

Exact Price of Anarchy for Polynomial Congestion Games ^{*}

Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann

Department of Computer Science, Electrical Engineering and Mathematics
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
{sebaland,masa,gairing,bm,fschopp}@uni-paderborn.de

Abstract. We show exact values for the price of anarchy of weighted and unweighted congestion games with polynomial latency functions. The given values also hold for weighted and unweighted *network* congestion games.

1 Introduction

Motivation and Framework. Large scale communication networks, like e.g. the internet, often lack a central regulation for several reasons: The size of the network may be too large, or the users may be free to act according to their private interests. Even cooperation among the users may be impossible due to the fact that users may not even know each other. Such an environment—where users neither obey some central control instance nor cooperate with each other—can be modeled as a *non-cooperative game* [18].

One of the most widely used solution concepts for non-cooperative games is the concept of *Nash equilibrium*. A Nash equilibrium is a state in which no player can improve his objective by unilaterally changing his strategy. A Nash equilibrium is called *pure* if all players choose a pure strategy, and *mixed* if players choose probability distributions over strategies.

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 players share a private cost function, defined as the sum (over their chosen resources) of functions in the number of players sharing this resource. Later Milchtaich [20] considered *weighted congestion games* as an extension to congestion games in which the players have weights and thus different influence on the congestion of the resources. Weighted congestion games provide us with a general framework for modeling any kind of non-cooperative resource sharing problem. A typical resource sharing problem is that of routing. In a routing game the strategy sets of the players correspond to paths in a network. Routing games where the demand of the players cannot be split among multiple paths are also called (*weighted*) *network congestion games*. Another model for selfish routing—the so called *Wardrop* model—was already

^{*} This work has been partially supported by the DFG-SFB 376 and by the European Union within the 6th Framework Programme under contract 001907 (DELIS).

studied in the 1950's (see e.g. [3, 29]) in the context of road traffic systems, where traffic flows can be split arbitrarily. The Wardrop model can be seen as a special network congestion game with infinitely many players each carrying a negligible demand.

In order to measure the degradation of social welfare due to the selfish behavior of the players, Koutsoupias and Papadimitriou [16] introduced a global objective function, usually coined as *social cost*. They defined the *price of anarchy*, also called coordination ratio, as the worst-case ratio between the value of social cost in a Nash equilibrium and that of some social optimum. Thus, the price of anarchy measures the extent to which non-cooperation approximates cooperation. The price of anarchy directly depends on the definition of social cost. Koutsoupias and Papadimitriou [16] considered a very simple weighted network congestion game on parallel links, now known as KP-model. For this model they defined the social cost as the expected maximum latency. For the Wardrop model, Roughgarden and Tardos [28] considered social cost as the *total latency*, which is a measure for the (weighted) total travel time. Awerbuch et al. [1] and Christodoulou and Koutsoupias [5] considered the total latency for congestion games with a finite number of players with non-negligible demands. In this setting, they show asymptotic bounds on the price of anarchy for weighted (and unweighted) congestion games with polynomial latency (cost) functions. Here, all polynomials are of maximum degree d and have non-negative coefficients. For the case of linear latency functions they give exact bounds on the price of anarchy.

Contribution and Comparison. In this work we prove *exact* bounds on the price of anarchy for unweighted and weighted congestion games with polynomial latency functions. We use the total latency as social cost measure. This improves on results by Awerbuch et al. [1] and Christodoulou and Koutsoupias [5], where non-matching upper and lower bounds are given.

We now describe our findings in more detail.

- For *unweighted congestion games* we show that the price of anarchy (PoA) is exactly

$$\text{PoA} = \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}},$$

where $k = \lfloor \Phi_d \rfloor$ and Φ_d is a natural generalization of the golden ratio to larger dimensions such that Φ_d is the solution to $(\Phi_d + 1)^d = \Phi_d^{d+1}$. Prior to this paper the best known upper and lower bounds were shown to be of the form $d^{d(1-o(1))}$ [5]. However, the term $o(1)$ still hides a gap between the upper and the lower bound.

- For *weighted congestion games* we show that the price of anarchy (PoA) is exactly

$$\text{PoA} = \Phi_d^{d+1}.$$

This result closes the gap between the so far best upper and lower bounds of $O(2^d d^{d+1})$ and $\Omega(d^{d/2})$ from [1].

We show that the above values on the price of anarchy also hold for the subclasses of unweighted and weighted network congestion games.

For our upper bounds we use a similar analysis as in [5]. The core of our analysis is to determine parameters c_1 and c_2 such that

$$y \cdot f(x+1) \leq c_1 \cdot x \cdot f(x) + c_2 \cdot y \cdot f(y) \quad (1)$$

for all polynomial latency functions of maximum degree d and for all reals $x, y \geq 0$. For the case of unweighted demands it suffices to show (1) for all integers x, y . In order to prove their upper bound Christodoulou and Koutsoupias [5] looked at (1) with $c_1 = \frac{1}{2}$ and gave an asymptotic estimate for c_2 . In our analysis we optimize both parameters c_1, c_2 . This optimization process requires new ideas and is non-trivial.

Table 1 shows a numerical comparison of our bounds with the previous results of Awerbuch et al. [1] and Christodoulou and Koutsoupias [5].

For $d \geq 2$, the table only gives the respective lower bounds that are given in the cited works (before any estimates are applied). Values in parentheses denote cases in which the bound for linear functions is better than the general case.

In [1, Theorem 4.3], a construction scheme for networks is described with price of anarchy approximating $\frac{1}{e} \sum_{k=1}^{\infty} \frac{k^d}{k!}$ which yields the d -th Bell number. In [5, Theorem 10], a network with price of anarchy $\frac{(N-1)^{d+2}}{N}$ is given, with N being the largest integer for which $(N-1)^{d+2} \leq N^d$ holds.

The column with the upper bound from [5] is computed by using (1) with $c_1 = \frac{1}{2}$ and optimizing c_2 with help of our analysis. Thus, the column shows the best possible bounds that can be shown with $c_1 = \frac{1}{2}$.

d	Φ_d	unweighted PoA			weighted PoA	
		Our exact result	Upper Bound [5]	Lower bound [5]	Our exact result	Lower bound [1]
1	1.618	2.5	2.5	2.5	2.618	2.618
2	2.148	9,583	10	(2.5)	9,909	(2.618)
3	2.630	41.54	47	(2.5)	47.82	5
4	3.080	267.6	269	21.33	277.0	15
5	3.506	1,514	2,154	42.67	1,858	52
6	3.915	12,345	15,187	85.33	14,099	203
7	4.309	98,734	169,247	170.7	118,926	877
8	4.692	802,603	1,451,906	14,762	1,101,126	4,140
9	5.064	10,540,286	20,241,038	44,287	11,079,429	21,147
10	5.427	88,562,706	202,153,442	132,860	120,180,803	115,975

Table 1. Comparison of our results to [5] and [1]

Related Work. The papers most closely related to our work are those of Awerbuch et al. [1] and Christodoulou and Koutsoupias [5, 4]. For (unweighted) congestion games and social cost defined as average private cost (which in this case is the same as total latency) it was shown that the price of anarchy of pure Nash equilibria is $\frac{5}{2}$ for linear latency functions and $d^{\Theta(d)}$ for polynomial latency functions of maximum degree d [1, 5]. The bound of $\frac{5}{2}$ for linear latency function also holds for the correlated and thus also for the mixed price of anarchy [4]. For *weighted* congestion games the mixed price of anarchy for total latency is $\frac{3+\sqrt{5}}{2}$ for linear latency functions and $d^{\Theta(d)}$ for polynomial latency functions [1].

The *price of anarchy* [24], also known as *coordination ratio*, was first introduced and studied by Koutsoupias and Papadimitriou [16]. As a starting point of their investigation they considered a simple weighted congestion game on parallel links, now known as KP-model. In the KP-model latency functions are linear and social cost is defined as the maximum expected congestion on a link. In this setting, there exist *tight* bounds on the price of anarchy of $\Theta(\frac{\log m}{\log \log m})$ for identical links [7, 15] and $\Theta(\frac{\log m}{\log \log \log m})$ [7] for

related links. The price of anarchy has also been studied for variations of the KP-model, namely for non-linear latency functions [6, 12], for the case of restricted strategy sets [2, 10], for the case of incomplete information [14] and for different social cost measures [11, 17]. In particular Lücking et al. [17] study the total latency (they call it quadratic social cost) for routing games on parallel links with linear latency functions. For this model they show that the price of anarchy is exactly $\frac{4}{3}$ for case of identical player weights and $\frac{9}{8}$ for the case of identical links and arbitrary player weights.

The class of *congestion games* was introduced by Rosenthal [25] and extensively studied afterwards (see e.g. [8, 20, 21]). In Rosenthal's model the strategy of each player is a subset of resources. Resource utility functions can be arbitrary but they only depend on the number of players sharing the same resource. Rosenthal showed that such games always admit a pure Nash equilibrium using a potential function. Monderer and Shapley [21] characterize games that possess a potential function as potential games and show their relation to congestion games. Milchtaich [20] considers weighted congestion games with player specific payoff functions and shows that these games do not admit a pure Nash equilibrium in general. Fotakis et al. [8, 9] consider the price of anarchy for symmetric weighted network congestion games in layered networks [8] and for symmetric (unweighted) network congestion games in general networks [9]. In both cases they define social cost as expected maximum latency. For a survey on weighted congestion games we refer to [13].

Inspired by the arisen interest in the price of anarchy Roughgarden and Tardos [28] re-investigated the Wardrop model and used the *total latency* as a social cost measure. In this context the price of anarchy was shown to be $\frac{4}{3}$ for linear latency functions [28] and $\Theta(\frac{d}{\log d})$ [26] for polynomial latency functions of maximum degree d . An overview on results for this model can be found in the recent book of Roughgarden [27].

Roadmap. The rest of this paper is organized as follows. In Section 2 we give an exact definition of weighted congestion games. We present exact bounds on the price of anarchy for unweighted congestion games in Section 3 and for weighted congestion games in Section 4. Due to lack of space we omit some of the proofs.

2 Notations

General. For all integers $k \geq 0$, we denote $[k] = \{1, \dots, k\}$, $[k]_0 = \{0, \dots, k\}$. For all integers $d > 0$, let $\Phi_d \in \mathbb{R}^+$ denote the number for which $(\Phi_d + 1)^d = \Phi_d^{d+1}$. Clearly, Φ_1 coincides with the *golden ratio*. Thus, Φ_d is a natural generalization of the golden ratio to larger dimensions.

Weighted Congestion Games. A *weighted congestion game* Γ is a tuple $\Gamma = (n, E, (w_i)_{i \in [n]}, (S_i)_{i \in [n]}, (f_e)_{e \in E})$. Here, n is the number of *players* (or *users*) and E is the finite set of *resources*. For every player $i \in [n]$, $w_i \in \mathbb{R}^+$ is the *weight* and $S_i \subseteq 2^E$ is the *strategy set* of player i . Denote $S = S_1 \times \dots \times S_n$ and $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$. For every resource $e \in E$, the *latency function* $f_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ describes the *latency* on resource e . We consider only *polynomial latency functions* with maximum degree d and non-negative coefficients, that is for all $e \in E$ 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 $j \in [d]_0$.

In a (unweighted) *congestion game*, the weights of all players are equal. Thus, the private cost of a player only depends on the *number* of players choosing the same resources.

Strategies and Strategy Profiles. A *pure strategy* for player $i \in [n]$ is some specific $s_i \in S_i$ 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 $\mathbf{s} = (s_1, \dots, s_n) \in S$ whereas a *mixed strategy profile* $\mathbf{P} = (P_1, \dots, P_n)$ is represented by an n -tuple of mixed strategies. For a mixed strategy profile \mathbf{P} denote by $p(\mathbf{s}) = \prod_{i \in [n]} p(i, s_i)$ the probability that the players choose the pure strategy profile $\mathbf{s} = (s_1, \dots, s_n)$. Following standard game theory notation, we denote $\mathbf{P}_{-i} = (P_1, \dots, P_{i-1}, P_{i+1}, \dots, P_n)$ as the (mixed) strategy profile of all players except player i and (\mathbf{P}_{-i}, Q_i) as the strategy profile that results from \mathbf{P} if player i deviates to strategy Q_i .

Private Cost. Fix any pure strategy profile \mathbf{s} , and denote by $l_e(\mathbf{s}) = \sum_{i \in [n], s_i \ni e} w_i$ the *load* on resource $e \in E$. The *private cost* of player $i \in [n]$ in a pure strategy profile \mathbf{s} is defined by $\text{PC}_i(\mathbf{s}) = \sum_{e \in s_i} f_e(l_e(\mathbf{s}))$. For a mixed strategy profile \mathbf{P} , the *private cost* of player $i \in [n]$ is

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

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

$$\begin{aligned} \text{SC}(\mathbf{P}) &= \sum_{\mathbf{s} \in S} p(\mathbf{s}) \sum_{e \in E} l_e(\mathbf{s}) \cdot f_e(l_e(\mathbf{s})) \\ &= \sum_{\mathbf{s} \in S} p(\mathbf{s}) \sum_{i \in [n]} \sum_{e \in s_i} w_i \cdot f_e(l_e(\mathbf{s})) \\ &= \sum_{i \in [n]} w_i \cdot \text{PC}_i(\mathbf{P}). \end{aligned}$$

The *optimum* associated with a weighted congestion game is defined by $\text{OPT} = \min_{\mathbf{P}} \text{SC}(\mathbf{P})$.

Nash Equilibria and Price of Anarchy. We are interested in a special class of (mixed) strategy profiles called Nash equilibria [22, 23] that we describe here. Given a weighted congestion game and an associated mixed strategy profile \mathbf{P} , a player $i \in [n]$ is *satisfied* if he can not improve his private cost by unilaterally changing his 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, $\text{PC}_i(\mathbf{P}) \leq \text{PC}_i(\mathbf{P}_{-i}, s_i)$ 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 . Depending on the type of strategy profile, we differ between *pure* and *mixed* Nash equilibria.

The *price of anarchy*, also called *coordination ratio* and denoted PoA, is the maximum value, over all instances Γ and Nash equilibria \mathbf{P} , of the ratio $\frac{\text{SC}(\mathbf{P})}{\text{OPT}}$.

3 Price of Anarchy for Unweighted Congestion Games

In this section, we prove the exact value for the price of anarchy of unweighted congestion games with polynomial latency functions. We start with two technical lemmas which are crucial for determining c_1 and c_2 in (1) and thus for proving the upper bound Theorem 1. In Theorem 2 we give a matching lower bound which also holds for unweighted network congestion games (Corollary 1).

Lemma 1. *Let $0 \leq c < 1$ and $d \in \mathbb{N}_0$ then*

$$\max_{x \in \mathbb{N}_0, y \in \mathbb{N}} \left\{ \left(\frac{x+1}{y} \right)^d - c \cdot \left(\frac{x}{y} \right)^{d+1} \right\} = \max_{x \in \mathbb{N}_0} \left\{ (x+1)^d - c \cdot x^{d+1} \right\}.$$

Lemma 2. *Let $d \in \mathbb{N}$ and*

$$\mathcal{F}_d = \{g_r^{(d)} : \mathbb{R} \rightarrow \mathbb{R} \mid g_r^{(d)}(x) = (r+1)^d - x \cdot r^{d+1}, r \in \mathbb{R}_{\geq 0}\}$$

be an infinite set of linear functions. Furthermore, let $\gamma(s, t)$ for $s, t \in \mathbb{R}_{\geq 0}$ and $s \neq t$ denote the intersection abscissa of $g_s^{(d)}$ and $g_t^{(d)}$. Then it holds for any $s, t, u \in \mathbb{R}_{\geq 0}$ with $s < t < u$ that $\gamma(s, t) > \gamma(s, u)$ and $\gamma(u, s) > \gamma(u, t)$.

Theorem 1. *For unweighted congestion games with polynomial latency functions of maximum degree d and non-negative coefficients, we have*

$$\text{PoA} \leq \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}}, \text{ where } k = \lfloor \Phi_d \rfloor.$$

Proof. Let $\mathbf{P} = (P_1, \dots, P_n)$ be a (mixed) Nash equilibrium and let $\mathbf{Q} = (Q_1, \dots, Q_n)$ be a pure strategy profile with optimum social cost. Since \mathbf{P} is a Nash equilibrium, player $i \in [n]$ cannot improve by switching from strategy P_i to strategy Q_i . Thus,

$$\begin{aligned} \text{PC}_i(\mathbf{P}) &= \sum_{s \in S} p(s) \sum_{e \in s_i} f_e(l_e(s)) \leq \text{PC}_i(\mathbf{P}_{-i}, Q_i) \\ &= \sum_{s \in S} p(s) \left[\sum_{e \in Q_i \cap s_i} f_e(l_e(s)) + \sum_{e \in Q_i \setminus s_i} f_e(l_e(s) + 1) \right] \\ &\leq \sum_{s \in S} p(s) \sum_{e \in Q_i} f_e(l_e(s) + 1). \end{aligned}$$

Summing up over all players $i \in [n]$ yields

$$\begin{aligned} \text{SC}(\mathbf{P}) &= \sum_{i \in [n]} \sum_{s \in S} p(s) \sum_{e \in s_i} f_e(l_e(s)) \leq \sum_{i \in [n]} \sum_{s \in S} p(s) \sum_{e \in Q_i} f_e(l_e(s) + 1) \\ &= \sum_{s \in S} p(s) \sum_{e \in E} l_e(\mathbf{Q}) \cdot f_e(l_e(s) + 1). \end{aligned}$$

Now, $l_e(\mathbf{Q})$ and $l_e(s)$ are both integer, since \mathbf{Q} and s are both pure strategy profiles. Thus, by choosing c_1, c_2 such that

$$y \cdot f(x+1) \leq c_1 \cdot x \cdot f(x) + c_2 \cdot y \cdot f(y) \quad (2)$$

for all polynomials f with maximum degree d and non-negative coefficients and for all $x, y \in \mathbb{N}_0$, we get

$$\text{SC}(\mathbf{P}) \leq \sum_{s \in S} p(s) \sum_{e \in E} [c_1 l_e(s) f_e(l_e(s)) + c_2 l_e(\mathbf{Q}) f_e(l_e(\mathbf{Q}))] = c_1 \cdot \text{SC}(\mathbf{P}) + c_2 \cdot \text{SC}(\mathbf{Q}).$$

With $c_1 < 1$ it follows that $\frac{\text{SC}(\mathbf{P})}{\text{SC}(\mathbf{Q})} \leq \frac{c_2}{1-c_1}$. Since \mathbf{P} is an arbitrary (mixed) Nash equilibrium we get

$$\text{PoA} \leq \frac{c_2}{1-c_1}. \quad (3)$$

In fact, c_1 and c_2 depend on the maximum degree d , however, for the sake of readability we omit this dependence in our notation.

We will now show how to determine constants c_1 and c_2 such that Inequality (2) holds and such that the resulting upper bound of $\frac{c_2}{1-c_1}$ is minimal. To do so, we will first show that it suffices to consider Inequality (2) with $y = 1$ and $f(x) = x^d$.

Since f is a polynomial of maximum degree d with non-negative coefficients, it is sufficient to determine c_1 and c_2 that fulfill (2) for $f(x) = x^r$ for all integers $0 \leq r \leq d$.

So let $f(x) = x^r$ for some $0 \leq r \leq d$. In this case (2) reduces to

$$y \cdot (x+1)^r \leq c_1 \cdot x^{r+1} + c_2 \cdot y^{r+1}. \quad (4)$$

For any given constant $0 \leq c_1 < 1$ let $c_2(r, c_1)$ be the minimum value for c_2 such that (4) holds, that is

$$c_2(r, c_1) = \max_{x \in \mathbb{N}_0, y \in \mathbb{N}} \left\{ \frac{y(x+1)^r - c_1 \cdot x^{r+1}}{y^{r+1}} \right\} = \max_{x \in \mathbb{N}_0, y \in \mathbb{N}} \left\{ \left(\frac{x+1}{y} \right)^r - c_1 \cdot \left(\frac{x}{y} \right)^{r+1} \right\}.$$

Note that (4) holds for any c_2 when $y = 0$. By Lemma 1 we have

$$c_2(r, c_1) = \max_{x \in \mathbb{N}_0} \{ (x+1)^r - c_1 \cdot x^{r+1} \}. \quad (5)$$

Now, $c_2(r, c_1)$ is the maximum of infinitely many linear functions in c_1 ; one for each $x \in \mathbb{N}_0$. Denote \mathcal{F}_r as the (infinite) set of linear functions defining $c_2(r, c_1)$:

$$\mathcal{F}_r := \{g_x^{(r)} : (0, 1) \rightarrow \mathbb{R} \mid g_x^{(r)}(c_1) = (x+1)^r - c_1 \cdot x^{r+1}, x \in \mathbb{N}_0\}$$

For the partial derivative of any function $(x, r, c_1) \mapsto g_x^{(r)}(c_1)$ we get

$$\begin{aligned} \frac{\partial((x+1)^r - c_1 \cdot x^{r+1})}{\partial r} &= (x+1)^r \cdot \ln(x+1) - c_1 \cdot x^{r+1} \cdot \ln(x) \\ &> \ln(x+1) [(x+1)^r - c_1 \cdot x^{r+1}] \geq 0, \end{aligned}$$

for $(x+1)^r - c_1 \cdot x^{r+1} \geq 0$, that is, for the positive range of the chosen function from \mathcal{F}_r . Thus, the positive range of $(x+1)^d - c_1 \cdot x^{d+1}$ dominates the positive range of $(x+1)^r - c_1 \cdot x^{r+1}$ for all $0 \leq r \leq d$. Since $c_2(r, c_1) > 0$ for all $0 \leq r \leq d$, it follows that $c_2(d, c_1) \geq c_2(r, c_1)$, for all $0 \leq r \leq d$. Thus, without loss of generality, we may assume that $f(x) = x^d$.

For $s, t \in \mathbb{R}_{\geq 0}$ and $s \neq t$ define $\gamma(s, t)$ as the intersection abscissa of $g_s^{(d)}$ and $g_t^{(d)}$ (as in Lemma 2). Now consider the intersection of the two functions $g_v^{(d)}$ and $g_{v+1}^{(d)}$ from \mathcal{F}_d for some $v \in \mathbb{N}$. We show that this intersection lies above all other functions from \mathcal{F}_d .

- First consider any function $g_z^{(d)}$ with $z > v + 1$. We have $g_z^{(d)}(0) > g_{v+1}^{(d)}(0) > g_v^{(d)}(0)$. Furthermore, by Lemma 2 we get $\gamma(v, z) < \gamma(v, v + 1)$. It follows that $g_v^{(d)}(\gamma(v, v + 1)) > g_z^{(d)}(\gamma(v, v + 1))$.
- Now consider any function $g_z^{(d)}$ with $z < v$. We have $g_{v+1}^{(d)}(0) > g_v^{(d)}(0) > g_z^{(d)}(0)$. Furthermore, by Lemma 2 we get $\gamma(v, z) > \gamma(v, v + 1)$. Again, it follows that $g_v^{(d)}(\gamma(v, v + 1)) > g_z^{(d)}(\gamma(v, v + 1))$.

Thus, all intersections of two consecutive linear functions from \mathcal{F}_d lie on $c_2(d, c_1)$.

By (3), any point that lies on $c_2(d, c_1)$ gives an upper bound on PoA. Let k be the largest integer such that $(k + 1)^d \geq k^{d+1}$, that is $k = \lfloor \Phi_d \rfloor$. Then $(k + 2)^d < (k + 1)^{d+1}$. Choose c_1 and c_2 at the intersection of the two lines from \mathcal{F}_d with $x = k$ and $x = k + 1$, that is $c_2 = (k + 1)^d - c_1 \cdot k^{d+1}$ and $c_2 = (k + 2)^d - c_1 \cdot (k + 1)^{d+1}$. Thus,

$$c_1 = \frac{(k + 2)^d - (k + 1)^d}{(k + 1)^{d+1} - k^{d+1}} \quad \text{and} \quad c_2 = \frac{(k + 1)^{2d+1} - (k + 2)^d \cdot k^{d+1}}{(k + 1)^{d+1} - k^{d+1}}.$$

Note that by the choice of k we have $0 < c_1 < 1$.

It follows that

$$\text{PoA} \leq \frac{(k + 1)^{2d+1} - k^{d+1}(k + 2)^d}{(k + 1)^{d+1} - (k + 2)^d + (k + 1)^d - k^{d+1}}.$$

This completes the proof of the theorem. \square

Theorem 2. *For unweighted congestion games with polynomial latency functions of maximum degree d and non-negative coefficients, we have*

$$\text{PoA} \geq \frac{(k + 1)^{2d+1} - k^{d+1}(k + 2)^d}{(k + 1)^{d+1} - (k + 2)^d + (k + 1)^d - k^{d+1}}, \quad \text{where } k = \lfloor \Phi_d \rfloor.$$

Proof. Given the maximum degree $d \in \mathbb{N}$ for the polynomial latency functions, we construct a congestion game for $n \geq k + 2$ players and $|E| = 2n$ facilities.

We divide the set E into two subsets $E_1 := \{g_1, \dots, g_n\}$ and $E_2 := \{h_1, \dots, h_n\}$. Each player i has two pure strategies, $P_i := \{g_{i+1}, \dots, g_{i+k}, h_{i+1}, \dots, h_{i+k+1}\}$ and $Q_i := \{g_i, h_i\}$ where $g_j := g_{j-n}$ and $h_j := h_{j-n}$ for $j > n$. I. e. $S_i = \{Q_i, P_i\}$.

Each of the facilities in E_1 share the latency function $x \mapsto ax^d$ for an $a \in \mathbb{R}_{>0}$ (yet to be determined) whereas the facilities in E_2 have latency $x \mapsto x^d$.

Obviously, the optimal allocation \mathbf{Q} is for every player i to choose Q_i . Now we determine a value for a such that the allocation $\mathbf{P} := (P_1, \dots, P_n)$ becomes a Nash Equilibrium, i.e., each player i is satisfied with \mathbf{P} , that is $\text{PC}_i(\mathbf{P}) \leq \text{PC}_i(\mathbf{P}_{-i}, Q_i)$ for all $i \in [n]$, or equivalently $k \cdot a \cdot k^d + (k + 1) \cdot (k + 1)^d \leq a \cdot (k + 1)^d + (k + 2)^d$. Resolving to the coefficient a gives

$$a \geq \frac{(k + 1)^{d+1} - (k + 2)^d}{(k + 1)^d - k^{d+1}} > 0. \quad (6)$$

Because $(k + 1)^d \neq k^{d+1}$, due to either $k + 1$ or k being odd and the other being even, a is well defined and positive. Now since for any player i the private costs are $\text{PC}_i(\mathbf{Q}) = a + 1$ and $\text{PC}_i(\mathbf{P}) = a \cdot k^{d+1} + (k + 1)^{d+1}$, it follows that

$$\frac{\text{SC}(\mathbf{P})}{\text{SC}(\mathbf{Q})} = \frac{\sum_{i \in [n]} \text{PC}_i(\mathbf{P})}{\sum_{i \in [n]} \text{PC}_i(\mathbf{Q})} = \frac{a \cdot k^{d+1} + (k + 1)^{d+1}}{a + 1}. \quad (7)$$

Provided that $(k+1)^d \geq k^{d+1}$, it is not hard to see that (7) is monotonically decreasing in a . Thus, we assume equality in (6), which then gives

$$\text{PoA} \geq \frac{\text{SC}(\mathbf{P})}{\text{SC}(\mathbf{Q})} = \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}}.$$

This completes the proof of the theorem. \square

Corollary 1. *The lower bound in Theorem 2 on PoA also holds for unweighted network congestion games.*

4 Price of Anarchy for Weighted Congestion Games

In this section, we prove the exact value for the price of anarchy of weighted congestion games with polynomial latency functions. The proof of the upper bound in Theorem 3 has a similar structure as the one for the unweighted case (cf. Theorem 1). In Theorem 4 we give a matching lower bound which also holds for weighted network congestion games (Corollary 2). Corollary 3 shows the impact of player weights to the price of anarchy.

Theorem 3. *For weighted congestion games with polynomial latency functions of maximum degree d and non-negative coefficients we have $\text{PoA} \leq \Phi_d^{d+1}$.*

Theorem 4. *For weighted congestion games with polynomial latency functions of maximum degree d and non-negative coefficients, we have $\text{PoA} \geq \Phi_d^{d+1}$.*

Proof. Given the maximum degree $d \in \mathbb{N}$ for the polynomial latency functions, set $k \geq \max\{\binom{d}{\lfloor d/2 \rfloor}, 2\}$. Note, that $\binom{d}{\lfloor d/2 \rfloor} = \max_{j \in [d]_0} \binom{d}{j}$. We construct a congestion game for $n = (d+1) \cdot k$ players and $|E| = n$ facilities.

We divide the set E into $d+1$ partitions: For $i \in [d]_0$, let $E_i := \{g_{i,1}, \dots, g_{i,k}\}$, with each $g_{i,j}$ sharing the latency function $x \mapsto a_i \cdot x^d$. The values of the coefficients a_i will be determined later. For simplicity of notation, set $g_{i,j} := g_{i,j-k}$ for $j > k$ in the following.

Similarly, we partition the set of players $[n]$: For $i \in [d]_0$, let $N_i := \{u_{i,1}, \dots, u_{i,k}\}$. The weight of each player in set N_i is Φ_d^i , so $w_{u_{i,j}} = \Phi_d^i$ for all $i \in [d]_0, j \in [k]$.

Now, for every set N_i , each player $u_{i,j} \in N_i$ has exactly two strategies:

$$Q_{u_{i,j}} := \{g_{i,j}\} \quad \text{and} \quad P_{u_{i,j}} := \begin{cases} \{g_{d,j+1}, \dots, g_{d,j+\binom{d}{i}}, g_{i-1,j}\} & \text{for } i = 1 \text{ to } d \\ \{g_{d,j+1}\} & \text{for } i = 0 \end{cases}$$

Now let $\mathbf{Q} := (Q_1, \dots, Q_n)$ and $\mathbf{P} := (P_1, \dots, P_n)$ be strategy profiles. The facilities in each set E_i then have the following loads for \mathbf{Q} and \mathbf{P} :

		load on every facility $e \in E_i$	
i	$l_e(\mathbf{Q})$	$l_e(\mathbf{P})$	
d	Φ_d^d	$\sum_{l=0}^d \binom{d}{l} \Phi_d^l = (\Phi_d + 1)^d = \Phi_d^{d+1}$	
0 to $d-1$	Φ_d^i	Φ_d^{i+1}	

For \mathbf{P} to become a Nash Equilibrium, we need to fulfill the following Nash inequalities for each set N_i of players:

i	Nash inequality to fulfill
1 to d	$\text{PC}_{u_{i,j}}(\mathbf{P}) = \binom{d}{i} \cdot a_d \cdot (\Phi_d^{d+1})^d + a_{i-1} \cdot (\Phi_d^i)^d$ $\leq a_i \cdot (\Phi_d^{i+1} + \Phi_d^i)^d = \text{PC}_{u_{i,j}}(\mathbf{P}_{-u_{i,j}}, Q_{u_{i,j}})$
0	$\text{PC}_{u_{0,j}}(\mathbf{P}) = a_d \cdot (\Phi_d^{d+1})^d \leq a_0 \cdot (\Phi_d + 1)^d = \text{PC}_{u_{0,j}}(\mathbf{P}_{-u_{0,j}}, Q_{u_{0,j}})$

Replacing “ \leq ” by “ $=$ ” yields a homogeneous system of linear equations, i.e., the system $B_d \cdot a = 0$ where B_d is the following $(d+1) \times (d+1)$ matrix:

$$B_d = \begin{pmatrix} -\Phi_d^{d^2+d+1} + \Phi_d^{d^2+d} & \Phi_d^{d^2} & 0 & \dots & \dots & 0 \\ \binom{d}{d-1} \Phi_d^{d^2+d} & -\Phi_d^{d^2+1} & \ddots & & & \vdots \\ \vdots & 0 & \ddots & & & \vdots \\ \vdots & \vdots & \ddots & \ddots & & \vdots \\ \binom{d}{i} \Phi_d^{d^2+d} & 0 & \dots & 0 & -\Phi_d^{id+d+1} & \Phi_d^{id} & 0 & \dots & 0 \\ \vdots & \vdots & & & 0 & \ddots & \ddots & & \vdots \\ \vdots & \vdots & & & \vdots & & \ddots & & 0 \\ \vdots & \vdots & & & \vdots & & \ddots & & \Phi_d^d \\ \Phi_d^{d^2+d} & 0 & \dots & 0 & \dots & 0 & \dots & 0 & -\Phi_d^{d^2+1} \end{pmatrix} \quad (8)$$

and $a := (a_d \dots a_0)^t$. Obviously, a solution to this system fulfills the initial Nash inequalities. Note that

$$(\Phi_d^{i+1} + \Phi_d^i)^d = (\Phi_d^i)^d \cdot (\Phi_d + 1)^d = \Phi_d^{id+d+1}.$$

Claim. The $(d+1) \times (d+1)$ matrix B_d from (8) has rank d .

Proof. We use the well-known fact from linear algebra that if a matrix C results from another matrix D by adding a multiple of one row (or column) to another row (or column, respectively) then $\text{rank}(C) = \text{rank}(D)$.

Now consider the matrix C_d that results from adding row j multiplied by the factor Φ_d^{-1} to row $j-1$, sequentially done for $j = d+1, d, \dots, 2$. Obviously, C_d is a lower triangular matrix with nonzero elements only in the first column and on the principal diagonal.

For the top left element of C_d we get

$$-\Phi_d^{d^2+d+1} + \sum_{j=0}^d \binom{d}{j} \Phi_d^{d^2+j} = \Phi_d^{d^2} \cdot \left(-\Phi_d^{d+1} + \underbrace{\sum_{j=0}^d \binom{d}{j} \Phi_d^j}_{(\Phi_d+1)^d} \right) = 0.$$

Since all elements on the principal diagonal of C_d —with the just shown exception of the first one—are nonzero, it is easy to see that C_d (and thus also B_d) has rank d . \square

By the above claim it follows that the column vectors of B_d are linearly dependent and thus there are—with degree of freedom 1—infinately many linear combinations of them yielding 0. In other words, $B_d \cdot a = 0$ has a one-dimensional solution space.

We now show (by induction over i) that all coefficients a_i , $i \in [d]_0$ must have the same sign and thus we can always find a valid solution. From the last equality, for $i = 0$, we have that a_d and a_0 must have the same sign. Now for $i = 1, \dots, d - 1$, it follows that a_i must have the same sign as a_{i-1} and a_d , for $(\Phi_d^{d+1})^d$, $(\Phi_d^i)^d$, and $(\Phi_d^{i+1} + \Phi_d^i)^d$ are all positive.

Choosing $a \neq 0$ with all components being positive, all coefficients of the latency functions are positive. We get,

$$\text{PoA} \geq \frac{\text{SC}(\mathbf{P})}{\text{SC}(\mathbf{Q})} = \frac{k \cdot \sum_{i=0}^d a_i (\Phi_d^{i+1})^{d+1}}{k \cdot \sum_{i=0}^d a_i (\Phi_d^i)^{d+1}} = \Phi_d^{d+1}.$$

□

Corollary 2. *The lower bound in Theorem 4 on PoA also holds for weighted network congestion games.*

Corollary 3. *The exact price of anarchy for unweighted congestion games*

$$\text{PoA} = \frac{(k+1)^{2d+1} - k^{d+1}(k+2)^d}{(k+1)^{d+1} - (k+2)^d + (k+1)^d - k^{d+1}},$$

where $k = \lfloor \Phi_d \rfloor$, is bounded by $\lfloor \Phi_d \rfloor^{d+1} \leq \text{PoA} \leq \Phi_d^{d+1}$.

References

1. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 57–66, 2005.
2. B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in Worst-Case Equilibria. *Proc. of the 1st Int. Workshop on Approximation and Online Algorithms (WAOA'03)*, LNCS 2909, pages 41–52, 2003.
3. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
4. G. Christodoulou and E. Koutsoupias. On The Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In *Proc. of the 13th Annual European Symposium on Algorithms (ESA'05)*, LNCS 3669, pages 59–70, 2005.
5. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 67–73, 2005.
6. A. Czumaj, P. Krysta, and B. Vöcking. Selfish Traffic Allocation for Server Farms. In *Proc. of the 34th Annual ACM Symposium on Theory of Computing (STOC'02)*, pages 287–296, 2002.
7. A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*, pages 413–420, 2002. Also accepted to *Journal of Algorithms* as Special Issue of SODA'02.
8. D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *Proc. of the 31st Int. Colloquium on Automata, Languages, and Programming (ICALP'04)*, LNCS 3142, pages 593–605, 2004.

9. D. Fotakis, S. Kontogiannis, and P. Spirakis. Symmetry in Network Congestion Games: Pure Equilibria and Anarchy Cost. In *Proc. of the 3rd Int. Workshop on Approximation and Online Algorithms (WAOA'05)*, 2005.
10. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing Nash Equilibria for Scheduling on Restricted Parallel Links. In *Proc. of the 36th Annual ACM Symposium on Theory of Computing (STOC'04)*, pages 613–622, 2004.
11. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The Price of Anarchy for Polynomial Social Cost. *Proc. of the 29th Int. Symposium on Mathematical Foundations of Computer Science (MFCS'04)*, LNCS 3153, pages 574–585, 2004.
12. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. *Proc. of the 31st Int. Colloquium on Automata, Languages, and Programming (ICALP'04)*, LNCS 3142, pages 645–657, 2004.
13. M. Gairing, T. Lücking, B. Monien, and K. Tiemann. Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture. In *Proc. of the 32nd Int. Colloquium on Automata, Languages, and Programming (ICALP'05)*, LNCS 3850, pages 51–65, 2005.
14. M. Gairing, B. Monien, and K. Tiemann. Selfish Routing with Incomplete Information. In *Proc. of the 17th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA'05)*, pages 203–212, 2005.
15. E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.
16. E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. *Proc. of the 16th Int. Symposium on Theoretical Aspects of Computer Science (STACS'99)*, LNCS 1563, pages 404–413, 1999.
17. T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. *Proc. of the 21st Int. Symposium on Theoretical Aspects of Computer Science (STACS'04)*, LNCS 2996, pages 547–558, 2004.
18. A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
19. R. D. McKelvey and A. McLennan. Computation of Equilibria in Finite Games. *Handbook of Computational Economics*, 1996.
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. *Proc. 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 *Proc. 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. *Int. Journal of Game Theory*, 2:65–67, 1973.
26. T. Roughgarden. How Unfair is Optimal Routing. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*, pages 203–204, 2002.
27. T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
28. T. Roughgarden and É. Tardos. How Bad Is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.
29. J. G. Wardrop. Some Theoretical Aspects of Road Traffic Research. In *Proc. of the Institute of Civil Engineers, Pt. II, Vol. 1*, pages 325–378, 1952.