Total Latency in Singleton Congestion Games*

Martin Gairing[1] and Florian Schoppmann[1,2]

[1] Faculty of Computer Science, Electrical Engineering and Mathematics,
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany.
[2] International Graduate School of Dynamic Intelligent Systems
{gairing,fschopp}@uni-paderborn.de

Abstract. In this work, we consider *singleton congestion games* as a special class of (weighted) congestion games where each players' strategy consists only of a *single* resource.
For singleton congestion games, we provide a collection of upper and lower bounds on the *price of anarchy* for multiple interesting cases. In our study, we distinguish between restricted and unrestricted strategy sets, between weighted and unweighted player weights, and between linear and polynomial latency functions.

1 Introduction

Motivation and Framework. The *price of anarchy,* also known as *coordination ratio,* has been defined in the seminal work by Koutsoupias and Papadimitriou [18] as a measure of the extent to which competition approximates cooperation. In general, the price of anarchy is the worst-case ratio between the value of a social objective function, usually coined as *social cost,* in some equilibrium state of a system, and that of some social optimum. Usually, the equilibrium state has been taken to be that of a *Nash equilibrium* [23] – a state in which no *user* or *player* wishes to unilaterally leave its own strategy in order to improve the value of its private objective function, also known as *private cost.* A Nash equilibrium is called *pure* if all players choose a pure strategy, and *mixed* if players choose probability distributions over strategies. The price of anarchy represents a rendezvous of Nash equilibrium, a concept fundamental to Game Theory, with *approximation,* an omnipresent concept in Theoretical Computer Science today (see, e.g., [28]).

Rosenthal [25] introduced a special class of non-cooperative games, now widely known as *congestion games.* Here, the strategy set of each player is a subset of the power set of given resources. The latency on each resource is described by a latency function in the number of players sharing this resource, and the private cost of a player is the sum of the latencies over its chosen resources. Milchtaich [20] considered *weighted congestion games* as an extension to congestion games in which the players have weights and thus different influence on the latency of the resources. Weighted congestion games provide us with a general framework for modeling any kind of non-cooperative resource sharing problem. One such classical resource sharing problem is that of *load balancing.*

In this paper, we study selfish load balancing games that we call (weighted or unweighted) *singleton congestion games.* Singleton congestion games are congestion games where each player's strategy consists only of a *single* resource. The strategy set of a player is *unrestricted* if it contains all resources, and *restricted* otherwise. For singleton congestion games, we further examine the price of anarchy and use the *total latency* (see e.g. [26]) as our social objective function. In our study, we distinguish between *polynomial*, *affine* and *linear* latency functions. All our latency functions have non-negative coefficients.

Related Work. The class of *congestion games* was introduced by Rosenthal [25] and extensively studied afterwards (see, e.g., [11,20,21]). In Rosenthal's model the strategy of each player is a subset

* This work has been partially supported by the European Union within the Integrated Project IST-15964 "Algorithmic Principles for Building Efficient Overlay Computers" (AEOLUS).

of resources. Resource latency functions can be arbitrary but they only depend on the number of players sharing the same resource. Later, Milchtaich [20] considered weighted congestion games as an extension to congestion games, where players have weights and thus different influence in the latency of a resource.

The *price of anarchy* [24], also known as *coordination ratio*, was first introduced and studied by Koutsoupias and Papadimitriou [18]. As a starting point of their investigation they considered a weighted singleton congestion game with m resources, unrestricted strategy sets and linear latency functions. They defined social cost as the expected maximum latency on a resource. This model is now known as KP-model. In this setting, there exist *tight* bounds on the price of anarchy of $\Theta(\frac{\log m}{\log\log m})$ for identical linear latency functions [9,17] and $\Theta(\frac{\log m}{\log\log\log m})$ for linear latency functions [9]. The price of anarchy has also been studied for variations of the KP-model, namely for non-linear latency functions [8], for the case of restricted strategy sets [3,12,13], for the case of incomplete information [15] and for player-specific latency functions [16].

Singleton congestion games with social cost defined as the total latency have been studied in [5,14,19,27]; see Table 1 for a comparison of their bounds on the price of anarchy. Such games always possess a pure Nash equilibrium when latency functions are non-decreasing [10]. Hence, also the *pure* price of anarchy is of interest, where Nash equilibria are restricted to be pure. Specifically, Lücking et al. [19] considered the case of linear latency functions and unrestricted strategy sets. Note that the upper bound for identical linear latency functions and arbitrary player weights follows by employing an additional argument from [15]. The case of polynomial latency functions, unrestricted strategy sets, and unweighted players was studied by Gairing et al. [14]. Finally, Suri et al. [27] and Caragiannis et al. [5] studied the case of restricted strategy sets with a focus on the pure price of anarchy. In particular, for weighted players and linear latency functions, Caragiannis et al. [5] give a lower bound that (for $m \to \infty$) matches a corresponding upper bound from [2].

For general (weighted) congestion games and social cost defined as the total latency, exact values for the price of anarchy have been given in [1,2,7]. In case of linear latency functions, the price of anarchy is exactly $\frac{5}{2}$ for unweighted congestion games [7] and $1+\Phi$ for weighted congestion games [2], where $\Phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. For polynomial latency functions (of maximum degree d and with non-negative coefficients), Aland et al. [1] proved exact values on the price of anarchy. In particular, they showed that for identical players the price of anarchy is exactly $\Upsilon(d) := \frac{(\lambda+1)^{2d+1}-\lambda^{d+1}(\lambda+2)^d}{(\lambda+1)^{d+1}-(\lambda+2)^d+(\lambda+1)^d-\lambda^{d+1}}$, where $\lambda = \lfloor \Phi_d \rfloor$ and Φ_d is a natural generalization of the golden ratio to larger dimensions such that Φ_d is the (only) positive real solution to $(x+1)^d = x^{d+1}$. For arbitrary players the price of anarchy increases slightly to Φ_d^{d+1} [1]. Note that all upper bounds on the price of anarchy for congestion games are also upper bounds for singleton congestion games.

Contribution and Comparison. In this work, we prove a collection of upper and lower bounds on the price of anarchy for multiple interesting classes of singleton congestion games. The new bounds are highlighted (by a gray background) in Table 1.

We first consider the case of *unrestricted strategy sets*. In particular, we show:

- For unweighted singleton congestion games with linear latency functions the price of anarchy is exactly $2 - \frac{1}{m}$ (Theorem 1 (b)). For affine latency functions, we prove a slightly weaker upper bound of 2 (Theorem 1 (a)).

  To get these results, we prove an upper bound on the total latency of any Nash equilibrium (Proposition 1) and a lower bound on the optimum total latency (Proposition 2). Both bounds also hold for *weighted* singleton congestion games with *affine* latency functions and may be of independent interest.

- For weighted singleton congestion games with polynomial latency functions of maximum degree d, we prove a lower bound on the pure price of anarchy which approaches the $(d+1)$-th Bell number B_{d+1} for $m \to \infty$ (Theorem 2). It is an interesting open problem to close the gap between this lower bound and the best known upper bound of Φ_d^{d+1} from [1], which was shown for general weighted congestion games. As a corollary, we obtain a lower bound for linear latency functions that approaches 2 as m goes to infinity (Corollary 1).
- For weighted singleton congestion games with linear latency functions, we give a lower bound of 2.036 on the price of anarchy (Theorem 3). Closing the gap between this lower bound and the upper bound of $1 + \Phi \approx 2.618$ from [2] remains tantalizingly open.

Next, we consider the case of *restricted strategy sets*. For polynomial latency functions of maximum degree d, we show:

- The exact value of $\Upsilon(d) := \frac{(\lambda+1)^{2d+1} - \lambda^{d+1}(\lambda+2)^d}{(\lambda+1)^{d+1} - (\lambda+2)^d + (\lambda+1)^d - \lambda^{d+1}}$ on the price of anarchy for unweighted congestion games from [1] also holds for unweighted singleton congestion games. To show this, we provide a lower bound that approaches the upper bound if m goes to infinity (Theorem 5). Specifically, we construct singleton congestion games having a recursive structure, thereby leading to Nash equilibria in which players' private costs may differ by orders of magnitude from their costs in an optimum state. A careful analysis then gives the desired result.
- The exact value of Φ_d^{d+1} on the price of anarchy for weighted congestion games from [1] also holds for weighted singleton congestion games. To show this, we again provide a lower bound that approaches the upper bound if m goes to infinity (Theorem 4).

Theorems 4 and 5 generalize the corresponding results from [5] to polynomial latency functions.

			$\mathsf{PoA_{pure}}$		PoA	
strategies	$f_e(x) =$	player	LB	UP	LB	UP
un-restricted	x	ident.	1		$2 - \frac{1}{m}$ [19]	
	x	arb.	$\frac{9}{8}$ [19]		$2 - \frac{1}{m}$ [19,15]	
	$a_e x$	ident.	$\frac{4}{3}$ [19]		$2 - \frac{1}{m}$ (T.1)	
	$a_e x$	arb.	$2 - o(1)$ (C.1)	$1 + \Phi$ [2]	2.036 (T.3)	$1 + \Phi$ [2]
	x^d	ident.	1		$B_{d+1} - o(1)$ [14]	B_{d+1} [14]
	$\sum_{j=0}^{d} a_{e,j} x^j$	arb.	$B_{d+1} - o(1)$ (T.2)	Φ_d^{d+1} [1]		
restricted	x	ident.	2.012 [27]	2.012 [5]		
	$a_e x$	ident.	$\frac{5}{2} - o(1)$ [5]	$\frac{5}{2}$ [27]	$\frac{5}{2} - o(1)$ [5]	$\frac{5}{2}$ [6]
	$\sum_{j=0}^{d} a_{e,j} x^j$	ident.	$\Upsilon(d) - o(1)$ (T.5)	$\Upsilon(d)$ [1]	$\Upsilon(d) - o(1)$ (T.5)	$\Upsilon(d)$ [1]
	$a_e x$	arb.	$1 + \Phi - o(1)$ [5]	$1 + \Phi$ [2]	$1 + \Phi - o(1)$ [5]	$1 + \Phi$ [2]
	$\sum_{j=0}^{d} a_{e,j} x^j$	arb.	$\Phi_d^{d+1} - o(1)$ (T.4)	Φ_d^{d+1} [1]	$\Phi_d^{d+1} - o(1)$ (T.4)	Φ_d^{d+1} [1]

Table 1. Lower bounds (LB) and upper bounds (UB) on the price of anarchy for singleton congestion games. The terms $o(1)$ are in m. B_{d+1} denotes the Bell number of order $d+1$.

Roadmap. The rest of this paper is organized as follows. In Section 2 we we give an exact definition of weighted singleton congestion games. We present our results for unrestricted strategy sets in Section 3 and for restricted strategy sets in Section 4. We conclude in Section 5.

2 Singleton Congestion Games

General. For all $k \in \mathbb{N}$ denote $[k] = \{1, \ldots, k\}$ and $[k]_0 = \{0, \ldots, k\}$. For a vector $\mathsf{v} = (v_1, \ldots, v_n)$, let $\mathsf{v}_{-\mathsf{i}} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$ and let $(\mathsf{v}_{-\mathsf{i}}, v_i') = (v_1, \ldots, v_{i-1}, v_i', v_{i+1}, \ldots, v_n)$.

The number of ways a set of k elements can be partitioned into nonempty subsets is called the k-th Bell Number [29], denoted by B_k. It is known (see e.g. [29, Identity (1.6.10)]) that for all $k \in \mathbb{N}_0$,

$$B_k = \frac{1}{\mathsf{e}} \sum_{j=0}^{\infty} \frac{j^k}{j!} \,. \tag{1}$$

Throughout, we denote by Φ_d a natural generalization of the golden ratio to larger dimensions such that Φ_d is the (only) positive real solution to $(x+1)^d = x^{d+1}$.

Instance. A *weighted singleton congestion game* Γ is a tuple

$$\Gamma = \left(n, m, (w_i)_{i \in [n]}, (S_i)_{i \in [n]}, (f_e)_{e \in [m]}\right) .$$

Here, n is the number of *players* and m is the number of *resources*. For every player $i \in [n]$, $w_i \in \mathbb{R}_{>0}$ is the *weight* and $S_i \subseteq [m]$ is the *strategy set* of player i. Denote by $W = \sum_{i \in [n]} w_i$ the *total weight* of the players. Strategy sets are *unrestricted* if $S_i = [m]$ for all $i \in [n]$; otherwise they are *restricted*. Denote $S = S_1 \times \ldots \times S_n$. For every resource $e \in [m]$, the *latency function* $f_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ defines the *latency* on resource e. We consider *polynomial latency functions* with maximum degree d and non-negative coefficients, that is, for each resource $e \in [m]$, the latency function is of the form $f_e(x) = \sum_{j=0}^{d} a_{e,j} \cdot x^j$ with $a_{e,j} \geq 0$ for all $j \in [d]_0$. For the special case of *affine latency functions*, we let $a_e := a_{e,1}$ and $b_e := a_{e,0}$, i.e., the latency function of resource $e \in [m]$ is $f_e(x) = a_e \cdot x + b_e$. Latency functions are *linear* if $b_e = 0$ for all resources $e \in [m]$.

In an *unweighted singleton congestion game*, the weights of all players are equal to 1. Thus, the latency on a resource only depends on the *number* of players choosing this resource.

Strategies and Strategy Profiles. A *pure strategy* for player $i \in [n]$ is some specific resource $s_i \in S_i$ whereas a *mixed strategy* $P_i = (p(i, s_i))_{s_i \in S_i}$ is a probability distribution over S_i, where $p(i, s_i)$ denotes the probability that player i chooses the pure strategy s_i.

A *pure strategy profile* is an n-tuple $\boldsymbol{s} = (s_1, \ldots, s_n) \in S$ whereas a *mixed strategy profile* $\mathbf{P} = (P_1, \ldots, P_n)$ is represented by an n-tuple of mixed strategies. For a mixed strategy profile $\mathbf{P}$ denote by $p(\boldsymbol{s}) = \prod_{i \in [n]} p(i, s_i)$ the probability that the players choose the pure strategy profile $\boldsymbol{s} = (s_1, \ldots, s_n)$.

Load and Private Cost. Fix a mixed strategy profile $\mathbf{P}$. Denote by $\delta_e(\mathbf{P}) = \sum_{i \in [n]} p(i, e) \cdot w_i$ the *expected load* on resource $e \in [m]$. In the same way, denote by $\delta_e^{-k}(\mathbf{P}) = \sum_{i \in [n], i \neq k} p(i, e) \cdot w_i$ the expected load of all players $i \in [n], i \neq k$, on resource $e \in [m]$. Observe that for a pure strategy profile $\boldsymbol{s}$, the (expected) load on resource $e \in E$ is $\delta_e(\boldsymbol{s}) = \sum_{i \in [n] : s_i = e} w_i$.

The *private cost* of player $i \in [n]$ in a pure strategy profile $\boldsymbol{s}$ is defined by the latency of the chosen resource. Thus,

$$\mathsf{PC}_i(\boldsymbol{s}) = f_{s_i}\left(\delta_{s_i}(\boldsymbol{s})\right) .$$

For a mixed strategy profile $\mathbf{P}$, the *private cost* of player $i \in [n]$ is

$$\mathsf{PC}_i(\mathbf{P}) = \sum_{\boldsymbol{s} \in S} p(\boldsymbol{s}) \cdot \mathsf{PC}_i(\boldsymbol{s}) \ .$$

Nash Equilibria. We are interested in a special class of (mixed) strategy profiles called Nash equilibria [22,23] that we describe here. Given a weighted singleton congestion game and an associated mixed strategy profile $\mathbf{P}$, player $i \in [n]$ is *satisfied* if it can not improve its private cost by unilaterally changing its strategy. Otherwise, player i is *unsatisfied*. The mixed strategy profile $\mathbf{P}$ is a *Nash equilibrium* if and only if all players $i \in [n]$ are satisfied, that is, $\mathsf{PC}_i(\mathbf{P}) \leq \mathsf{PC}_i(\mathbf{P}_{-i}, s_i)$ for all $i \in [n]$ and $s_i \in S_i$.

Note that if this inequality holds for all pure strategies $s_i \in S_i$ of player i, then it also holds for all mixed strategies over S_i. Also, note that in a Nash equilibrium $\mathbf{P}$, for all players $i \in [n]$, $\mathsf{PC}_i(\mathbf{P}) = \mathsf{PC}_i(\mathbf{P}_{-i}, s_i)$ for all $s_i \in S_i$ where $p(i, s_i) > 0$. Depending on the type of strategy profile, we differ between *pure* and *mixed* Nash equilibria.

For a weighted singleton congestion game Γ, denote by $\mathcal{NE}(\Gamma)$ the set of all (mixed) Nash equilibria and by $\mathcal{NE}_{\mathsf{pure}}(\Gamma)$ the set of all pure Nash equilibria for Γ.

Social Cost. Associated with a weighted singleton congestion game Γ and a mixed strategy profile $\mathbf{P}$ is the *social cost* $\mathsf{SC}(\Gamma, \mathbf{P})$ as a measure of social welfare. In particular we use the expected total latency [26], that is,

$$\begin{aligned}\mathsf{SC}(\Gamma, \mathbf{P}) &= \sum_{\boldsymbol{s} \in S} p(\boldsymbol{s}) \sum_{e \in [m]} \delta_e(\boldsymbol{s}) \cdot f_e(\delta_e(\boldsymbol{s})) \\ &= \sum_{\boldsymbol{s} \in S} p(\boldsymbol{s}) \sum_{i \in [n]} w_i \cdot f_{s_i}(\delta_{s_i}(\boldsymbol{s})) \\ &= \sum_{i \in [n]} w_i \cdot \mathsf{PC}_i(\mathbf{P}).\end{aligned}$$

The *optimum total latency* associated with a weighted singleton congestion game Γ is the least possible social cost, over all pure strategy profiles $\boldsymbol{s} \in S$. Thus,

$$\mathsf{OPT}(\Gamma) = \min_{\boldsymbol{s} \in S} \mathsf{SC}(\Gamma, \boldsymbol{s})\,.$$

Price of Anarchy. Let $\mathcal{G}$ be a class of weighted singleton congestion games. The *price of anarchy*, also called *coordination ratio* and denoted by PoA, is the supremum, over all instances $\Gamma \in \mathcal{G}$ and Nash equilibria $\mathbf{P} \in \mathcal{NE}(\Gamma)$, of the ratio $\frac{\mathsf{SC}(\Gamma, \mathbf{P})}{\mathsf{OPT}(\Gamma)}$. Thus,

$$\mathsf{PoA}(\mathcal{G}) = \sup_{\Gamma \in \mathcal{G}, \mathbf{P} \in \mathcal{NE}(\Gamma)} \frac{\mathsf{SC}(\Gamma, \mathbf{P})}{\mathsf{OPT}(\Gamma)}$$

Define $\mathsf{PoA}_{\mathsf{pure}}(\mathcal{G})$ accordingly.

3 Unrestricted Strategy Sets

In this section, we consider the case of unrestricted strategy sets. We start with a small technical lemma which has an easy proof. Afterwards, we show an upper bound on the social cost of any Nash equilibrium (Proposition 1) and a lower bound on the optimum total latency (Proposition 2). Both bounds hold for weighted singleton congestion games with affine latency functions.

Lemma 1. *If $x_j, y_j > 0$ for all $j \in [k]$ then*

$$\min_{j \in [k]} \frac{x_j}{y_j} \leq \frac{\sum_{j \in [k]} x_j}{\sum_{j \in [k]} y_j}\,.$$

Proposition 1. *Let Γ be a weighted singleton congestion game with unrestricted strategy sets, affine latency functions and associated Nash equilibrium $\mathbf{P}$. Then, for all subsets of resources $\mathcal{M} \subseteq [m]$,*

$$SC(\Gamma, \mathbf{P}) \leq \sum_{i \in [n]} w_i \cdot \frac{W + (|\mathcal{M}| - 1) w_i + \sum_{j \in \mathcal{M}} \frac{b_j}{a_j}}{\sum_{j \in \mathcal{M}} \frac{1}{a_j}}\,.$$

Proof. Since $\mathbf{P}$ is a Nash equilibrium, we have $\mathsf{PC}_i(\mathbf{P}) \le \mathsf{PC}_i(\mathbf{P}_{-i}, j)$ for all players $i \in [n]$ and resources $j \in [m]$. So,

$$\mathsf{PC}_i(\mathbf{P}) = \min_{j\in[m]}\{f_j(\delta_j^{-i}(\mathbf{P}) + w_i)\} \le \min_{j\in\mathcal{M}}\{f_j(\delta_j^{-i}(\mathbf{P}) + w_i)\},$$

where the inequality follows from the fact that $\mathcal{M} \subseteq [m]$. It follows that

$$\begin{aligned}\mathsf{SC}(\Gamma, \mathbf{P}) &= \sum_{i\in[n]} w_i \cdot \mathsf{PC}_i(\mathbf{P}) \le \sum_{i\in[n]} w_i \cdot \min_{j\in\mathcal{M}}\{f_j(\delta_j^{-i}(\mathbf{P}) + w_i)\} \\ &= \sum_{i\in[n]} w_i \cdot \min_{j\in\mathcal{M}}\{a_j(\delta_j^{-i}(\mathbf{P}) + w_i) + b_j\} = \sum_{i\in[n]} w_i \cdot \min_{j\in\mathcal{M}}\left\{\frac{\delta_j^{-i}(\mathbf{P}) + w_i + \frac{b_j}{a_j}}{\frac{1}{a_j}}\right\}.\end{aligned}$$

Applying Lemma 1 yields

$$\mathsf{SC}(\Gamma, \mathbf{P}) \le \sum_{i\in[n]} w_i \cdot \frac{\sum_{j\in\mathcal{M}}\left(\delta_j^{-i}(\mathbf{P}) + w_i + \frac{b_j}{a_j}\right)}{\sum_{j\in\mathcal{M}} \frac{1}{a_j}} \le \sum_{i\in[n]} w_i \cdot \frac{W + (|\mathcal{M}| - 1)w_i + \sum_{j\in\mathcal{M}} \frac{b_j}{a_j}}{\sum_{j\in\mathcal{M}} \frac{1}{a_j}}.$$

To see the last inequality, note that $\sum_{j\in\mathcal{M}} \delta_j^{-i}(\mathbf{P}) \le \sum_{j\in[m]} \delta_j^{-i}(\mathbf{P}) = W - w_i$. This completes the proof of the proposition. □

We proceed with a technical lemma that holds already for a more general model (see e.g. [4,26]).

Lemma 2 ([4,26]). *Let $f_1, \dots, f_m$ be semi-convex latency functions. For all $j \in [m]$, define $f_j^* = \frac{\mathsf{d}}{\mathsf{d}x}(x \cdot f_j(x))$. Define $X = \{\boldsymbol{x} \in \mathbb{R}^m_{\ge 0} : \sum_{j\in[m]} x_j = W\}$. Then $\boldsymbol{x}^* \in \arg\min_{\boldsymbol{x}\in X}\{\sum_{j\in[m]} x_j \cdot f_j(x_j)\}$ if and only if $f_{j_1}^*(x_{j_1}^*) \le f_{j_2}^*(x_{j_2}^*)$ for all $j_1, j_2 \in [m]$ with $x_{j_1}^* > 0$.*

We are now ready to prove:

Proposition 2. *Let Γ be a weighted singleton congestion game with unrestricted strategy sets and affine latency functions Let $\boldsymbol{s}$ be an associated pure strategy profile with optimum total latency and let $\mathcal{M} = \{e : \delta_e(\boldsymbol{s}) > 0\}$. Define $X = \{\boldsymbol{x} \in \mathbb{R}^{\mathcal{M}}_{\ge 0} : \sum_{j\in\mathcal{M}} x_j = W\}$ and let $\boldsymbol{x}^* \in \arg\min_{\boldsymbol{x}\in X}\{\sum_{j\in\mathcal{M}} x_j \cdot f_j(x_j)\}$. Denote $\mathcal{M}^* = \{j \in \mathcal{M} : x_j^* > 0\}$. Then,*

$$\mathit{OPT}(\Gamma) = \mathit{SC}(\Gamma, \boldsymbol{s}) \ge \frac{W^2 + \frac{W}{2} \cdot \sum_{j\in\mathcal{M}^*} \frac{b_j}{a_j}}{\sum_{j\in\mathcal{M}^*} \frac{1}{a_j}}.$$

Proof. We have

$$\begin{aligned}\mathsf{SC}(\Gamma, \boldsymbol{s}) &= \sum_{j\in[m]} f_j(\delta_j(\boldsymbol{s})) \cdot \delta_j(\boldsymbol{s}) = \sum_{j\in\mathcal{M}} f_j(\delta_j(\boldsymbol{s})) \cdot \delta_j(\boldsymbol{s}) \ge \sum_{j\in\mathcal{M}} f_j(x_j^*) \cdot x_j^* \\ &= \sum_{j\in\mathcal{M}^*} f_j(x_j^*) \cdot x_j^* = \sum_{j\in\mathcal{M}^*} \left(a_j \cdot x_j^* + b_j\right) \cdot x_j^* = \sum_{j\in\mathcal{M}^*} \frac{x_j^* + \frac{b_j}{a_j}}{\frac{1}{a_j}} \cdot x_j^*\end{aligned}$$

where the inequality follows by the definition of $\boldsymbol{x}^*$. Lemma 2 implies that for all resources $j_1, j_2 \in \mathcal{M}^*$, we have $2a_{j_1}x_{j_1}^* + b_{j_1} = 2a_{j_2}x_{j_2}^* + b_{j_2}$, or equivalently

$$\frac{x_{j_1}^* + \frac{1}{2} \cdot \frac{b_{j_1}}{a_{j_1}}}{\frac{1}{a_{j_1}}} = \frac{x_{j_2}^* + \frac{1}{2} \cdot \frac{b_{j_2}}{a_{j_2}}}{\frac{1}{a_{j_2}}}.$$

This implies that for all resources $j \in \mathcal{M}^*$,

$$\frac{x_j^* + \frac{1}{2} \cdot \frac{b_j}{a_j}}{\frac{1}{a_j}} = \frac{\sum_{k \in \mathcal{M}^*}(x_k^* + \frac{1}{2} \cdot \frac{b_k}{a_k})}{\sum_{k \in \mathcal{M}^*} \frac{1}{a_k}} = \frac{W + \frac{1}{2} \cdot \sum_{k \in \mathcal{M}^*} \frac{b_k}{a_k}}{\sum_{k \in \mathcal{M}^*} \frac{1}{a_k}} \,.$$

We get

$$\begin{aligned} \mathsf{SC}(\Gamma, \boldsymbol{s}) &\geq \sum_{j \in \mathcal{M}^*} \frac{x_j^* + \frac{b_j}{a_j}}{\frac{1}{a_j}} \cdot x_j^* \geq \sum_{j \in \mathcal{M}^*} \frac{x_j^* + \frac{1}{2} \cdot \frac{b_j}{a_j}}{\frac{1}{a_j}} \cdot x_j^* \\ &= \frac{W + \frac{1}{2} \cdot \sum_{k \in \mathcal{M}^*} \frac{b_k}{a_k}}{\sum_{k \in \mathcal{M}^*} \frac{1}{a_k}} \cdot \sum_{j \in \mathcal{M}^*} x_j^* = \frac{W^2 + \frac{W}{2} \cdot \sum_{k \in \mathcal{M}^*} \frac{b_k}{a_k}}{\sum_{k \in \mathcal{M}^*} \frac{1}{a_k}} \,. \end{aligned}$$

This completes the proof of the proposition. □

We are now equipped with all tools to prove the following upper bounds on the price of anarchy:

Theorem 1. *Let $\mathcal{G}_a$ be the class of unweighted singleton congestion games with at most m resources, unrestricted strategy sets and affine latency functions and $\mathcal{G}_b$ be the subset of $\mathcal{G}_a$ with linear latency functions. Then,*

(a) $\mathsf{PoA}(\mathcal{G}_a) < 2$
(b) $\mathsf{PoA}(\mathcal{G}_b) \leq 2 - \frac{1}{m}$

Proof. Consider an arbitrary $\Gamma \in \mathcal{G}_a$ with associated Nash equilibrium $\mathbf{P}$. Define $\mathcal{M}^*$ as in Proposition 2. Then, by Proposition 2,

$$\mathsf{OPT}(\Gamma) \geq \frac{n^2 + \frac{n}{2} \cdot \sum_{j \in \mathcal{M}^*} \frac{b_j}{a_j}}{\sum_{j \in \mathcal{M}^*} \frac{1}{a_j}}.$$

Furthermore, by Proposition 1,

$$\mathsf{SC}(\Gamma, \mathbf{P}) \leq \frac{n^2 + n \cdot (|\mathcal{M}^*| - 1) + n \cdot \sum_{j \in \mathcal{M}^*} \frac{b_j}{a_j}}{\sum_{j \in \mathcal{M}^*} \frac{1}{a_j}}.$$

Since $|\mathcal{M}^*| \leq n$, we get

$$\begin{aligned} \frac{\mathsf{SC}(\Gamma, \mathbf{P})}{\mathsf{OPT}(\Gamma)} &\leq \frac{n^2 + n \cdot (|\mathcal{M}^*| - 1) + n \cdot \sum_{j \in \mathcal{M}^*} \frac{b_j}{a_j}}{n^2 + \frac{n}{2} \cdot \sum_{j \in \mathcal{M}^*} \frac{b_j}{a_j}} \leq 1 + \frac{n^2 \cdot \frac{|\mathcal{M}^*| - 1}{|\mathcal{M}^*|} + \frac{n}{2} \cdot \sum_{j \in \mathcal{M}^*} \frac{b_j}{a_j}}{n^2 + \frac{n}{2} \cdot \sum_{j \in \mathcal{M}^*} \frac{b_j}{a_j}} \\ &< 2 \,, \end{aligned} \tag{2}$$

proving (a). To prove (b), recall that if $\Gamma \in \mathcal{G}_b$ then $b_j = 0$ for all $j \in [m]$. Thus (2) reduces to

$$\frac{\mathsf{SC}(\Gamma, \mathbf{P})}{\mathsf{OPT}(\Gamma)} \leq 1 + \frac{n^2 \cdot \frac{|\mathcal{M}^*| - 1}{|\mathcal{M}^*|}}{n^2} = 2 - \frac{1}{|\mathcal{M}^*|} \,.$$

The claim follows since $|\mathcal{M}^*| \leq m$. □

We proceed with a lower bound on the pure price of anarchy.

Theorem 2. *Let $\mathcal{G}$ be the class of weighted singleton congestion games with unrestricted strategy sets and polynomial latency functions of maximum degree d. Then*

$$\mathit{PoA}_{\mathsf{pure}}(\mathcal{G}) \geq B_{d+1} .$$

Proof. For some parameter $k \in \mathbb{N}$ define the following weighted singleton congestion game $\Gamma(k)$ with unrestricted strategy sets and polynomial latency functions:

- There are $k+1$ disjoint sets $\mathcal{M}_0, \ldots, \mathcal{M}_k$ of resources. Set $\mathcal{M}_j, j \in [k]_0$, consists of $|\mathcal{M}_j| = 2^{k-j} \cdot \frac{k!}{j!}$ resources sharing the polynomial latency function $f_e(x) = 2^{-jd} \cdot x^d$ for all resources $e \in \mathcal{M}_j$.
- There are k disjoint sets of players $\mathcal{N}_1, \ldots, \mathcal{N}_k$. Set $\mathcal{N}_j, j \in [k]$ consists of $|\mathcal{N}_j| = |\mathcal{M}_{j-1}| = 2^{k-(j-1)} \cdot \frac{k!}{(j-1)!}$ players with weight $w_i = 2^{j-1}$ for all players $i \in \mathcal{N}_j$.

Observe that $|\mathcal{M}_j| = 2^{k-j} \cdot \frac{k!}{j!} = 2^{k-(j+1)} \cdot \frac{k!}{(j+1)!} \cdot 2(j+1) = |\mathcal{M}_{j+1}| \cdot 2(j+1)$.

On the one hand, let $\boldsymbol{s}$ be a pure strategy profile that assigns exactly $2j$ players from $\mathcal{N}_j$ to each resource in $\mathcal{M}_j$ for $j \in [k]_0$. Then, for all resources $e \in \mathcal{M}_j, j \in [k]$ we have $\delta_e(\boldsymbol{s}) = 2j \cdot 2^{j-1} = j \cdot 2^j$ and $f_e(\delta_e(\boldsymbol{s})) = 2^{-jd} \cdot (j \cdot 2^j)^d = j^d$. It is now easy to check that $\boldsymbol{s}$ is a Nash equilibrium for $\Gamma(k)$ with

$$\begin{aligned}\mathsf{SC}(\Gamma(k), \boldsymbol{s}) &= \sum_{e \in [m]} \delta_e(\boldsymbol{s}) \cdot f_e(\delta_e(\boldsymbol{s})) = \sum_{j \in [k]_0} \sum_{e \in \mathcal{M}_j} \delta_e(\boldsymbol{s}) \cdot f_e(\delta_e(\boldsymbol{s})) \\ &= \sum_{j \in [k]_0} |\mathcal{M}_j| \cdot 2^j \cdot j^{d+1} = \sum_{j \in [k]_0} 2^{k-j} \cdot \frac{k!}{j!} \cdot 2^j \cdot j^{d+1} = 2^k \cdot k! \sum_{j \in [k]_0} \frac{j^{d+1}}{j!}\end{aligned}$$

On the other hand, let $\boldsymbol{s}^*$ be a strategy profile that assigns each player $\mathcal{N}_j$ to a separate resource in $\mathcal{M}_{j-1}$. Then, for all resources $e \in \mathcal{M}_j, j \in [k-1]_0$ we have $\delta_e(\boldsymbol{s}^*) = 2^j$ and $f_e(\delta_e(\boldsymbol{s}^*)) = 2^{-jd} \cdot (2^j)^d = 1$. So

$$\mathsf{SC}(\Gamma(k), \boldsymbol{s}^*) = \sum_{e \in [m]} \delta_e(\boldsymbol{s}^*) \cdot f_e(\delta_e(\boldsymbol{s}^*)) = \sum_{e \in [m]} 2^j = \sum_{j \in [k-1]_0} |\mathcal{M}_j| \cdot 2^j = 2^k \cdot k! \sum_{j \in [k-1]_0} \frac{1}{j!}$$

It follows that

$$\mathsf{PoA}_{\mathsf{pure}}(\mathcal{G}) \geq \sup_{k \in \mathbb{N}} \frac{\mathsf{SC}(\Gamma(k), \boldsymbol{s})}{\mathsf{SC}(\Gamma(k), \boldsymbol{s}^*)} \geq \lim_{k \to \infty} \frac{\mathsf{SC}(\Gamma(k), \boldsymbol{s})}{\mathsf{SC}(\Gamma(k), \boldsymbol{s}^*)} = \frac{\sum_{j=1}^{\infty} \frac{j^{d+1}}{j!}}{\sum_{j=0}^{\infty} \frac{1}{j!}} = \frac{1}{\mathsf{e}} \sum_{j=1}^{\infty} \frac{j^{d+1}}{j!} = B_{d+1} ,$$

by Equation (1), since $d + 1 > 0$. □

As an immediate consequence, we obtain:

Corollary 1. *Let $\mathcal{G}$ be the class of weighted singleton congestion games with unrestricted strategy sets and affine latency functions. Then*

$$\mathit{PoA}_{\mathsf{pure}}(\mathcal{G}) \geq 2 .$$

We close this section with a lower bound on the pure price of anarchy for weighted congestion games with linear latency functions.

Theorem 3. *Let $\mathcal{G}$ be the class of weighted singleton congestion games with unrestricted strategy sets and linear latency functions. Then*

$$\mathit{PoA}(\mathcal{G}) > 2.036 .$$

Proof. For some parameter $w \in \mathbb{R}_{>0}$ (to be determined later), define the weighted singleton congestion game $\Gamma(w)$ with 5 players, 5 resources and unrestricted strategy sets as follows:

- $w_1 = w$ and $w_i = 1$ for $i \in \{2, \ldots, 5\}$
- $f_1(x) = \frac{w}{w+4} \cdot x$ and $f_j(x) = x$ for $j \in \{2, \ldots, 5\}$

Define the pure strategy profile $\boldsymbol{s}$ where $s_i = i$ for all players $i \in [n]$. Then,

$$\mathsf{SC}(\Gamma(w), \boldsymbol{s}) = \frac{w^3}{w+4} + 4 \,.$$

Let $\mathbf{P}$ be the mixed strategy profile where:

- $p(1,1) = p$ and $p(1,j) = \frac{1-p}{4}$ for $j \in \{2, \ldots, 5\}$
- $p(i,1) = 1$ for $i \in \{2, \ldots, 5\}$

It is easy to check that $\mathbf{P}$ is a Nash equilibrium for $p \leq \frac{w^2-8w+16}{5w^2+4w}$. Now

$$\mathsf{SC}(\Gamma(w), \mathbf{P}) = p \cdot w \cdot (4+w) + (1-p) \cdot \left(\frac{16w}{w+4} + w^2\right) = p\frac{4w^2}{w+4} + \frac{16w}{w+4} + w^2 \,,$$

which is monotone increasing in p. So choose $p = \frac{w^2-8w+16}{5w^2+4w}$. Observe that for all $w > 0$,

$$\mathsf{PoA}(\mathcal{G}) \geq \frac{\mathsf{SC}(\Gamma(w), \mathbf{P})}{\mathsf{SC}(\Gamma(w), \boldsymbol{s})} \,.$$

Choosing $w = 3.258$ yields the claimed lower bound. □

4 Restricted Strategy Sets

It is known [1] that the price of anarchy for general weighted congestion games with polynomial latency functions of maximum degree d is exactly Φ_d^{d+1}. The next theorem shows that asymptotically the lower bound is already achieved with singleton congestion games.

Theorem 4. *Let $\mathcal{G}$ be the class of weighted singleton congestion games with restricted strategy sets and polynomial latency functions of maximum degree d. Then*

$$\mathit{PoA}(\mathcal{G}) = \mathit{PoA}_{\mathsf{pure}}(\mathcal{G}) = \Phi_d^{d+1} \,.$$

Proof. The upper bound $\mathsf{PoA}(\mathcal{G}) \leq \Phi_d^{d+1}$ follows from [1], hence we only need to show the lower bound. For some parameter $n \in \mathbb{N}$, define the following weighted singleton congestion game $\Gamma(n)$ with n players and $n+1$ resources. The weight of player $i \in [n]$ is $w_i = \Phi_d^i$ and the latency function of resource $j \in [n+1]$ is

$$f_j(x) = \begin{cases} \Phi_d^{-(d+1)\cdot(n-1)} \cdot x^d & \text{if } j = n \\ \Phi_d^{-(d+1)\cdot j} \cdot x^d & \text{otherwise}\,. \end{cases}$$

Each player $i \in [n]$ only has two available resources in its strategy set: $S_i = \{i, i+1\}$.

Let $\boldsymbol{s} := (i)_{i=1}^n \in S$. Then $\boldsymbol{s}$ is a Nash Equilibrium: Clearly, player n experiences the same utility for both strategies n and $n+1$; moreover, for any player $i \in [n-1]$ we have

$$\mathsf{PC}_i(\boldsymbol{s}_{-i}, i+1) = f_{i+1}(w_i + w_{i+1}) = f_{i+1}(\Phi_d^i + \Phi_d^{i+1}) = \frac{(\Phi_d^i + \Phi_d^{i+1})^d}{\Phi_d^{(d+1)\cdot(i+1)}} = \frac{(\Phi_d^i(\Phi_d + 1))^d}{\Phi_d^{(d+1)\cdot(i+1)}}$$
$$= \frac{\Phi_d^{id}}{\Phi_d^{(d+1)\cdot i}} = f_i(\Phi_d^i) = f_i(w_i) = \mathsf{PC}_i(\boldsymbol{s}) \,.$$

Consequently,

$$\mathsf{SC}(\Gamma(n), \boldsymbol{s}) = \sum_{i=1}^n w_i \cdot f_i(w_i) = \sum_{i=1}^n \Phi_d^i \cdot \frac{\Phi_d^{id}}{\Phi_d^{(d+1)\cdot i}} = n \,.$$

Now let $\boldsymbol{s}^* := (i+1)_{i=1}^n \in S$. Then,

$$\mathsf{SC}(\Gamma(n), \boldsymbol{s}^*) = \sum_{i=1}^n w_i \cdot f_{i+1}(w_i) = \sum_{i=1}^{n-1} \Phi_d^i \cdot \frac{\Phi_d^{id}}{\Phi_d^{(d+1)\cdot(i+1)}} + \Phi_d^n \cdot \frac{\Phi_d^{nd}}{\Phi_d^{(d+1)\cdot n}} = (n-1) \cdot \frac{1}{\Phi_d^{d+1}} + 1 \,.$$

We get $\mathsf{PoA}(\mathcal{G}) \geq \sup_{n\in\mathbb{N}} \left\{ \frac{\mathsf{SC}(\Gamma(n),\boldsymbol{s})}{\mathsf{SC}(\Gamma(n),\boldsymbol{s}^*)} \right\} = \Phi_d^{d+1}$. □

Before we can show a corresponding result for the case of unweighted singleton congestion games, we state a simple technical lemma which has an easy proof.

Lemma 3. *For all $n \in \mathbb{N} \setminus \{1\}$ it holds that $n^n > (n+1)^{n-1}$.*

Theorem 5. *Let $\mathcal{G}$ be the class of unweighted singleton congestion games with restricted strategy sets and polynomial latency functions of maximum degree d. Moreover, let $\lambda = \lfloor \Phi_d \rfloor$. Then*

$$\mathit{PoA}(\mathcal{G}) = \mathit{PoA}_{\mathsf{pure}}(\mathcal{G}) = \Upsilon(d) = \frac{(\lambda+1)^{2d+1} - \lambda^{d+1}(\lambda+2)^d}{(\lambda+1)^{d+1} - (\lambda+2)^d + (\lambda+1)^d - \lambda^{d+1}} \,.$$

Proof. For some parameter $k \in \mathbb{N}$, define an unweighted singleton congestion game $\Gamma(k)$. For simplicity of description, we borrow the representation introduced by [5] which makes use of a "game graph" $G = (N, A)$: Resources correspond to nodes and players correspond to arcs. Every player has exactly two strategies, namely choosing one or the other of its adjacent nodes.

The game graph is a tree which is constructed as follows: At the root there is a complete $(d+1)$-ary tree with $k+1$ levels. Each leaf of this tree is then the root of a complete d-ary tree the leafs of which are again the root of a complete $(d-1)$-ary tree; and so on. This recursive definition stops with the unary trees. For an example of this construction, see Figure 1.

Altogether, the game graph consists of $(d+1) \cdot k + 1$ levels. We let level 0 denote the root level. Thus, clearly, the nodes on level $i \cdot k$, where $i \in [d]_0$, are the root of a complete $(d+1-i)$-ary subtree (as indicated by the hatched shapes).

For any resource on level $(d+1-i) \cdot k + j$, where $i \in [d+1]$ and $j \in [k-1]_0$, let the latency function be $f_{i,j} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$,

$$f_{i,j}(x) := \left[\prod_{l=i+1}^{d+1} \frac{l}{l+1} \right]^{d\cdot(k-1)} \cdot \left(\frac{i}{i+1} \right)^{dj} \cdot x^d \,.$$

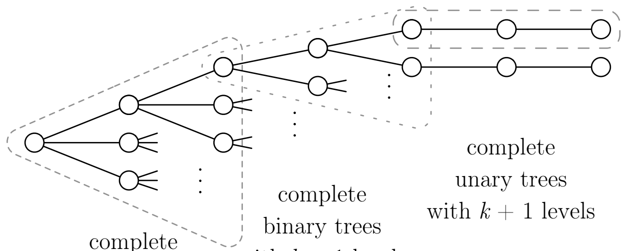

Fig. 1. The game graph for $d=2$ and $k=2$

Note that by construction there are exactly $\left[\prod_{l=i+1}^{d+1} l\right]^k \cdot i^j$ resources on level $(d+1-i)\cdot k+j$. The resources on level $(d+1)\cdot k$ have the same latency function $f_{0,0} := f_{1,k-1}$ as those on level $(d+1)\cdot k-1$.

Let $\boldsymbol{s}$ denote the strategy profile in $\Gamma(k)$ where each player uses the resource which is closer to the root. Since $f_{i+1,k-1} = f_{i,0}$ for $i \in [d+1]_0$ and $f_{i,j}(i) = f_{i,j+1}(i+1)$ for all $i \in [d+1]$ and $j \in [k-2]_0$, it is easy to verify that in $\boldsymbol{s}$ no player has an incentive to switch to its other strategy farther away from the root. Hence, $\boldsymbol{s}$ is a Nash equilibrium. We get

$$\mathsf{SC}(\boldsymbol{s}) = \sum_{i=1}^{d+1}\sum_{j=0}^{k-1}\left[\prod_{l=i+1}^{d+1} l^k\right]\cdot i^j \cdot i \cdot f_{i,j}(i) = \sum_{i=1}^{d+1}\sum_{j=0}^{k-1}\left[\prod_{l=i+1}^{d+1} l^k \cdot \left(\frac{l}{l+1}\right)^{d\cdot(k-1)}\right]\cdot\left(\frac{i}{i+1}\right)^{dj}\cdot i^{j+d+1}$$

$$= \sum_{i=1}^{d+1} i^{d+1}\cdot\left[\prod_{l=i+1}^{d+1} l\right]\cdot\left[\prod_{l=i+1}^{d+1}\frac{l^{d+1}}{(l+1)^d}\right]^{k-1}\cdot\sum_{j=0}^{k-1}\left(\frac{i^{d+1}}{(i+1)^d}\right)^j .$$

Since

$$\sum_{j=0}^{k-1}\left(\frac{i^{d+1}}{(i+1)^d}\right)^j = (i+1)^d\cdot\frac{\left(\frac{i^{d+1}}{(i+1)^d}\right)^k - 1}{i^{d+1}-(i+1)^d} = \frac{i^{d+1}}{i^{d+1}-(i+1)^d}\cdot\left(\frac{i^{d+1}}{(i+1)^d}\right)^{k-1} + \frac{(i+1)^d}{(i+1)^d - i^{d+1}} ,$$

$\mathsf{SC}(\boldsymbol{s})$ can be written as a weighted sum of terms raised to the power of $(k-1)$,

$$\mathsf{SC}(\boldsymbol{s})$$

$$= \sum_{i=2}^{d+1}\left(i^{d+1}\cdot\left[\prod_{l=i+1}^{d+1} l\right]\cdot\frac{i^{d+1}}{i^{d+1}-(i+1)^d} + (i-1)^{d+1}\cdot\left[\prod_{l=i}^{d+1} l\right]\cdot\frac{i^d}{i^d-(i-1)^{d+1}}\right)\cdot\left[\prod_{l=i}^{d+1}\frac{l^{d+1}}{(l+1)^d}\right]^{k-1}$$

$$+ \left[\prod_{l=2}^{d+1} l\right]\cdot\frac{1}{1-2^d}\cdot\left[\prod_{l=1}^{d+1}\frac{l^{d+1}}{(l+1)^d}\right]^{k-1} + (d+1)^{d+1}\cdot\frac{(d+2)^d}{(d+2)^d-(d+1)^{d+1}} .$$

Let $\boldsymbol{s}^*$ be the profile where each player uses the resource farther away from the root. Similarly to before, we get

$$\mathsf{SC}(\boldsymbol{s}^*) = \sum_{i=1}^{d+1}\sum_{j=0}^{k-1}\left[\prod_{l=i+1}^{d+1} l^k\right]\cdot i^j\cdot f_{i,j}(1) + \left[\prod_{l=1}^{d+1} l^k\right]\cdot f_{0,0}(1) - 1$$

$$= \sum_{i=1}^{d+1} \left[\prod_{l=i+1}^{d+1} l\right] \cdot \left[\prod_{l=i+1}^{d+1} \frac{l^{d+1}}{(l+1)^d}\right]^{k-1} \cdot \sum_{j=0}^{k-1} \left(\frac{i^{d+1}}{(i+1)^d}\right)^j + \left[\prod_{l=1}^{d+1} l\right] \cdot \left[\prod_{l=1}^{d+1} \frac{l^{d+1}}{(l+1)^d}\right]^{k-1} - 1$$

$$= \sum_{i=2}^{d+1} \left(\left[\prod_{l=i+1}^{d+1} l\right] \cdot \frac{i^{d+1}}{i^{d+1} - (i+1)^d} + \left[\prod_{l=i}^{d+1} l\right] \cdot \frac{i^d}{i^d - (i-1)^{d+1}} \right) \cdot \left[\prod_{l=i}^{d+1} \frac{l^{d+1}}{(l+1)^d}\right]^{k-1}$$

$$+ \left[\prod_{l=2}^{d+1} l\right] \cdot \frac{2 - 2^d}{1 - 2^d} \cdot \left[\prod_{l=1}^{d+1} \frac{l^{d+1}}{(l+1)^d}\right]^{k-1} + \frac{(d+2)^d}{(d+2)^d - (d+1)^{d+1}} - 1\,.$$

Consequently, the quotient $\frac{\mathsf{SC}(\boldsymbol{s})}{\mathsf{SC}(\boldsymbol{s}^*)}$ is of the form

$$\frac{\sum_{i=0}^{d+1} \beta_i \cdot \alpha_i^{k-1}}{\sum_{i=0}^{d+1} \gamma_i \cdot \alpha_i^{k-1}}$$

where $\beta_i, \gamma_i \in \mathbb{Q}$, $\alpha_0 = 1$, and $\alpha_i = \prod_{l=i}^{d+1} \frac{l^{d+1}}{(l+1)^d} = \frac{i^{d+1}}{(d+2)^d} \cdot \prod_{l=i+1}^{d+1} l$ for all $i \in [d+1]$. In order to find the largest α_i for $i \in [d+1]_0$, consider the following equivalencies: For all $i \in [d]$, we have

$$\alpha_{i+1} > \alpha_i \iff (i+1)^{d+1} \cdot \prod_{l=i+2}^{d+1} l > i^{d+1} \cdot \prod_{l=i+1}^{d+1} l = i^{d+1} \cdot (i+1) \cdot \prod_{l=i+2}^{d+1} l$$
$$\iff (i+1)^d > i^{d+1}\,.$$

Moreover, $\alpha_1 = \frac{(d+1)!}{(d+2)^d} < 1$ and $\alpha_{d+1} = \frac{(d+1)^{d+1}}{(d+2)^d} > 1$, where the last inequality is due to Lemma 3. Let $\lambda := \lfloor \Phi_d \rfloor$. Then, $(\lambda+1)^d > \lambda^{d+1}$ but $(\lambda+2)^d < (\lambda+1)^{d+1}$, so $\lambda \in [d]$. Hence, $\alpha_{\lambda+1}$ is maximal, i.e., $\alpha_{\lambda+1} > \alpha_i$ for all $i \in [d+1]_0$, $i \neq \lambda+1$. Using standard calculus we get

$$\lim_{k\to\infty} \frac{\sum_{i=0}^{d+1} \beta_i \cdot \alpha_i^{k-1}}{\sum_{i=0}^{d+1} \gamma_i \cdot \alpha_i^{k-1}} = \frac{\beta_{\lambda+1}}{\gamma_{\lambda+1}}\,.$$

Inserting for $\beta_{\lambda+1}$ and $\gamma_{\lambda+1}$ as well as subsequent canceling of common terms in numerator and denominator then yields

$$\lim_{k\to\infty} \frac{\mathsf{SC}(\Gamma(k), \boldsymbol{s})}{\mathsf{SC}(\Gamma(k), \boldsymbol{s}^*)} = \frac{\frac{(\lambda+1)^{2d+2}}{(\lambda+1)^{d+1} - (\lambda+2)^d} + \frac{\lambda^{d+1} \cdot (\lambda+1)^{d+1}}{(\lambda+1)^d - \lambda^{d+1}}}{\frac{(\lambda+1)^{d+1}}{(\lambda+1)^{d+1} - (\lambda+2)^d} + \frac{(\lambda+1)^{d+1}}{(\lambda+1)^d - \lambda^{d+1}}} = \frac{(\lambda+1)^{2d+1} - \lambda^{d+1} \cdot (\lambda+2)^d}{(\lambda+1)^{d+1} - (\lambda+2)^d + (\lambda+1)^d - \lambda^{d+1}}\,.$$

Note that the denominator is non-zero. The theorem follows. □

5 Conclusion

In this paper, we presented a collection of upper and lower bounds on the price of anarchy for singleton congestion games. In some cases we determined the exact value, while for other cases there is still a (small) gap between the upper and lower bounds. Closing these gaps – in particular those for weighted singleton congestion games with unrestricted strategy sets and linear latency functions – remains a challenging open problem that deserves further investigation.

We found it very surprising that both upper bounds on the price of anarchy from [1] – proved for general congestion games with polynomial latency functions – are already exact for the case of singleton strategy sets and pure Nash equilibria.

References

1. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. In B. Durand and W. Thomas, editors, *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS'06)*, Lecture Notes in Computer Science, Vol. 3884, Springer Verlag, pages 218–229, 2006.
2. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 57–66, 2005.
3. B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in Worst-Case Equilibria. In K. Jansen and R. Solis-Oba, editors, *Proceedings of the 1st International Workshop on Approximation and Online Algorithms (WAOA'03)*, Lecture Notes in Computer Science, Vol. 2909, Springer Verlag, pages 41–52, 2003.
4. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
5. I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight Bounds for Selfish and Greedy Load Balancing. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP'06), Part I*, Lecture Notes in Computer Science, Vol. 4051, Springer Verlag, pages 311–322, 2006.
6. G. Christodoulou and E. Koutsoupias. On The Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In S. Leonardi G. S. Brodal, editor, *Proceedings of the 13th Annual European Symposium on Algorithms (ESA'05)*, Lecture Notes in Computer Science, Vol. 3669 Springer Verlag, pages 59–70, 2005.
7. 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'05)*, pages 67–73, 2005.
8. A. Czumaj, P. Krysta, and B. Vöcking. Selfish Traffic Allocation for Server Farms. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC'02)*, pages 287–296, 2002.
9. A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. *ACM Transactions on Algorithms*, 3(1):Article No. 4, 2007.
10. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In P. Widmayer, F. T. Ruiz, R. M. Bueno, M. Hennessy, S. Eidenbenz, and R. Conejo, editors, *Proceedings of the 29th International Colloquium on Automata, Languages, and Programming (ICALP'02)*, Lecture Notes in Computer Science, Vol. 2380, Springer Verlag, pages 123–134, 2002.
11. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. In J. Diaz, J. Karhumäki, A. Lepistö, and D. Sannella, editors, *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming (ICALP'04)*, Lecture Notes in Computer Science, Vol. 3142, Springer Verlag, pages 593–605, 2004.
12. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing Nash Equilibria for Scheduling on Restricted Parallel Links. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC'04)*, pages 613–622, 2004.
13. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The Price of Anarchy for Restricted Parallel Links. *Parallel Processing Letters*, 16(1):117–131, 2006.
14. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. In J. Diaz, J. Karhumäki, A. Lepistö, and D. Sannella, editors, *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming (ICALP'04)*, Lecture Notes in Computer Science, Vol. 3142, Springer Verlag, pages 645–657, 2004.
15. M. Gairing, B. Monien, and K. Tiemann. Selfish Routing with Incomplete Information. In *Proceedings of the 17th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA'05)*, pages 203–212, 2005. Also accepted to *Theory of Computing Systems*.
16. M. Gairing, B. Monien, and K. Tiemann. Routing (Un-) Splittable Flow in Games with Player-Specific Linear Latency Functions. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP'06), Part I*, Lecture Notes in Computer Science, Vol. 4051, Springer Verlag, pages 501–512, 2006.
17. E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.
18. E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In C. Meinel and S. Tison, editors, *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, Lecture Notes in Computer Science, Vol. 1563, Springer Verlag, pages 404–413, 1999.
19. T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. In V. Diekert and M. Habib, editors, *Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science (STACS'04)*, Lecture Notes in Computer Science, Vol. 2996, Springer Verlag, pages 547–558, 2004.
20. I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

21. D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.
22. J. F. Nash. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences of the United States of America*, 36:48–49, 1950.
23. J. F. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.
24. C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 749–753, 2001.
25. R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
26. T. Roughgarden and É. Tardos. How Bad Is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.
27. S. Suri, C. D. Tóth, and Y. Zhou. Selfish Load Balancing and Atomic Congestion Games. In *Proceedings of the 16th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA'04)*, pages 188–195, 2004.
28. V. Vazirani. *Approximation Algorithms*. Springer Verlag, 2001.
29. H.S. Wilf. *Generatingfunctionology*. Academic Press, 1994.