

The Price of Stability of Weighted Congestion Games

George Christodoulou¹


Department of Computer Science, University of Liverpool, Liverpool, UK
G.Christodoulou@liverpool.ac.uk

Martin Gairing

Department of Computer Science, University of Liverpool, Liverpool, UK
gairing@liverpool.ac.uk


Yiannis Giannakopoulos²

Department of Mathematics, TU Munich, Munich, Germany
Yiannis.Giannakopoulos@tum.de

 <https://orcid.org/0000-0003-2382-1779>

Paul G. Spirakis

Department of Computer Science, University of Liverpool, Liverpool, UK
Computer Engineering and Informatics Department, University of Patras, Patras, Greece
P.Spirakis@liverpool.ac.uk

 <https://orcid.org/0000-0001-5396-3749>

Abstract

We give exponential lower bounds on the Price of Stability (PoS) of weighted congestion games with polynomial cost functions. In particular, for any positive integer d we construct rather simple games with cost functions of degree at most d which have a PoS of at least $\Omega(\Phi_d)^{d+1}$, where $\Phi_d \sim d/\ln d$ is the unique positive root of equation $x^{d+1} = (x+1)^d$. This essentially closes the huge gap between $\Theta(d)$ and Φ_d^{d+1} and asymptotically matches the Price of Anarchy upper bound. We further show that the PoS remains exponential even for singleton games. More generally, we also provide a lower bound of $\Omega((1+1/\alpha)^d/d)$ on the PoS of α -approximate Nash equilibria, even for singleton games. All our lower bounds extend to network congestion games, and hold for mixed and correlated equilibria as well.

On the positive side, we give a general upper bound on the PoS of α -approximate Nash equilibria, which is sensitive to the range W of the player weights and the approximation parameter α . We do this by explicitly constructing a novel approximate potential function, based on Faulhaber's formula, that generalizes Rosenthal's potential in a continuous, analytic way. From the general theorem, we deduce two interesting corollaries. First, we derive the existence of an approximate pure Nash equilibrium with PoS at most $(d+3)/2$; the equilibrium's approximation parameter ranges from $\Theta(1)$ to $d+1$ in a smooth way with respect to W . Secondly, we show that for unweighted congestion games, the PoS of α -approximate Nash equilibria is at most $(d+1)/\alpha$.

2012 ACM Subject Classification Theory of computation \rightarrow Algorithmic game theory, Theory of computation \rightarrow Quality of equilibria, Theory of computation \rightarrow Network games

Keywords and phrases Congestion games, price of stability, Nash equilibrium, approximate equilibrium, potential games

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.150

¹ Supported by EPSRC grant EP/M008118/1.

² Supported by the Alexander von Humboldt Foundation with funds from the German Federal Ministry of Education and Research (BMBF).



© George Christodoulou, Martin Gairing, Yiannis Giannakopoulos, and Paul G. Spirakis; licensed under Creative Commons License CC-BY

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).

Editors: Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella; Article No. 150; pp. 150:1–150:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Related Version A full version of this paper is available at [26], <https://arxiv.org/abs/1802.09952>.

1 Introduction

In the last 20 years, a central strand of research within Algorithmic Game Theory has focused on understanding and quantifying the inefficiency of equilibria compared to centralized, optimal solutions. There are two standard concepts that measure this inefficiency. The Price of Anarchy (PoA) [34] which takes the worst-case perspective, compares the worst-case equilibrium with the system optimum. It is a very robust measure of performance. On the other hand, the Price of Stability (PoS) [44, 5], which is also the focus of this work, takes an optimistic perspective, and uses the best-case equilibrium for this comparison. The PoS is an appropriate concept to analyse the ideal solution that we would like our protocols to produce.

The initial set of problems that arose from the Price of Anarchy theory have now been resolved. The most rich and well-studied among these models are, arguably, the atomic and non-atomic variants of congestion games (see [38, Ch. 18] for a detailed discussion). This class of games is very descriptive and captures a large variety of scenarios where users compete for resources, most prominently routing games. The seminal work of Roughgarden and Tardos [42, 43] gave the answer for the non-atomic variant, where each player controls a negligible amount of traffic. Awerbuch et al. [6], Christodoulou and Koutsoupias [17] resolved the Price of Anarchy for atomic congestion games with affine latencies, generalized by Aland et al. [3] to polynomials; this led to the development of Roughgarden’s Smoothness Framework [41] which extended the bounds to general cost functions, but also distilled and formulated previous ideas to bound the Price of Anarchy in an elegant, unified framework. At the computational complexity front, we know that even for simple congestion games, finding a (pure) Nash equilibrium is a PLS-complete problem [21, 2].

Allowing the players to have different loads, gives rise to the class of *weighted* congestion games [40]; this is a natural and very important generalization of congestion games, with numerous applications in routing and scheduling. Unfortunately though, an immediate dichotomy between weighted and unweighted congestion games occurs: the former may *not* even have pure Nash equilibria [36, 25, 27, 30]; as a matter of fact, it is a strongly NP-hard problem to even determine if that’s the case [20]. Moreover, in such games there does not, in general, exist a potential function [37, 31], which is the main tool for proving equilibrium existence in the unweighted case.

As a result, a sharp contrast with respect to our understanding of the two aforementioned inefficiency notions arises. The Price of Anarchy has been studied in depth and general techniques for providing tight bounds are known. Moreover, the asymptotic behaviour of weighted and unweighted congestion games with respect to the Price of Anarchy is identical; it is $\Theta(d/\log d)^d$ for both classes when latencies are polynomials of degree at most d [3].

The situation for the Price of Stability though, is completely different. For unweighted games we have a good understanding³ and the values are much lower than the Price of

³ Much work has been done on the PoS for network design games, which is though not so closely related to our work. This problem was first studied by Anshelevich et al. [5] who showed a tight bound of H_n , the harmonic number of the number of players n , for directed networks. Finding tight bounds on undirected networks is still a long-standing open problem (see, e.g., [23, 9, 35]). Recently, Bilò et al. [10] (asymptotically) resolved the question for broadcast networks. For the weighted variant of this problem, Albers [4] showed a lower bound of $\Omega(\log W/\log \log W)$, where W is the sum of the players’

■ **Table 1** Previous results (*left table*) regarding the Price of Anarchy and Stability for unweighted and weighted congestion games, with polynomial latency functions of maximum degree d . Φ_d is the unique positive solution of $(x + 1)^d = x^{d+1}$ and $\Phi_d = \Theta(d/\log d)$. Tight answers were known for all settings, except for the Price of Stability of the weighted case where only trivial bounds existed. In this paper (*right table*) we essentially close this gap by showing a lower bound of $\Omega(\Phi_d)^{d+1}$ (Theorem 1), even for network games, which is exponential even for singleton games Theorem 5.

		PoA	PoS	PoS lower bound	
unweighted		$[\Phi_d]^{d+1}$ [3]	$\Theta(d)$ [15]	general	$\Omega(\Phi_d)^{d+1}$
	weighted	Φ_d^{d+1} [3]	$[\Theta(d), \Phi_d^{d+1}]$	singleton	$\Omega(2^d/d)$
				α -approximate equilibria	$\Omega((1 + 1/\alpha)^d/d)$

Anarchy values, and also *tight*; approximately 1.577 for affine functions [16, 14], and $\Theta(d)$ [15] for polynomials. For weighted games though there is a huge gap; the current state of the art lower bound is $\Theta(d)$ and the upper bound is $\Theta(d/\ln d)^d$. These previous results are summarized at the left of Table 1.

The main focus of this work is precisely to deal with this lack of understanding, and to determine the Price of Stability of weighted congestion games. What makes this problem challenging is that the only general known technique for showing upper bounds for the Price of Stability is the potential method, which is applicable only to potential games. In a nutshell, the idea of this method is to use the global minimizer of Rosenthal’s potential [39] as an equilibrium refinement. This equilibrium is also a pure Nash equilibrium and can serve as an upper bound of the Price of Stability. Interestingly, it turns out that, for several classes of potential games, this technique actually provides the tight answer (see for example [5, 16, 14, 15]). However, as already mentioned above, unlike their unweighted counterparts, weighted congestion games are not potential games;⁴ so, a completely fresh approach is required.

One way to override the aforementioned limitations of non-existence of pure Nash equilibria, but also their computational hardness, is to consider *approximate* equilibria. In this direction, Hansknecht et al. [29] have shown that $(d + 1)$ -approximate pure Nash equilibria always exist in weighted congestion games with polynomial latencies of maximum degree d , while, in the negative side, there exist games that do not have 1.153-approximate pure Nash equilibria. Notice here, that these results do not take into account computational complexity considerations; if we insist in polynomial-time algorithms for actually finding those equilibria, then the currently best approximation parameter becomes $d^{O(d)}$ [12, 13, 22].

1.1 Our Results

We provide lower and upper bounds on the Price of Stability for the class of weighted congestion games with polynomial latencies with nonnegative coefficients. We consider both exact and approximate equilibria. Our lower bounds are summarized at the right of Table 1.

weights. See [10] and references therein for a thorough discussion of those results.
⁴ For the special case of weighted congestion games with linear latency functions, a potential does exist [25] and this was used by [8] to provide a PoS upper bound of 2.

Lower Bound for Weighted Congestion Games. In our main result in Theorem 1, we resolve a long-standing open problem by providing almost tight bounds for the Price of Stability of weighted congestion games with polynomial latency functions. We construct an instance having a Price of Stability of $\Omega(\Phi_d)^{d+1}$, where d is the maximum degree of the latencies and $\Phi_d \sim \frac{d}{\ln d}$ is the unique positive solution of equation $(x+1)^d = x^{d+1}$.

This bound essentially closes the previously huge gap between $\Theta(d)$ and Φ_d^{d+1} for the PoS of weighted congestion games. The previously best lower and upper bounds were rather trivial: the lower bound corresponds to the PoS results of Christodoulou and Gairing [15] for the unweighted case (and thus, it is also a valid lower bound for the general weighted case as well) and the upper bound comes from the Price of Anarchy results of Aland et al. [3] (PoA, by definition, upper-bounds PoS).

We stress that, although as mentioned before, weighted congestion games do not always possess pure equilibria, our lower bound construction involves a *unique* equilibrium occurring by iteratively eliminating strongly dominated strategies. As a result, this lower bound holds not only for pure, but mixed and correlated equilibria as well.

Singleton Games. Next we switch to the class of singleton congestion games, where a pure strategy for each player is a *single* resource. This class is very well-studied as, on one hand, it abstracts scheduling environments, and on the other, it has very attractive equilibrium properties; unlike general weighted congestion games, there exists an (ordinal) lexicographic potential [24, 32], thus implying the existence of *pure* Nash equilibria. It is important to note that, the tight lower bounds for the Price of Anarchy of general weighted congestion games, hold also for the class of singleton games [14, 7, 11].

Nevertheless, even for this special class, we show in Theorem 5 an exponential lower bound of $\Omega(2^d/d)$. The previous best upper and lower bounds were the same as those of the general case, namely $\Theta(d)$ and Φ_d^{d+1} , respectively. As a matter of fact, this new lower bound comes as a corollary of a more general result that we show in Theorem 5, that extends to approximate equilibria and gives a lower bound of $\Omega((1+1/\alpha)^d/d)$ on the PoS of α -approximate equilibria, for any (multiplicative) approximation parameter $\alpha \in [1, d]$. Setting $\alpha = 1$ we recover the special case of exact equilibria and the aforementioned exponential lower bound on the standard, exact notion of the PoS. Notice here that, as we show in [26, Appendix D], the optimal solution (which, in general, is not an equilibrium) itself constitutes a $(d+1)$ -approximate equilibrium with a (trivially) optimal PoS of 1.

Positive Results for Approximate Equilibria. In light of the above results, in Section 4, we turn our attention to identifying environments with more structure or flexibility with respect to the underlying solution concept, for which we can hope for improved quality of equilibria. Both our lower bound constructions discussed above use players' weights that form a geometric sequence. In particular the ratio W of the largest over the smallest weight is equal to w^n (for some $w > 1$), which grows very large as the number of players $n \rightarrow \infty$. On the other hand, for games where the players have equal weights, i.e. $W = 1$, we know that the PoS is at most $d+1$. It is therefore natural to ask how the performance of the good equilibria captured by the notion of PoS varies with respect to W . In Theorem 9, we are able to give a general upper bound for α -approximate equilibria which is sensitive to this parameter W and to α . This general theorem has two immediate, interesting corollaries.

Firstly (Corollary 10), by allowing the ratio W to range in $[1, \infty)$, we derive the existence of an α -approximate pure Nash equilibrium with PoS at most $(d+3)/2$; the equilibrium's approximation parameter α ranges from $\Theta(1)$ to $d+1$ in a smooth way with respect to W .

This is of particular importance in settings where player weights are not very far away from each other (that is, W is small). Secondly (Corollary 11), by setting $W = 1$ and allowing α to range, we get an upper bound of $\frac{d+1}{\alpha}$ for the α -approximate PoS of *unweighted* congestion games which, to the best of our knowledge, was not known before, degrading gracefully from $d + 1$ (which is the actual PoS of exact equilibria in the unweighted case [15]) down to the optimal value of 1 if we allow $(d + 1)$ -approximate equilibria (which in fact can be achieved by the optimum solution itself; see [26, Appendix D]).

Our Techniques. An advantage of our main lower bound (Theorem 1) is the simplicity of the underlying construction, as well as its straightforward adaptation to network games (see Section 3.1.1)). However, fine-tuning the parameters of the game (player weights and latency functions), to ensure uniqueness of the equilibrium at the “bad” instance, was a technically involved task. This was in part due to the fact that, in order to guarantee uniqueness (via iteratively dominant strategies), each player interacts with a window of μ other players. This μ depends on d in a delicate way (see [26, Fig. 1] and Lemma 2); it has to be an integer but, at the same time, needs also to balance nicely with the algebraic properties of Φ_d . Moreover we needed to provide deeper insights on the asymptotic, analytic behaviour of Φ_d , and to explore some new algebraic characteristics of Φ_d (see, e.g., [26, Lemma 7]).

In order to derive our upper bounds, we need to define a novel *approximate potential function* [18, 29]. First, in Lemma 6, we identify clear algebraic sufficient conditions for the existence of approximate equilibria with good social-cost guarantees, and then explicitly define (see the (8) and the proof of Theorem 9 in [26]) a function that satisfies them. This continuous function, which is defined in the entire space of positive reals, essentially generalizes that of Rosenthal’s in a smooth way: by setting $W = \alpha = 1$, we recover exactly the first significant terms of the well known Rosenthal potential [39] polynomial, with which one can demonstrate the usual PoS results for the unweighted case (see, e.g. [16]). The simple, analytic way in which this function is defined, is the very reason why we can handle both the approximation parameter α of the equilibrium and the ratio W of the weights in a smooth manner while at the same time providing good PoS guarantees.

It is important to stress that, by the purely analytical way in which our approximate potential function is defined, in principle it can also incorporate more general cost functions than polynomials; so, we believe that this technique may be of independent interest. We point towards that direction in [26, Appendix C].

Due to space limitations, all omitted proofs and various supplementary figures can be found in the full version of our paper [26].

2 Model and Notation

Weighted Congestion Games. A *weighted congestion game* consists of a finite, nonempty set of players N and resources (or facilities) E . Each player $i \in N$ has a real *weight* $w_i > 0$ and a *strategy set* $S_i \subseteq 2^E$. Associated with each resource $e \in E$ is a *cost* (or *latency*) *function* $c_e : (0, \infty) \rightarrow [0, \infty)$. In this paper we mainly focus on polynomial cost functions with maximum degree $d \geq 0$ and nonnegative coefficients; that is, every cost function is of the form $c_e(x) = \sum_{j=0}^d a_{e,j} \cdot x^j$, with $a_{e,j} \geq 0$ for all j . In the following, whenever we refer to polynomial cost functions we mean cost functions of this particular form.

A *pure strategy profile* is a choice of strategies $\mathbf{s} = (s_1, s_2, \dots, s_n) \in S = S_1 \times \dots \times S_n$ by the players. We use the standard game-theoretic notation $\mathbf{s}_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$, $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$, such that $\mathbf{s} = (s_i, \mathbf{s}_{-i})$. Given a pure strategy profile

\mathbf{s} , we define the *load* $x_e(\mathbf{s})$ of resource $e \in E$ as the total weight of players that use resource e on \mathbf{s} , i.e., $x_e(\mathbf{s}) = \sum_{i \in N: e \in s_i} w_i$. The *cost* player i is defined by $C_i(\mathbf{s}) = \sum_{e \in s_i} c_e(x_e(\mathbf{s}))$.

A *singleton* weighted congestion game is a special form of congestion games where the strategies of all players consist only of single resources; that is, for all players $i \in N$, $|s_i| = 1$ for all $s_i \in S_i$. In a weighted *network* congestion games the resources E are given as the edge set of some directed graph $G = (V, E)$, and each player $i \in N$ has a source $o_i \in V$ and destination $t_i \in V$ node; then, the strategy set S_i of each player is implicitly given as the edge sets of all directed $o_i \rightarrow t_i$ paths in G .

Nash Equilibria. A pure strategy profile \mathbf{s} is a pure *Nash equilibrium* if and only if for every player $i \in N$ and for all $s'_i \in S_i$, we have $C_i(\mathbf{s}) \leq C_i(s'_i, \mathbf{s}_{-i})$. Similarly a strategy profile is an α -*approximate pure Nash equilibrium*, for $\alpha \geq 1$, if $C_i(\mathbf{s}) \leq \alpha \cdot C_i(s'_i, \mathbf{s}_{-i})$ for all players $i \in N$ and $s'_i \in S_i$. As discussed in the introduction, weighted congestion games do not always admit pure Nash equilibria. However, by Nash's theorem they have mixed Nash equilibria. A tuple $\sigma = (\sigma_1, \dots, \sigma_N)$ of independent probability distributions over players' strategy sets is a *mixed Nash equilibrium* if

$$\mathbb{E}_{\mathbf{s} \sim \sigma} [C_i(\mathbf{s})] \leq \mathbb{E}_{\mathbf{s}_{-i} \sim \sigma_{-i}} [C_i(s'_i, \mathbf{s}_{-i})]$$

holds for every $i \in N$ and $s'_i \in S_i$. Here σ_{-i} is a product distribution of all σ_j 's with $j \neq i$, and \mathbf{s}_{-i} denotes a strategy profile drawn from this distribution. We use $\text{NE}(G)$ to denote the set of all mixed Nash equilibria of a game G .

Social Cost and Price of Stability. Fix a weighted congestion game G . The *social cost* of a pure strategy profile \mathbf{s} is the weighted sum of the players' costs

$$C(\mathbf{s}) = \sum_{i \in N} w_i \cdot C_i(\mathbf{s}) = \sum_{e \in E} x_e(\mathbf{s}) \cdot c_e(x_e(\mathbf{s})).$$

Denote by $\text{OPT}(G) = \min_{\mathbf{s} \in S} C(\mathbf{s})$ the *optimum social cost* over all strategy profiles $\mathbf{s} \in S$. Then, the *Price of Stability (PoS)* of G is the social cost of the best-case Nash equilibrium over the optimum social cost:

$$\text{PoS}(G) = \min_{\sigma \in \text{NE}(G)} \frac{\mathbb{E}_{\mathbf{s} \sim \sigma} [C(\mathbf{s})]}{\text{OPT}(G)}.$$

The Price of Stability of α -approximate Nash equilibria is defined accordingly. The PoS for a class \mathcal{G} of games is the worst (i.e., largest) PoS among all games in the class, that is, $\text{PoS}(\mathcal{G}) = \sup_{G \in \mathcal{G}} \text{PoS}(G)$. For example, our focus in this paper is determining the Price of Stability for the class \mathcal{G} of weighted congestion games with polynomial cost functions.

Finally, notice that, by using a straightforward scaling argument, it is without loss with respect to the PoS metric to analyse games with player weights in $[1, \infty)$; if not, divide all w_i 's with $\min_i w_i$ and scale cost functions accordingly.

3 Lower Bounds

In this section, we present our lower bound constructions. In Section 3.1 we present the general lower bound and then in Section 3.2 the lower bound for singleton games.

3.1 General Congestion Games

The next theorem presents our main negative result on the Price of Stability of weighted congestion games with polynomial latencies of degree d , that essentially matches the Price of Anarchy upper bound of Φ_d^{d+1} from Aland et al. [3]. Our result, shows a strong separation of the Price of Stability of weighted and unweighted congestion games, where the Price of Stability is at most $d + 1$ [15]. This is in sharp contrast to the Price of Anarchy of these two classes, where the respective bounds are essentially the same.

We will need to introduce some notation. Let $\Phi_d = \Theta\left(\frac{d}{\ln d}\right)$ be the unique positive root of the equation $(x + 1)^d = x^{d+1}$ and let β_d be a parameter such that $\beta_d \geq 0.38$ for any d , $\lim_{d \rightarrow \infty} \beta_d = \frac{1}{2}$. A plot of its values can be seen in [26, Fig. 1].

► **Theorem 1.** *The Price of Stability of weighted congestion games with polynomial latency functions of degree at most $d \geq 9$ is at least $(\beta_d \Phi_d)^{d+1}$.*

► **Lemma 2.** *For any positive integer d define*

$$c_d = \frac{1}{d} \left[d \frac{\ln(\Phi_d^{1+2/d} - \Phi_d) - \ln(\Phi_d^{1+2/d} - \Phi_d - 1)}{\ln \Phi_d} \right] \quad \text{and} \quad \beta_d = 1 - \Phi_d^{-c_d}, \quad (1)$$

Then $\Phi_d^{d+2} \leq \left(\Phi_d + \frac{1}{\beta_d}\right)^d$; furthermore, for all $d \geq 9$ we have: $d \cdot c_d \geq 3$, $0.38 \leq \beta_d \leq \frac{1}{2}$ and $\lim_{d \rightarrow \infty} \beta_d = \frac{1}{2}$.

Proof of Theorem 1. We now move on to the description of our congestion game instance. Fix some integer⁵ $d \geq 9$. Our instance consists of $n + \mu$ players and $n + \mu + 1$ facilities, where $\mu \equiv c \cdot d$ for some real $c \geq \frac{3}{d}$ (to be specifically determined later on, see (1)) such that $\mu \geq 3$ is an integer. You can think of n as a very large integer, since at the end we will take $n \rightarrow \infty$. Every player $i = 1, 2, \dots, n + \mu$ has a weight of $w_i = w^i$, where $w = 1 + \frac{1}{\Phi_d}$.

It will be useful for subsequent computations to notice that $w^d = \left(1 + \frac{1}{\Phi_d}\right)^d = \frac{(\Phi_d + 1)^d}{\Phi_d^d} = \frac{\Phi_d^{d+1}}{\Phi_d^d} = \Phi_d$ and $w^{d+1} = w^d \cdot w = \Phi_d \left(1 + \frac{1}{\Phi_d}\right) = \Phi_d + 1$. Let also define

$$\alpha = \alpha(\mu) \equiv \sum_{j=1}^{\mu} w^{-j} = \frac{1 - w^{-\mu}}{w - 1} = \frac{1 - (w^d)^{-c}}{w - 1} = \frac{1 - \Phi_d^{-c}}{1 + \frac{1}{\Phi_d} - 1} = \Phi_d (1 - \Phi_d^{-c}) = \beta_d \Phi_d,$$

where $\beta_d \equiv 1 - \Phi_d^{-c} \in (0, 1)$. Observe that

$$w^{-\mu} = 1 - \beta_d \Phi_d (w - 1) = 1 - \beta_d \Phi_d \left(1 + \frac{1}{\Phi_d} - 1\right) = 1 - \beta_d$$

and furthermore, for every $i \geq \mu + 1$, $\sum_{j=i-\mu}^{i-1} w_j = \sum_{j=1}^{\mu} w^{i-j} = \alpha \cdot w^i$, $\sum_{j=i-\mu}^i w_j = (\alpha + 1) \cdot w^i$ and $\sum_{\ell=1}^{\infty} w^{-\ell} = \frac{1}{w-1} = \frac{1}{1 + \frac{1}{\Phi_d} - 1} = \Phi_d$.

The facilities have latency functions

$$\begin{aligned} c_j(t) &= (\alpha + 1)^d, & \text{if } j &= 1, \dots, \mu, \\ c_j(t) &= w^{-j(d+1)} t^d, & \text{if } j &= \mu + 1, \dots, \mu + n, \\ c_{n+\mu+1}(t) &= 0, \end{aligned}$$

where for simplicity we use j instead e_j to refer to the j -th facility.

⁵ For polynomial latencies of smaller degrees $d \leq 8$ we can instead apply the simpler lower-bound instance for singleton games given in Section 3.2.

Every player i has two available strategies, s_i^* and \tilde{s}_i . Eventually we will show that the profile \mathbf{s}^* corresponds to the optimal solution, while $\tilde{\mathbf{s}}$ corresponds to the *unique* Nash equilibrium of the game. Informally, at the former the player chooses to stay at her “own” i -th facility, while at the latter she chooses to deviate and play the μ following facilities $i + 1, \dots, i + \mu$. However, special care shall be taken for the boundary cases of the first μ and last μ players, so for any player i we formally define $S_i = \{s_i^*, \tilde{s}_i\}$ where $s_i^* = \{i\}$ and

$$\tilde{s}_i = \begin{cases} \{\mu + 1, \dots, \mu + i\}, & \text{if } i = 1, \dots, \mu, \\ \{i + 1, \dots, i + \mu\}, & \text{if } i = \mu + 1, \dots, n, \\ \{i + 1, \dots, n + \mu + 1\}, & \text{if } i = n + 1, \dots, n + \mu. \end{cases}$$

These two outcomes, \mathbf{s}^* and $\tilde{\mathbf{s}}$, are shown in [26, Fig. 2].

Notice here that any facility j cannot get a load greater than the sum of the weights of the previous μ players plus the weight of the j -th player. So, for any strategy profile \mathbf{s} :

$$x_j(\mathbf{s}) \leq \sum_{\ell=j-\mu}^j w_\ell = (\alpha + 1)w^j \quad \text{for all } j \geq \mu + 1 \quad (2)$$

Next we will show that the strategy profile $\tilde{\mathbf{s}} = (\tilde{s}_1, \dots, \tilde{s}_{n+\mu})$ is the *unique* Nash equilibrium of our congestion game. We do that by proving that

1. It is a strongly *dominant* strategy for any player $i = 1, \dots, \mu$ to play \tilde{s}_i .
2. For any $i = \mu + 1, \dots, n + \mu$, given that every player $k < i$ has chosen to play \tilde{s}_k , then it is a strongly *dominant* strategy for player i to deviate to \tilde{s}_i as well.

For the first condition, fix some player $i \leq \mu$ and a strategy profile \mathbf{s}_{-i} for the other players and observe that by choosing \tilde{s}_i , player i incurs a cost of at most

$$\begin{aligned} C_i(\tilde{s}_i, \mathbf{s}_{-i}) &= \sum_{j \in \tilde{s}_i} c_j(x_j(\tilde{s}_i)) \leq \sum_{\ell=\mu+1}^{\mu+i} c_\ell((\alpha + 1)w^\ell) \\ &= \sum_{\ell=d+1}^{d+i} w^{-\ell(d+1)} (\alpha + 1)^d w^{\ell d} = (\alpha + 1)^d \sum_{\ell=d+1}^{d+i} w^{-\ell} \\ &< (\alpha + 1)^d w^{-d} \sum_{\ell=1}^{\infty} w^{-\ell} = (\alpha + 1)^d \frac{1}{\Phi_d} = (\alpha + 1)^d = C_i(s_i^*, \mathbf{s}_{-i}), \end{aligned}$$

where in the first inequality we used the bound from (2).

For the second condition, we will consider the deviations of the remaining players.⁶ Fix now some $i = \mu + 1, \dots, n$ and assume a strategy profile $\mathbf{s}_{-i} = (\tilde{s}_1, \dots, \tilde{s}_{i-1}, s_{i+1}, \dots, s_{n+\mu})$ for the remaining players. If player i chooses strategy s_i^* she will experience a cost of

$$C_i(s_i^*, \mathbf{s}_{-i}) = c_i \left(\sum_{\ell=i-\mu}^i w_\ell \right) = c_i((\alpha + 1)w^i) = w^{-i(d+1)} (\alpha + 1)^d w^{id} = (\alpha + 1)^d w^{-i}.$$

It remains to show that

$$C_i(\tilde{s}_i, \mathbf{s}_{-i}) < C_i(s_i^*, \mathbf{s}_{-i}) = (\alpha + 1)^d w^{-i}. \quad (3)$$

⁶ For the remaining last μ players $i = n + 1, \dots, n + \mu$ the proof is similar to the text, and as a matter of fact easier, since when these players deviate to \tilde{s}_i they also use the final “dummy” facility $n + \mu + 1$ that has zero cost.

The cost $C_i(\tilde{s}_i, \mathbf{s}_{-i})$ is complicated to bound immediately, for any profile \mathbf{s}_{-i} . Instead, we will resort to the following claim which characterizes the profile \mathbf{s}_{-i} where this cost is maximized.⁷

► **Claim 3.** *There exists a profile \mathbf{s}' with*

1. $s'_j = s_j$ for all $j \leq i$ and $i > i + \mu$
2. $s'_{i+\mu} = s_{i+\mu}^*$
3. *there exists some $k \in \{i + 1, \dots, i + \mu - 1\}$ such that*

$$s'_j = \tilde{s}_j \quad \text{for all } j \in \{i + 1, \dots, i + \mu - 1\} \setminus \{k\},$$

that dominates \mathbf{s} , i.e. $C_i(\tilde{s}_i, \mathbf{s}_{-i}) \leq C_i(\tilde{s}_i, \mathbf{s}'_{-i})$.

By use of Claim 3, it remains to show

$$C_i(\tilde{s}_i, \mathbf{s}'_{-i}) < (\alpha + 1)^d w^{-i}, \tag{4}$$

just for the special case of profiles \mathbf{s}' that are described in Claim 3. We do this in [26, Appendix A.4].

Summarizing, we proved that indeed $\tilde{\mathbf{s}}$ is the *unique* Nash equilibrium of our congestion game. Finally, to conclude with lower-bounding the Price of Stability, let us compute the social cost on profiles $\tilde{\mathbf{s}}$ and \mathbf{s}^* . On \mathbf{s}^* , any facility j (except the last one) gets a load equal to the weight of player j , so

$$\begin{aligned} C(\mathbf{s}^*) &= \sum_{j=1}^{n+\mu} w_j c_j(w_j) \\ &= \sum_{j=1}^{\mu} w^j (\alpha + 1)^d + \sum_{j=\mu+1}^{n+\mu} w^j w^{-j(d+1)} (w^j)^d \\ &= (\alpha + 1)^d \sum_{j=1}^{\mu} w^j + \sum_{j=\mu+1}^{\mu+n} 1 \\ &= (\alpha + 1)^d w \frac{w^{\mu} - 1}{w - 1} + n \\ &= n + (\beta \Phi_d + 1)^d \left(1 + \frac{1}{\Phi_d} \right) \frac{\frac{1}{1-\beta} - 1}{1 + \frac{1}{\Phi_d} - 1} \\ &= n + (\beta \Phi_d + 1)^d (\Phi_d + 1) \frac{\beta}{1 - \beta} \\ &\leq n + \frac{\beta}{1 - \beta} (\Phi_d + 1)^{d+1}. \end{aligned}$$

On the other hand, at the unique Nash equilibrium $\tilde{\mathbf{s}}$ each facility $j \geq \mu + 1$ receives a load equal to the sum of the weights of the previous μ players, i.e.

$$x_j(\tilde{\mathbf{s}}) = \sum_{\ell=j-\mu}^{j-1} w_{\ell} = \alpha w^j$$

⁷ For an explanatory figure and a proof of this claim, see the full version of our paper [26, Fig. 3].

so

$$C(\tilde{\mathbf{s}}) \geq \sum_{j=\mu+1}^{n+\mu} x_j(\tilde{\mathbf{s}}) c_j(x_j(\tilde{\mathbf{s}})) = \sum_{j=\mu+1}^{n+\mu} w^{-j(d+1)} (\alpha w^j)^{d+1} = \alpha^{d+1} \sum_{j=\mu+1}^{\mu+n} 1 = \alpha^{d+1} n.$$

By taking n arbitrarily large we get a lower bound on the Price of Stability of

$$\lim_{n \rightarrow \infty} \frac{C(\tilde{\mathbf{s}})}{C(\mathbf{s}^*)} \geq \lim_{n \rightarrow \infty} \frac{\alpha^{d+1} n}{n + \frac{\beta}{1-\beta} (\Phi_d + 1)^{d+1}} = \alpha^{d+1} = (\beta \Phi_d)^{d+1},$$

where from Lemma 2 we know that $\frac{1}{3} \leq \beta = \frac{1}{2} - o(1)$. ◀

3.1.1 Network Games

Due to the rather simple structure of the players' strategy sets in the lower bound construction of Theorem 1, it can be readily extended to network games as well:

► **Proposition 4.** *Theorem 1 applies also to network weighted congestion games.*

3.2 Singleton Games

In this section we give an exponential lower bound for *singleton* weighted congestion games with polynomial latency functions. The following theorem handles also approximate equilibria and provides a lower bound on the Price of Stability in a very strong sense; even if one allows for the best approximate equilibrium with approximation factor $\alpha = o\left(\frac{d}{\ln d}\right)$, then its cost is lower-bounded by $\omega(d)$ times the optimal cost.⁸ In other words, in order to achieve linear guarantees on the Price of Stability, one has to consider $\Omega\left(\frac{d}{\ln d}\right)$ -approximate equilibria—almost linear in d ; this shows that our positive result in Corollary 10, of the following Section 4.3, is essentially tight. This is furthermore complemented by [26, Appendix D], where we show that the socially optimum profile is a $(d + 1)$ -approximate equilibrium (achieving an optimal Price of Stability of 1).

► **Theorem 5.** *For any positive integer d and any real $\alpha \in [1, d)$, the α -approximate (mixed) Price of Stability of weighted (singleton) congestion games with polynomial latencies of degree at most d is at least $\frac{1}{c(d+1)} \left(1 + \frac{1}{\alpha}\right)^{d+1}$. In particular, for the special case of $\alpha = 1$, we derive that the Price of Stability of exact equilibria is $\Omega(2^d/d) = (2 - o(1))^{d+1}$.*

4 Upper Bounds

The negative results of the previous sections, involve constructions where the ratio W of the largest to smallest weight can be exponential in d . In the main theorem (Theorem 9) of this section we present an analysis which is sensitive to this parameter W , and identify conditions under which the performance of approximate equilibria can be significantly improved.

Our upper bound approach is based on the design of a suitable approximate potential function and has three main steps. First, in Section 4.1, we set up a framework for the definition of this function by identifying conditions that, on the one hand, certify the existence of an approximate equilibrium and, on the other, provide guarantees about its efficiency.

⁸ To see this, just take any upper bound of $\frac{d+1}{c \ln(d+1)}$ on α , for a constant $c > 2$. Then, the lower bound in Theorem 5 becomes $\Omega(d^{c-1})$.

Then, in Section 4.2, by use of the Euler-Maclaurin summation formula we present a general form of an approximate potential function, which extends Rosenthal’s potential for weighted congestion games (see also [26, Appendix C]). Finally, in Section 4.3, we deploy this potential for polynomial latencies. Due to its analytic description, our potential differs from other extensions of the Rosenthal’s potential that have appeared in previous work, and we believe that this contribution might be of independent interest, and applied to other classes of latency functions.

4.1 The Potential Method

In the next lemma we lay the ground for the design and analysis of approximate potential functions, by supplying conditions that not only provide guarantees for the existence of approximate equilibria, but also for their performance with respect to the social optimum. In the premises of the lemma, we give conditions on the resource functions ϕ_e , having in mind that $\Phi(\mathbf{s}) = \sum_{e \in E} \phi_e(x_e(\mathbf{s}))$ will eventually serve as the “approximate” potential function.

► **Lemma 6.** *Consider a weighted congestion game with latency functions c_e , for each facility $e \in E$, and player weights w_i , for each player $i \in N$. If there exist functions $\phi_e : [0, \infty) \rightarrow \mathbb{R}$ and parameters $\alpha_1, \alpha_2, \beta_1, \beta_2 > 0$ such that for any facility e and player weight $w \in \{w_1, \dots, w_n\}$*

$$\alpha_1 \leq \frac{\phi_e(x+w) - \phi_e(x)}{w \cdot c_e(x+w)} \leq \alpha_2, \quad \text{for all } x \geq 0, \quad (5)$$

and

$$\beta_1 \leq \frac{\phi_e(x)}{x \cdot c_e(x)} \leq \beta_2, \quad \text{for all } x \geq \min_n w_n, \quad (6)$$

then our game has an $\frac{\alpha_2}{\alpha_1}$ -approximate pure Nash equilibrium which, furthermore, has Price of Stability at most $\frac{\beta_2}{\beta_1}$.

4.2 Faulhaber’s Potential

In this section we propose an approximate potential function, which is based on the following classic number-theoretic result, known as Faulhaber’s formula⁹, which states that for any positive integers n, m ,

$$\begin{aligned} \sum_{k=1}^n k^m &= \frac{1}{m+1} \sum_{j=0}^m (-1)^j \binom{m+1}{j} B_j n^{m+1-j} \\ &= \frac{1}{m+1} n^{m+1} + \frac{1}{2} n^m + \frac{1}{m+1} \sum_{j=2}^m \binom{m+1}{j} B_j n^{m+1-j}, \end{aligned} \quad (7)$$

where the coefficients B_j are the usual Bernoulli numbers.¹⁰ In particular, this shows that the sum of the first n powers of m can be expressed as a polynomial of n with degree $m+1$. Furthermore, this sum corresponds to the well-known potential of Rosenthal [39] for *unweighted* congestion games when the latency function is the monomial $x \mapsto x^m$.

⁹ See, e.g., [33, p. 287] or [19, p. 106]).

¹⁰ See, e.g., [28, Chapter 6.5] or [1, Chapter 23]. The first Bernoulli numbers are: $B_0 = 1, B_1 = -1/2, B_2 = 1/6, B_3 = 0, B_4 = -1/30, \dots$. Also, we know that $B_j = 0$ for all *odd* integers $j \geq 3$.

Based on the above observation, we go beyond just integer values of n , and generalize this idea to all positive reals; in that way, we design a “potential” function that can handle different player weights and, furthermore, incorporate in a more powerful, analytically smooth way, approximation factors with respect to both the Price of Stability, as well as the approximation parameter of the equilibrium (in the spirit of Lemma 6). A natural way to do that is to directly generalize (7) and simply define, for any real $x \geq 0$ and positive integer m ,

$$S_m(x) \equiv \frac{1}{m+1}x^{m+1} + \frac{1}{2}x^m, \quad (8)$$

keeping just the first two significant terms.¹¹ For the special case of $m = 0$ we set $S_0(y) = y$.

For any nonnegative integer m we define the function $A_m : [1, \infty) \rightarrow (0, \infty)$ with

$$A_m(x) \equiv \left[\frac{S_m(x)}{x^{m+1}} \right]^{-1} = \left(\frac{1}{m+1} + \frac{1}{2x} \right)^{-1} = \frac{2(m+1)x}{2x + m + 1}. \quad (9)$$

Observe that A_m is strictly increasing (in x) for all $m \geq 1$,

$$A_m(1) = \frac{2(m+1)}{m+3} \in [1, 2), \quad \text{and} \quad \lim_{x \rightarrow \infty} A_m(x) = m + 1. \quad (10)$$

For the special case of $m = 0$ we simply have $A_0(x) = 1$ for all $x \geq 0$. A graph of these functions can be found in [26, Fig. 6]. Since A_m is strictly increasing for $m \geq 1$, its inverse function, $A_m^{-1} : [2\frac{m+1}{m+3}, m+1] \rightarrow [1, \infty)$, is well-defined and also strictly increasing for all $m \geq 1$.

The following two lemmas describe some useful properties regarding the algebraic behaviour, and the relation among, functions A_m and S_m :

► **Lemma 7.** *Fix any reals $y \geq x \geq 1$. Then the sequences $\frac{A_m(x)}{m+1}$ and $\frac{A_m(x)}{A_m(y)}$ are decreasing, and sequence $A_m(x)$ is increasing (with respect to m).*

► **Lemma 8.** *Fix any integer $m \geq 0$ and reals $\gamma, w \geq 1$. Then*

$$\frac{\gamma^{m+1}}{A_m(\gamma w)} \leq \frac{S_m(\gamma(x+w)) - S_m(\gamma x)}{w(x+w)^m} \leq \gamma^{m+1}, \quad \text{for all } x \geq 0, \quad (11)$$

and

$$\frac{\gamma^{m+1}}{m+1} \leq \frac{S_m(\gamma x)}{x^{m+1}} \leq \frac{\gamma^{m+1}}{A_m(\gamma)}, \quad \text{for all } x \geq 1. \quad (12)$$

4.3 The Upper Bound

Now we are ready to state our main positive result:

► **Theorem 9.** *At any congestion game with polynomial latency functions of degree at most $d \geq 1$ and player weights ranging in $[1, W]$, for any $\frac{2(d+1)W}{2W+d+1} \leq \alpha \leq d+1$ there exists an α -approximate pure Nash equilibrium that, furthermore, has Price of Stability at most*

$$1 + \left(\frac{d+1}{\alpha} - 1 \right) W.$$

¹¹See [26, Sec. 4.4] for further discussion on this choice.

Observe that, as the approximation parameter α increases, the Price of Stability decreases, in a smooth way, from $\frac{d+3}{2}$ down to the optimal value of 1. Furthermore, notice how the interval within which α ranges, shrinks as the range of player weights W grows; in particular, its left boundary $\frac{2(d+1)W}{2W+d+1}$ goes from $2\frac{d+1}{d+3}$ (for $W = 1$) up to $d + 1$ (for $W \rightarrow \infty$).

As a result, Theorem 9 has two interesting corollaries, one for $\alpha = \frac{2(d+1)W}{2W+d+1}$ and one for $W = 1$ (unweighted games):

► **Corollary 10.** *At any congestion game with polynomial latencies of degree at most $d \geq 1$ where player weights lie within the range $[1, W]$, there is an $\frac{2(d+1)W}{2W+d+1}$ -approximate pure Nash equilibrium with Price of Stability at most $\frac{d+3}{2}$.*

It is interesting to point out here that, in light of Theorem 5, the above result of Corollary 10 is essentially asymptotically tight as far as the Price of Stability is concerned (see the discussion preceding Theorem 5).

► **Corollary 11.** *At any unweighted congestion game with polynomial latencies of degree at most $d \geq 1$, the Price of Stability of α -approximate equilibria is at most $\frac{d+1}{\alpha}$, for any $2\frac{d+1}{d+3} \leq \alpha \leq d + 1$.*

References

- 1 Milton Abramowitz and Irene A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York, 9th printing, 10th gpo printing edition, 1964. URL: <http://people.maths.ox.ac.uk/~macdonald/aands/index.html>.
- 2 Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM*, 55(6):1–22, dec 2008. doi:10.1145/1455248.1455249.
- 3 Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing*, 40(5):1211–1233, jan 2011. doi:10.1137/090748986.
- 4 Susanne Albers. On the value of coordination in network design. *SIAM Journal on Computing*, 38(6):2273–2302, 2009.
- 5 Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, jan 2008. doi:10.1137/070680096.
- 6 Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. *SIAM Journal on Computing*, 42(1):160–177, jan 2013. doi:10.1137/070702370.
- 7 Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden. Weighted congestion games: The price of anarchy, universal worst-case examples, and tightness. *ACM Trans. Econ. Comput.*, 2(4):141–1423, oct 2014. doi:10.1145/2629666.
- 8 Vittorio Bilò. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. *Theory of Computing Systems*, 62, dec 2017. doi:10.1007/s00224-017-9826-1.
- 9 Vittorio Bilò, Ioannis Caragiannis, Angelo Fanelli, and Gianpiero Monaco. Improved Lower Bounds on the Price of Stability of Undirected Network Design Games. *Theory of Computing Systems*, 52(4):668–686, may 2013. doi:10.1007/s00224-012-9411-6.
- 10 Vittorio Bilò, Michele Flammini, and Luca Moscardelli. The price of stability for undirected broadcast network design with fair cost allocation is constant. *Games and Economic Behavior*, 2014.

- 11 Vittorio Bilò and Cosimo Vinci. On the impact of singleton strategies in congestion games. In *25th Annual European Symposium on Algorithms, ESA 2017, September 4-6, 2017, Vienna, Austria*, volume 87 of *LIPICs*, pages 17:1–17:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.ESA.2017.17.
- 12 Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. Efficient computation of approximate pure Nash equilibria in congestion games. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, oct 2011. doi:10.1109/focs.2011.50.
- 13 Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. Approximate pure Nash equilibria in weighted congestion games: Existence, efficient computation, and structure. *ACM Trans. Econ. Comput.*, 3(1):2:1–2:32, 2015. doi:10.1145/2614687.
- 14 Ioannis Caragiannis, Michele Flammini, Christos Kaklamanis, Panagiotis Kanellopoulos, and Luca Moscardelli. Tight bounds for selfish and greedy load balancing. *Algorithmica*, 61(3):606–637, 2010. doi:10.1007/s00453-010-9427-8.
- 15 George Christodoulou and Martin Gairing. Price of stability in polynomial congestion games. *ACM Transactions on Economics and Computation*, 4(2):1–17, dec 2015. doi:10.1145/2841229.
- 16 George Christodoulou and Elias Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Algorithms - ESA 2005, 13th Annual European Symposium*, pages 59–70, 2005.
- 17 George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *STOC '05: Proceedings of the 37th annual ACM symposium on Theory of computing*, pages 67–73, New York, NY, USA, 2005. ACM. doi:10.1145/1060590.1060600.
- 18 George Christodoulou, Elias Koutsoupias, and Paul G. Spirakis. On the performance of approximate equilibria in congestion games. *Algorithmica*, 61(1):116–140, 2011. doi:10.1007/s00453-010-9449-2.
- 19 John H. Conway and Richard K. Guy. *The Book of Numbers*. Springer New York, 1996. doi:10.1007/978-1-4612-4072-3.
- 20 Juliane Dunkel and Andreas S. Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Mathematics of Operations Research*, 33(4):851–868, 2008.
- 21 Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*, STOC '04, pages 604–612, New York, NY, USA, 2004. doi:10.1145/1007352.1007445.
- 22 Matthias Feldotto, Martin Gairing, Grammateia Kotsialou, and Alexander Skopalik. Computing approximate pure Nash equilibria in Shapley value weighted congestion games. In *Web and Internet Economics*, pages 191–204. Springer International Publishing, 2017. doi:10.1007/978-3-319-71924-5_14.
- 23 Amos Fiat, Haim Kaplan, Meital Levy, Svetlana Olonetsky, and Ronen Shabo. On the price of stability for designing undirected networks with fair cost allocations. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming: 33rd International Colloquium (ICALP)*, pages 608–618, 2006. doi:10.1007/11786986_53.
- 24 Dimitris Fotakis, Spyros Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. *Theoretical Computer Science*, 410(36):3305–3326, 2009. doi:10.1016/j.tcs.2008.01.004.
- 25 Dimitris Fotakis, Spyros Kontogiannis, and Paul Spirakis. Selfish unsplitable flows. *Theoretical Computer Science*, 348(2):226–239, 2005.

- 26 Yiannis Giannakopoulos, George Christodoulou, Martin Gairing, and Paul Spirakis. The price of stability of weighted congestion games. *CoRR*, abs/1802.09952, 2018. arXiv: arXiv:1802.09952.
- 27 M. Goemans, Vahab Mirrokni, and A. Vetta. Sink equilibria and convergence. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 142–151, Oct 2005. doi:10.1109/SFCS.2005.68.
- 28 Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1989.
- 29 Christoph Hansknecht, Max Klimm, and Alexander Skopalik. Approximate pure Nash equilibria in weighted congestion games. In Klaus Jansen, José D. P. Rolim, Nikhil R. Devanur, and Cristopher Moore, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014)*, volume 28 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 242–257. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2014. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.242.
- 30 Tobias Harks and Max Klimm. On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research*, 37(3):419–436, 2012. doi:10.1287/moor.1120.0543.
- 31 Tobias Harks, Max Klimm, and Rolf H Möhring. Characterizing the existence of potential functions in weighted congestion games. *Theory Comput Syst*, 49:46–70, 2011. doi:10.1007/s00224-011-9315-x.
- 32 Tobias Harks, Max Klimm, and Rolf H Möhring. Strong equilibria in games with the lexicographical improvement property. *International Journal of Game Theory*, 42(2):461–482, 2012. doi:10.1007/s00182-012-0322-1.
- 33 Donald E. Knuth. Johann Faulhaber and sums of powers. *Mathematics of Computation*, 61(203):277–277, sep 1993. doi:10.1090/s0025-5718-1993-1197512-7.
- 34 Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *16th Annual Symposium on Theoretical Aspects of Computer Science*, STACS '99, pages 404–413, 1999.
- 35 Euiwoong Lee and Katrina Ligett. Improved bounds on the price of stability in network cost sharing games. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 607–620. ACM, 2013. doi:10.1145/2482540.2482562.
- 36 Lavy Libman and Ariel Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, Aug 2001. doi:10.1023/A:1016770831869.
- 37 Dov Monderer and Lloyd S. Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.
- 38 Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- 39 Robert W Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- 40 Robert W Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.
- 41 Tim Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32:1–32:42, 2015. doi:10.1145/2806883.
- 42 Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, mar 2002. doi:10.1145/506147.506153.
- 43 Tim Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, may 2004. doi:10.1016/j.geb.2003.06.004.
- 44 Andreas S. Schulz and Nicolás Stier Moses. On the performance of user equilibria in traffic networks. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete*

150:16 The Price of Stability of Weighted Congestion Games

Algorithms, SODA '03, pages 86–87, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.