5}$$

Proof Sketch. Given $m\in\mathbb{N}$, let $h(m)\in[m]$. We prove that, if $h(m)/m$ approaches $\lambda^*(x)$ for $m\to\infty$, the strategy profiles $\boldsymbol{\sigma}^* := \boldsymbol{\sigma}^*(mx+y+(m-h-1)inv(x),m)$ and $\boldsymbol{\sigma} := \boldsymbol{\sigma}(m,x,h(m))$ defined in the proof of Theorem 5, can be enforced as an optimal strategy profile and an ϵ-approximate pure Nash equilibrium for the relative game, respectively. Thus, by using similar arguments to those exploited to obtain (4), we get

$$\lim_{m\to\infty} \sup_{x\geq 0, y: inv(x)\leq y\leq x} \frac{\mathsf{SUM}(\boldsymbol{\sigma})}{\mathsf{SUM}(\boldsymbol{\sigma}^*)} = \sup_{x\in\mathbb{R}_{>0}} upp(x,\lambda^*(x)),$$

thus concluding the proof. ◀

5.1 Polynomial Latency Functions

By exploiting (5), we derive exact bounds on the price of anarchy of weighted symmetric load balancing games with identical resources having polynomial latency functions. In Figure 4, we show a comparison between the cases of general and identical resources with respect to the price of anarchy for games with polynomial latency functions.

▶ **Theorem 7.** *Let $\mathcal{P}(d)$ be the class of polynomial latency functions of maximum degree d. Then,* $\mathsf{PoA}_0(\mathcal{P}(d)) = \frac{d^d(2^{d+1}-1)^{d+1}}{2^d(d+1)^{d+1}(2^d-1)^d} \in \Theta\left((2+o(1))^d\right)$.

d	Identical Resources	General Resources	d	Identical Resources	General Resources
1	1.125	2.618	6	7.544	14,099
2	1.412	9.909	7	12.866	118,926
3	1.946	47.82	8	22.478	1,101,126
4	2.895	277	9	39.984	11,079,429
5	4.571	1,858	∞	$\Theta\left((2+o(1))^d\right)$	$\left(\Theta\left(\frac{d}{\log d}\right)\right)^{d+1}$

Figure 4 The price of anarchy of weighted symmetric load balancing games with polynomial latency functions of maximum degree d: a comparison between the cases of identical and general resources.

5.2 Lower Bounds for Exact One-Round Walks

The following construction gives a class of lower bounding instances for exact one-round walks generated by selfish/cooperative players in load balancing games with identical resources having latency function f. Fix $n \in \mathbb{N}$ and a sequence of integers $1 = o_1 \leq o_2 \leq \ldots \leq o_n$. Let $E = E_0 \supset E_1 \supset E_2 \supset \ldots \supset E_n \supset E_{n+1} = \emptyset$ be a sequence of sets of resources such that $(|E_{i-1}| - |E_i|)o_i = |E_i|$ (observe that such a sequence exists). For any $i \in [n]$, we have $|E_i|$ players of type i whose set of strategies is E_{i-1}. Suppose that players enter the game in non-decreasing order with respect to their type. One can easily prove that the strategy profile $\boldsymbol{\sigma}$ in which each player of type i selects a different resource $e \in E_i$ is a possible outcome for an exact one-round walk generated by selfish/cooperative players. Consider the strategy profile in which, for any resource $e \in E_{i-1} \setminus E_i$, there are exactly o_i players of type i selecting e. We get:

$$\mathsf{CR}_0^s(\{f\}) \geq \frac{\mathsf{SUM}(\boldsymbol{\sigma})}{\mathsf{SUM}(\boldsymbol{\sigma}^*)} = \frac{\sum_{i=1}^n (|E_i| - |E_{i+1}|) i f(i)}{\sum_{i=1}^n (|E_{i-1}| - |E_i|) o_i f(o_i)}. \tag{6}$$

For linear latency functions, by using $n = 10^{13}$ and $o_i = \left\lfloor \frac{i44411}{100000} + 1 + \left\lfloor \frac{\sqrt{i}}{7} \right\rfloor \right\rfloor$, by (6), we get a lower bound of at least 4.0009 which improves the currently known lower bound of 4 given in [10]. We conjecture that a tight class of lower bounding instances for linear and more general polynomial latency functions is given by the union of all the instances described above, over all values of $n \in \mathbb{N}$ and all sequences $(o_i)_{i \in [n]}$.

6 Open Problems

Our work leaves two open problems. The first is to understand whether better performance is possible for approximate one-round walks in weighted symmetric load balancing games (we conjecture this is not the case), while the second is to give upper bounds on the performance of one-round walks in weighted and unweighted load balancing games with identical resources.

References

1 H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. *Journal of ACM*, 55(6), 2008.

2 H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. *Theoretical Computer Science*, 410(17):1552–1563, 2009.

3 S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing*, 40(5):1211–1233, 2011.

4 B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.

5 B. Awerbuch, Y. Azar, E. F. Grove, M.-Y. Kao, P. Krishnan, and J. S. Vitter. Load balancing in the L_p norm. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 383–391, 1995.

6 K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: price of anarchy, universal worst-case examples, and tightness. *ACM Transactions on Economics and Computation*, 2(4):1–23, 2014.

7 V. Bilò. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. In *Proceedings of the 10th Workshop on Approximation and Online Algorithms (WAOA)*, volume 7846 of *LNCS*, pages 229–241, 2012.

8 V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. Performances of one-round walks in linear congestion games. *Theory of Computing Systems*, 49(1):24–45, 2011.

9 I. Caragiannis. Better bounds for online load balancing on unrelated machines. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 972–981, 2008.

10 I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. *Algorithmica*, 61(3):606–637, 2011.

11 G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.

12 G. Christodoulou, E. Koutsoupias, and P. G. Spirakis. On the performance of approximate equilibria in congestion games. *Algorithmica*, 61(1):116–140, 2011.

13 G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. *Theoretical Computer Science*, 438:13–27, 2012.

14 J. Correa, J. de Jong, B. de Keijzer, and M. Uetz. The curse of sequentiality in routing games. In *Proceedings of the 11th International Conference on Web and Internet Economics (WINE)*, volume 9470 of *LNCS*, pages 258–271, 2015.

15 J. de Jong, M. Klimm, and M. Uetz. Efficiency of equilibria in uniform matroid congestion games. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, volume 9928 of *LNCS*, pages 105–116, 2016.

16 A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004.

17 D. Fotakis. Stackelberg strategies for atomic congestion games. *Theory of Computing Systems*, 47(1):218–249, 2010.

18 D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348:226–239, 2005.

19 M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The price of anarchy for polynomial social cost. *Theoretical Computer Science*, 369(1–3):116–135, 2006.

20 M. Gairing and F. Schoppmann. Total latency in singleton congestion games. In *Proceedings of the Third International Workshop on Internet and Network Economics (WINE)*, volume 4858 of *LNCS*, pages 381–387, 2007.

21 T. Harks and M. Klimm. On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research*, 37(3):419–436, 2012.

22 E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1653 of *LNCS*, pages 404–413, 1999.

23 T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. *Theoretical Computer Science*, 406(3):187–2006, 2008.

24 J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Science*, 36(1):48–49, 1950.

25 R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

26 T. Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of ACM*, 62(5):32:1–32:42, 2015.

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