

The Price of Anarchy for Polynomial Social Cost^{*}

Martin Gairing¹, Thomas Lücking¹, Marios Mavronicolas², and
Burkhard Monien¹

¹ Faculty of Computer Science, Electrical Engineering and Mathematics,
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany.

Email: {gairing,luck,bm}@uni-paderborn.de

² Department of Computer Science,
University of Cyprus, P. O. Box 20537, Nicosia CY-1678, Cyprus.

Email: mavronic@ucy.ac.cy

Abstract. In this work, we consider an interesting variant of the well-studied *KP model* [18] for *selfish routing* that reflects some influence from the much older *Wardrop model* [31]. In the new model, user *traffics* are still *unsplittable*, while *social cost* is now the expectation of the sum, over all links, of a certain polynomial evaluated at the total *latency* incurred by all users choosing the link; we call it *polynomial social cost*. The polynomials that we consider have non-negative coefficients. We are interested in evaluating *Nash equilibria* in this model, and we use the *Price of Anarchy* as our evaluation measure. We prove the *Fully Mixed Nash Equilibrium Conjecture* for identical users and two links, and establish an approximate version of the conjecture for arbitrary many links. Moreover, we give upper bounds on the Price of Anarchy.

1 Introduction

Motivation and Framework. The *Price of Anarchy*, also known as *coordination ratio*, has been defined in a 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* [24] – a state in which no *user* wishes to leave unilaterally its own strategy in order to improve the value of its private objective function, also known as *individual cost*. So, the Price of Anarchy represents a rendezvous of Nash equilibrium, a concept fundamental to Game Theory, with *approximation*, an ubiquitous concept in Theoretical Computer Science today (see, e.g., [30]).

* This work has been partially supported by the DFG-Sonderforschungsbereich 376 Massive Parallelität: Algorithmen, Entwurfsmethoden, Anwendungen, by the EU within the 6th Framework Programme under contract 001907 “Dynamically Evolving, Large Scale Information Systems” (DELIS), by research funds at University of Cyprus, and by the VEGA grant No. 2/3164/23.

The setting considered by Koutsoupias and Papadimitriou [18], now widely known as the *KP model*, involves a network of m parallel *links* and a collection of n selfish *users* with *traffics* w_1, \dots, w_n which are *unsplittable*. The *latency* incurred at a link is the total traffic of users choosing it. *Expected latency* makes sense in case users' *strategies* are *mixed*, i.e., when they are probability distributions over the links. The individual cost of a user over a link is the expected latency it would experience had it chosen the link. In a Nash equilibrium, each user is minimizing its individual cost. The social cost is the expectation of maximum (over links) latency; the social optimum is the least possible maximum latency. Recently, there has been a lot of flourishing interest and attention into the KP model; see, e.g., [3, 9, 11, 15, 17, 22]. The resulting research is overviewed in three authoritative surveys [7, 12, 16].

A much older theoretical model for selfish routing in transportation networks due to Wardrop [31] dates back to the 1950s. In this model, users' traffics are infinitesimally *splittable*, a feature that rules out mixed strategies from consideration. In addition, social cost is defined here as the sum of the individual costs (those being taken as the sum of latencies, determined by convex functions of the total traffics, along paths). Wardrop's model, and, in particular, the incurred Price of Anarchy, has been reinvestigated in many recent works [19, 26–29].

On a middle ground, some recent research [14, 20] has attempted to understand the dependence of the Price of Anarchy on the particular assumptions made in the adopted model for selfish routing. In this direction, [14, 20] has each considered a hybridization of the KP model and the Wardrop model. In this paper, we further address this line of research by introducing and studying a new, yet interesting variant of the KP model that reflects some influence from the Wardrop model.

In the proposed model, traffics are unsplittable, which makes it interesting to study mixed strategies, and the network consists of parallel links as in the KP model; however, we define social cost as the expectation of a certain *sum* of link latencies, as opposed to expectation of *maximum* latency adopted in the KP model. The idea of summing (over links) comes from the Wardrop model, although social cost is defined there as the sum of individual costs (over users). An additional influence from the Wardrop model lies in modeling each latency cost function as a certain polynomial of total traffic of users choosing the link. So, social cost is now the expectation of a sum of polynomial functions, and we call it *polynomial social cost*. The polynomials that we consider have non-negative coefficients.

Restricted to *pure* Nash equilibria, our model was already studied for monotone functions in [8], and, restricted to polynomials of degree 2 in [20]. Restricted to identical users, it is a particular instance of what is known as *congestion game* [23, 25]. See below for more on related work.

Contribution and Significance. We study a natural conjecture asserting that the fully mixed Nash equilibrium \mathbf{F} , that is, the Nash equilibrium where each user chooses each link with non-zero probability, is the *worst* Nash equilibrium with respect to social cost:

Conjecture 1 (Fully Mixed Nash Equilibrium Conjecture). For any traffic vector \mathbf{w} and for any Nash equilibrium \mathbf{P} , $\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P}) \leq \text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{F})$.

In the following, we refer to this conjecture as the FMNE Conjecture. Moreover, we investigate the Price of Anarchy. Our results address the following two important special cases of the problem:

1. The case of identical users: To analyze this case, we first show that the polynomial social cost of the *fully mixed Nash equilibrium* is a certain combinatorial sum of *Stirling numbers of the second kind*. Moreover, the polynomial social cost of any (mixed) Nash equilibrium can be bounded from above by a sum of binomial cost functions, which again turns out to be certain combinatorial sum of *Stirling numbers of the second kind*. By considering terms of the same power pairwise, this enables us to prove that the polynomial social cost of any Nash equilibrium is bounded by the polynomial social cost of the fully mixed Nash equilibrium for the case of 2 links. This implies the validity of the *Fully Mixed Nash Equilibrium Conjecture* for this special case, in our model. Equipped now with the validity of the Fully Mixed Nash Equilibrium Conjecture, we prove an *exact* bound of $\frac{1}{2}(2^{d-1} + 1)$ on the Price of Anarchy for the special case where the polynomial is just the d th power, for any integer $d \geq 1$. We also obtain an upper bound on the Price of Anarchy for the general case of polynomials with non-negative coefficients.

We then generalize these results to the case of many links. Using similar techniques as in the case of 2 links, we prove that the polynomial social cost of any arbitrary (mixed) Nash equilibrium is no more than $(1 + \frac{1}{n-1})^d$ times the polynomial social cost of the fully mixed Nash equilibrium. This implies the validity of an approximate version of the Fully Mixed Nash Equilibrium Conjecture, in our model. Again, making use of this result, we prove that the Price of Anarchy is bounded from above by $(1 + \frac{1}{n-1})^d \cdot B_d$, where B_d is the d th Bell number. We also obtain an upper bound on the Price of Anarchy for the general case of polynomials with non-negative coefficients.

2. The case of pure Nash equilibria: We finally consider the special case of *pure Nash equilibria*. Through elegant combinatorial arguments, we prove a *tight* bound of $\frac{(2^d-1)^d}{(d-1)(2^d-2)^{d-1}} (\frac{d-1}{d})^d$ on the Price of Anarchy for the same special case where the polynomial is just the d th power, for any integer $d \geq 2$.

Related Work and Comparison. The KP model has been studied extensively in the recent years; see, e.g., [9, 11, 15, 17, 20, 22]. Mavronicolas and Spirakis [22] introduced the fully mixed Nash equilibrium. They showed that, in case it exists, the fully mixed Nash equilibrium is unique. Gairing et al. [15] were the first who stated explicitly the Fully Mixed Nash Equilibrium Conjecture. Up to now, the conjecture could be proved only for several particular cases of the KP model [15, 21]. A proof of the conjecture will enable the derivation of upper bounds on Price of Anarchy via studying the fully mixed Nash equilibrium.

Lücking et al. [20] considered a hybridization of the KP model with the Wardrop model that adopts *quadratic social cost*, defined as the sum of weighted individual costs. Their model is a special case of our model where the polynomial

is just the second power. In this model, Lücking et al. [20] proved the FMNE Conjecture in case of identical users and identical links.

Gairing et al. [14] introduced yet another model for selfish routing that adopts arbitrary convex latency cost functions and the social cost as the sum of individual costs. This model is identical to our model exactly in the case of identical users and linear cost functions. However, in general the models differ in the definition of individual cost. Hence, the models also differ in the Nash conditions, and thus the sets of mixed Nash equilibria are different. Gairing et al. [14] proved again the FMNE Conjecture in their model.

Tight bounds on the Price of Anarchy for the KP model have been shown to be $\Theta(\frac{\log m}{\log \log m})$ if all links are identical [9, 17] and $\Theta(\frac{\log m}{\log \log \log m})$ if links may have varying capacities [9]. The Price of Anarchy has also been studied for a network creation game in [10] and a network design game in [2].

The Wardrop model was already studied in the 1950's [5, 31] in the context of road traffic systems. Wardrop [31] introduced the concept of equilibrium to describe user behavior in this kind of traffic networks. For a survey of the early work on this model see [4]. A lot of subsequent work on this model has been motivated by Braess's Paradox [6]. Inspired by the new interest in the Price of Anarchy, Roughgarden and Tardos [27–29] re-investigated the Wardrop model. For a survey of results we refer to [12] and references therein.

Road Map. The rest of this paper is organized as follows. Section 2 summarizes some mathematical preliminaries that will be used throughout the paper. Section 3 introduces our formal model and recalls some preliminary facts. In Section 4, we consider the case of identical users and two links. In Section 5, we turn our attention to the more general case of identical users. The case of pure Nash equilibria is treated in Section 6.

2 Mathematical Preliminaries

For all $i \in \mathbb{N}$ we denote $[i] = \{1, \dots, i\}$.

Falling Factorials, Stirling Numbers and Bell Numbers. For any pair of integers $x \geq 1$ and $i \geq 0$, the i th *falling factorial* $x^{\underline{i}}$ is given by $x^{\underline{i}} = x \cdot (x - 1) \cdots (x - i)$. For any pair of integers $d \geq 1$ and $0 \leq i \leq d$, the *Stirling number of the second kind* $S(d, i)$ counts the number of partitions of a set with d elements into exactly i *blocks* (non-empty subsets). In particular, $S(d, 0)$ is taken to be 0, while $S(d, 1) = S(d, d) = 1$. Also, for all $d \geq 2$, $S(d, 2) = 2^{d-1} - 1$. It is known that for all integers $d \geq 1$, it holds that

$$x^d = \sum_{1 \leq i \leq d} S(d, i) \cdot x^{i-1}. \quad (1)$$

This implies that the Stirling numbers of the second kind are the *connecting coefficients* between the sequence of powers and the sequence of falling factorials. For any integer $d \geq 1$, the *Bell number* B_d counts the number of partitions of a set with d elements into blocks. Thus, $B_d = \sum_{1 \leq i \leq d} S(d, i)$. For a textbook introduction to falling factorials, Stirling numbers of the second kind and Bell numbers, see [1, Chapters II & III].

Binomial Cost Functions. We now give the definition and a technical lemma on the *binomial cost function* $H(p, g)$. We will see later that polynomial social cost of a mixed Nash equilibrium can be expressed as sum of binomial cost functions.

Definition 1. For a set of r probabilities $p = (p_1, \dots, p_r)$ and a function $g : \mathbb{N} \rightarrow \mathbb{R}$ define

$$H(p, g) = \sum_{A \subseteq [r]} \prod_{k \in A} p_k \prod_{k \notin A} (1 - p_k) \cdot g(|A|).$$

In the same way, we define a function $H(\tilde{p}, r, g)$ by replacing p with a vector of r probabilities all equal to \tilde{p} . In case that $g(\delta) = \delta^d$, we write $H_d(p)$ and $H_d(\tilde{p}, r)$, respectively.

Note, that in case of equal probabilities

$$H(\tilde{p}, r, g) = \sum_{1 \leq k \leq r} \binom{r}{k} \tilde{p}^k (1 - \tilde{p})^{r-k} \cdot g(k). \quad (2)$$

Throughout the paper, we will use several times the following technical fact which has been proved in [14].

Lemma 1 (Gairing et al. [14]). Let g be convex and define $p = (p_1, \dots, p_r)$ and $\tilde{p} = \frac{\sum_{i \in [r]} p_i}{r}$. Then $H(p, g) \leq H(\tilde{p}, r, g)$.

The following proposition, proved in the appendix, will be extensively used in the remainder of the paper.

Proposition 1. For the binomial cost function $H_d(p, r)$ we have

$$H_d(p, r) = \sum_{1 \leq i \leq d} p^i \cdot S(d, i) \cdot r^{\underline{i-1}}.$$

3 Model and Preliminaries

General. We consider a *network* consisting of a set of m parallel *links* $1, 2, \dots, m$ from a *source* node to a *destination* node. Each of n *network users* $1, 2, \dots, n$, or *users* for short, wishes to route a particular amount of traffic along a (non-fixed) link from source to destination. Denote w_i the *traffic* of user $i \in [n]$. In the model of *identical users*, all user traffics are equal to 1; user traffics may vary arbitrarily in the model of *arbitrary users*. Define the $n \times 1$ *traffic vector* \mathbf{w} in the natural way. Assume throughout that $m \geq 2$ and $n \geq 2$.

A *pure strategy* for user $i \in [n]$ is some specific link. A *mixed strategy* for user $i \in [n]$ is a probability distribution over pure strategies; thus, a mixed strategy is a probability distribution over the set of links. The *support* of the mixed strategy for user $i \in [n]$, denoted $\text{support}(i)$, is the set of those pure strategies (links) to which i assigns positive probability. The *view* of a link $j \in [m]$, denoted $\text{view}(j)$, is the set of users which assign their traffics to j with non-zero probability. A *pure strategy profile* is represented by an n -tuple $\langle \ell_1, \ell_2, \dots, \ell_n \rangle \in [m]^n$; a *mixed*

strategy profile is represented by an $n \times m$ probability matrix \mathbf{P} of nm probabilities p_{ij} , $i \in [n]$ and $j \in [m]$, where p_{ij} is the probability that user i chooses link j .

For a probability matrix \mathbf{P} , the support of the mixed strategy for user $i \in [n]$ is the set $\{j \in [m] \mid p_{ij} > 0\}$. A mixed strategy profile \mathbf{P} is *fully mixed* [22, Section 2.2] if for all users $i \in [n]$ and links $j \in [m]$, $p_{ij} > 0$. Throughout, we will cast a pure strategy profile as a special case of a mixed strategy profile.

System and Cost Measures. For a pure strategy profile $\langle \ell_1, \ell_2, \dots, \ell_n \rangle$, the *latency cost for user $i \in [n]$* is $\lambda_i = \sum_{k: \ell_k = \ell_i} w_k$; that is, the latency of the link it chooses. For a mixed strategy profile \mathbf{P} , denote δ_j the *actual traffic* on link $j \in [m]$; so, δ_j is a random variable. For each link $j \in [m]$, define the *expected latency* Λ_j as the *expected traffic* on link j ; thus, $\Lambda_j = \mathcal{E}(\delta_j) = \sum_{i \in [n]} p_{ij} w_i$. For a mixed strategy profile \mathbf{P} , the *expected latency cost* for user $i \in [n]$ on link $j \in [m]$, denoted λ_{ij} , is the expectation, over all random choices of the remaining users, of the latency cost for user i had its traffic been assigned to link j ; thus, $\lambda_{ij} = w_i + \sum_{k=1, k \neq i} p_{kj} w_k = (1 - p_{ij})w_i + \Lambda_j$. For each user $i \in [n]$, the *minimum expected latency cost*, denoted λ_i , is the minimum, over all links $j \in [m]$, of the expected latency cost for user i on link j ; thus, $\lambda_i = \min_{j \in [m]} \lambda_{ij}$.

Let $\pi^d(\delta) = \sum_{0 \leq t \leq d} a_t \delta^t$ be a polynomial of degree $d \geq 1$ with non-negative coefficients (that is, $a_t \geq 0$ for all $0 \leq t \leq d$). Associated with a traffic vector \mathbf{w} , a *polynomial cost function* $\pi^d(\delta)$, and a mixed strategy profile \mathbf{P} is the *polynomial social cost*, or *social cost* for short, denoted $\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P})$, which is the expectation of the sum of the polynomial $\pi^d(\delta)$ evaluated at the incurred link latencies:

$$\begin{aligned} \text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P}) &= \mathcal{E} \left(\sum_{j \in [m]} \pi^d \left(\sum_{k: \ell_k = j} w_k \right) \right) \\ &= \sum_{j \in [m]} \sum_{A \subseteq [n]} \left(\prod_{i \in A} p_{ij} \right) \left(\prod_{i \notin A} (1 - p_{ij}) \right) \pi^d \left(\sum_{k: \ell_k = j} w_k \right). \end{aligned}$$

If we restrict to the polynomial $\pi^d(\delta) = \delta^d$, we write $\text{SC}_{\delta^d}(\mathbf{w}, \mathbf{P})$. Note that

$$\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P}) = \sum_{0 \leq t \leq d} a_t \cdot \text{SC}_{\delta^t}(\mathbf{w}, \mathbf{P}). \quad (3)$$

Moreover, if we restrict to identical users, the formula for social cost reduces to

$$\begin{aligned} \text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P}) &= \sum_{j \in [m]} \sum_{A \subseteq [n]} \left(\prod_{i \in A} p_{ij} \right) \left(\prod_{i \notin A} (1 - p_{ij}) \right) \pi^d(|A|) \\ &= \sum_{j \in [m]} H((p_{1j}, \dots, p_{nj}), \pi^d(\delta_j)). \end{aligned} \quad (4)$$

The *optimum* associated with a traffic vector \mathbf{w} , denoted $\text{OPT}_{\pi^d(\delta)}(\mathbf{w})$, is the *least possible* sum of the polynomial $\pi^d(\delta)$ evaluated at the incurred link latencies. Note that while $\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P})$ is defined in relation to a mixed strategy profile \mathbf{P} , $\text{OPT}_{\pi^d(\delta)}(\mathbf{w})$ refers to the *optimum* pure strategy profile.

Nash Equilibria. We are interested in a special class of mixed strategies called Nash equilibria [24] that we describe below. Formally, the probability matrix \mathbf{P} is a *Nash equilibrium* [18, Section 2] if for all users $i \in [n]$ and links $j \in [m]$, $\lambda_{ij} = \lambda_i$ if $I_{ij} = 1$, and $\lambda_{ij} \geq \lambda_i$ if $I_{ij} = 0$. Thus, each user assigns its traffic with positive probability only on links (possibly more than one of them) for which its expected latency cost is minimized; this implies that there is no incentive for a user to unilaterally deviate from its mixed strategy in order to avoid links on which its expected latency cost is higher than necessary. We will refer to this conditions as *Nash conditions*. Depending on the type of strategy profile, we differ between *pure*, *mixed* and *fully mixed* Nash equilibria. Recall that in our model, the fully mixed Nash equilibrium \mathbf{F} exists always, and that it is unique (see [22]).

Lemma 2 (Mavronicolas and Spirakis [22]). *Consider the model of arbitrary users. Then, there exists a unique fully mixed Nash equilibrium \mathbf{F} with $f_{ij} = 1/m$, for any user $i \in [n]$ and link $j \in [m]$.*

Price of Anarchy. The *Price of Anarchy*, also called *coordination ratio*, is the maximum value, over all traffic vectors \mathbf{w} and Nash equilibria \mathbf{P} of the ratio $\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P}) / \text{OPT}_{\pi^d(\delta)}(\mathbf{w})$.

4 Identical Users and Two Links

In this section we turn our attention to the case of identical users and two links. We first use Proposition 1 and Lemma 3 to prove the FMNE Conjecture (Theorem 1). Making use of this result, we then give a tight bound on the Price of Anarchy for the special case where the polynomial cost function is just the d th power (Theorem 2). We close with an upper bound on the Price of Anarchy for general polynomial cost functions (Corollary 1).

Lemma 3 (Lücking et al. [20]). *Consider the model of arbitrary users, and let $j, k \in [m]$. If $\text{view}(j) \subsetneq \text{view}(k)$, then $\Lambda_j > \Lambda_k$.*

Theorem 1. *Consider the model of identical users and two links. Then, for any Nash equilibrium \mathbf{P} , $\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P}) \leq \text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{F})$.*

Proof. Fix any Nash equilibrium \mathbf{P} . We can identify three sets of users in \mathbf{P} : $\mathcal{U}_1 = \{i : \text{support}(i) = \{1\}\}$, $\mathcal{U}_2 = \{i : \text{support}(i) = \{2\}\}$ and $\mathcal{U}_{12} = \{i : \text{support}(i) = \{1, 2\}\}$. Without loss of generality let $|\mathcal{U}_1| \leq |\mathcal{U}_2|$. Denote $u = |\mathcal{U}_1|$, $v = |\mathcal{U}_2| - u$ and $r = |\mathcal{U}_{12}|$.

The proof of the theorem is structured as follows: we first prove that the claim holds if \mathbf{P} is a pure Nash equilibrium. Then we consider the case where $u = 0$. We show that it suffices to consider terms of the same power pairwise. To do so, we proceed by induction on i , $1 \leq i \leq d$, and we show that the claim holds for i th powers. For the case where $u > 0$, we introduce a Nash equilibrium \mathbf{Q} as a perturbation of \mathbf{F} with a special structure that is similar to the structure of \mathbf{P} such that $\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{Q}) \leq \text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{F})$. Due to this special structure, \mathbf{Q}

can be compared with \mathbf{P} more easily. To compare the two, we use the fact that the claim holds for $u = 0$. We now continue with the details of the formal proof:

First, let \mathbf{P} be a pure Nash equilibrium, that is, $r = 0$. Due to the Nash conditions, $|\mathcal{U}_1|$ and $|\mathcal{U}_2|$ differ by 1 if n is odd, and $|\mathcal{U}_1| = |\mathcal{U}_2|$ otherwise. Thus, each pure Nash equilibrium \mathbf{P} has unique social cost $\text{SC}_{\pi^a(\delta)}(\mathbf{w}, \mathbf{P}) = \text{OPT}_{\pi^a(\delta)}(\mathbf{w})$, proving the claim for pure Nash equilibria.

Now let \mathbf{P} be an arbitrary (non-pure) Nash equilibrium.

Case 1: $u = 0$: Denote by A_1 and A_2 the latency on link 1 and 2, respectively. Due to Lemma 3, $\text{view}(1) \subsetneq \text{view}(2)$ implies $A_1 > A_2$. First assume $r = 1$. Then, the Nash condition for the mixed user $1 = 1 + v$ implies $v = 0$. This contradicts our assumption that $n \geq 2$. So, assume that $2 \leq r \leq n - 1$ (mixed) users are assigned to both links, and that $n - r$ (pure) users are assigned to link 2. Consider any arbitrary mixed user $i \in \mathcal{U}_{12}$. Then, the Nash condition implies

$$A_1 - p_{i1} + 1 = A_2 - p_{i2} + 1 = A_2 + p_{i1} \quad \Rightarrow \quad p_{i1} = \frac{A_1 - A_2 + 1}{2}.$$

Since this holds for each mixed user, we write $p(1)$ and $p(2)$ for the probabilities of each mixed user on link 1 and 2, respectively. Since $p(2) = 1 - p(1)$, the Nash condition

$$(r - 1) \cdot p(1) + 1 = (r - 1) \cdot p(2) + n - r + 1 = (r - 1) \cdot (1 - p(1)) + n - r + 1$$

for the mixed users implies

$$p(1) = \frac{n - 1}{2(r - 1)} = \frac{1}{2} + \frac{n - r}{2(r - 1)} \quad \text{and} \quad p(2) = (1 - p(1)) = \frac{1}{2} - \frac{n - r}{2(r - 1)}.$$

Thus, we can write

$$A_1 = r \cdot p(1) = r \cdot \left(\frac{1}{2} + \frac{n - r}{2(r - 1)} \right) = \frac{n}{2} + \frac{n - r}{2(r - 1)} = \alpha + \beta,$$

where $\alpha = \frac{n}{2}$ and $\beta = \frac{n - r}{2(r - 1)}$ with $0 < \beta < \frac{1}{2}$. Moreover, $A_2 = n - A_1 = \alpha - \beta$. By Equation (3), we can write $\text{SC}_{\pi^a(\delta)}(\mathbf{w}, \mathbf{P}) = \sum_{0 \leq t \leq d} a_t \cdot \text{SC}_{\delta^t}(\mathbf{w}, \mathbf{P})$ for any Nash equilibrium \mathbf{P} . Since the coefficients a_t are non-negative, it suffices to prove that $\text{SC}_{\delta^t}(\mathbf{w}, \mathbf{P}) \leq \text{SC}_{\delta^t}(\mathbf{w}, \mathbf{F})$ for all $1 \leq t \leq d$.

Consider the average probabilities $\tilde{p}(1) = \frac{A_1}{r} = \frac{\alpha + \beta}{r}$ and $\tilde{p}(2) = \frac{A_2}{n} = \frac{\alpha - \beta}{n}$ on link 1 and 2, respectively. Since δ^t is convex, we can apply Lemma 1 to derive the upper bound

$$\begin{aligned} \text{SC}_{\delta^t}(\mathbf{w}, \mathbf{P}) &\stackrel{(4)}{=} H_t(\underbrace{(p(1), \dots, p(1))}_r) + H_t(\underbrace{(p(2), \dots, p(2))}_r, \underbrace{1, \dots, 1}_{n-r}) \\ &\stackrel{\text{Lem. 1}}{\leq} H_t(\tilde{p}(1), r) + H_t(\tilde{p}(2), n) = H_t\left(\frac{\alpha + \beta}{r}, r\right) + H_t\left(\frac{\alpha - \beta}{n}, n\right) \\ &\stackrel{\text{Prop. 1}}{=} \sum_{1 \leq i \leq t} \left(\frac{\alpha + \beta}{r}\right)^i \cdot S(t, i) \cdot r^{i-1} + \sum_{1 \leq i \leq t} \left(\frac{\alpha - \beta}{n}\right)^i \cdot S(t, i) \cdot n^{i-1} \\ &= \sum_{1 \leq i \leq t} S(t, i) \cdot \left[(\alpha + \beta)^i \cdot \frac{r^{i-1}}{r^i} + (\alpha - \beta)^i \cdot \frac{n^{i-1}}{n^i} \right] \end{aligned}$$

Moreover, by Lemma 2, we have

$$\text{SC}_{\delta^t}(\mathbf{w}, \mathbf{F}) \stackrel{(4)}{=} 2 \cdot H_t\left(\frac{1}{2}, n\right) = 2 \cdot H_t\left(\frac{\alpha}{n}, n\right) \stackrel{Prop. 1}{=} \sum_{1 \leq i \leq t} S(t, i) \left[2\alpha^i \cdot \frac{n^{i-1}}{n^i} \right].$$

In order to prove the claim it suffices to show that

$$\Delta = 2\alpha^i \cdot \frac{n^{i-1}}{n^i} - \left[(\alpha + \beta)^i \cdot \frac{r^{i-1}}{r^i} + (\alpha - \beta)^i \cdot \frac{n^{i-1}}{n^i} \right] \geq 0$$

for all $i \geq 1$. For $i = 1$, we have $\Delta = 2\alpha - [\alpha + \beta + \alpha - \beta] = 0$, proving the basis case. So assume $i \geq 2$. We have

$$\begin{aligned} \Delta = 2\alpha^i \cdot \frac{n^{i-1}}{n^i} - & \left[(\alpha + \beta) \cdot \frac{r - (i-1)}{r} \cdot \underbrace{(\alpha + \beta)^{i-1} \cdot \frac{r^{i-2}}{r^{i-1}}}_{\geq 0} \right. \\ & \left. + (\alpha - \beta) \cdot \frac{n - (i-1)}{n} \cdot \underbrace{(\alpha - \beta)^{i-1} \cdot \frac{n^{i-2}}{n^{i-1}}}_{\geq 0} \right]. \end{aligned}$$

The function

$$f(r) = (\alpha + \beta) \cdot \frac{r - (i-1)}{r} = \left(\frac{(n-1)r}{2(r-1)} \right) \cdot \frac{r - (i-1)}{r} = \frac{n-1}{2} \cdot \frac{r - (i-1)}{r-1}$$

is monotonic increasing in r for all $i \geq 2$, and thus

$$\begin{aligned} \Delta & \geq 2\alpha^i \cdot \frac{n^{i-1}}{n^i} - \left[\alpha \cdot \frac{n - (i-1)}{n} \cdot (\alpha + \beta)^{i-1} \cdot \frac{r^{i-2}}{r^{i-1}} \right. \\ & \quad \left. + \alpha \cdot \frac{n - (i-1)}{n} \cdot (\alpha - \beta)^{i-1} \cdot \frac{n^{i-2}}{n^{i-1}} \right] \\ & = 2\alpha^i \cdot \frac{n^{i-1}}{n^i} - \alpha \cdot \frac{n - (i-1)}{n} \cdot \left[(\alpha + \beta)^{i-1} \cdot \frac{r^{i-2}}{r^{i-1}} + (\alpha - \beta)^{i-1} \cdot \frac{n^{i-2}}{n^{i-1}} \right] \\ & \stackrel{ind.}{\geq} 2\alpha^i \cdot \frac{n^{i-1}}{n^i} - \alpha \cdot \frac{n - (i-1)}{n} \cdot 2\alpha^{i-1} \cdot \frac{n^{i-2}}{n^{i-1}} = 0, \end{aligned}$$

proving the inductive claim. This finishes the case where $u = 0$.

Case 2: $u > 0$: Consider the following mixed Nash equilibrium \mathbf{Q} . On both links, there are u pure users, and the remaining $n - 2u$ users have probabilities $q(1) = q(2) = \frac{1}{2}$. Since $\pi^d(\delta)$ is convex and due to Lemma 1, we have $\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{F}) \geq \text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{Q})$. As in case 1, it suffices to show that $\text{SC}_{\delta^t}(\mathbf{w}, \mathbf{Q}) \geq \text{SC}_{\delta^t}(\mathbf{w}, \mathbf{P})$ for all $1 \leq t \leq d$. Denote $\tilde{n} = n - 2u$. If $\tilde{n} = 1$, then $\text{SC}_{\delta^t}(\mathbf{w}, \mathbf{P}) = \text{SC}_{\delta^t}(\mathbf{w}, \mathbf{Q}) = (u+1)^t + u^t$. Thus, assume $\tilde{n} \geq 2$. We have

$$\text{SC}_{\delta^t}(\mathbf{w}, \mathbf{Q}) \stackrel{(4),(2)}{=} 2 \cdot \sum_{0 \leq k \leq \tilde{n}} \binom{\tilde{n}}{k} q(1)^k (1-q(1))^{\tilde{n}-k} (k+u)^t = \text{SC}_{(\delta+u)^t}(\tilde{\mathbf{w}}, \tilde{\mathbf{Q}}),$$

and by Lemma 1, Equation (4) and Equation (2)

$$\begin{aligned} \text{SC}_{\delta^t}(\mathbf{w}, \mathbf{P}) &\leq \sum_{0 \leq k \leq r} \binom{r}{k} p(1)^k (1-p(1))^{r-k} (k+u)^t \\ &\quad + \sum_{0 \leq k \leq \tilde{n}} \binom{\tilde{n}}{k} (1-p(1))^k p(1)^{\tilde{n}-k} (k+u)^t \\ &= H(p(1), r, (\delta+u)^t) + H(p(2), \tilde{n}, (\delta+u)^t) = \text{SC}_{(\delta+u)^t}(\tilde{\mathbf{w}}, \tilde{\mathbf{P}}) \end{aligned}$$

where $\tilde{\mathbf{w}}$ is an instance with $\tilde{n} \geq 2$ identical users, $\tilde{\mathbf{Q}}$ is the fully mixed Nash equilibrium of this instance, and $\tilde{\mathbf{P}}$ is a (mixed) Nash equilibrium where with the same structure as in case 1, that is, there are no pure users on link 1. Since $(\delta+u)^t$ is a polynomial of degree t with non-negative coefficients, we can apply case 1, proving the claim. \square

Theorem 2. *Consider the model of identical users and two links. Then*

$$\max_{\mathbf{w}, \mathbf{P}} \frac{\text{SC}_{\delta^d}(\mathbf{w}, \mathbf{P})}{\text{OPT}_{\delta^d}(\mathbf{w})} = \frac{1}{2} (2^{d-1} + 1).$$

Corollary 1. *Consider the model of identical users and two links. Then*

$$\max_{\mathbf{w}, \mathbf{P}} \frac{\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P})}{\text{OPT}_{\pi^d(\delta)}(\mathbf{w})} \leq \frac{1}{2} (2^d + d - 1).$$

5 Identical Users

In this section we first prove the validity of an approximate version of the FMNE Conjecture for the model of identical users (Theorem 3). Equipped with this result, we then give a bound on the Price of Anarchy for the special case where the polynomial cost function is just the d th power (Theorem 4), using similar techniques as in [14]. We close with an upper bound on the Price of Anarchy for general polynomial cost functions (Corollary 2).

Theorem 3. *Consider the model of identical users. Then, for any Nash equilibrium \mathbf{P} ,*

$$\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P}) \leq \left(1 + \frac{1}{n-1}\right)^d \cdot \text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{F}).$$

Theorem 4. *Consider the model of identical users. Then*

$$\max_{\mathbf{w}, \mathbf{P}} \frac{\text{SC}_{\delta^d}(\mathbf{w}, \mathbf{P})}{\text{OPT}_{\delta^d}(\mathbf{w})} \leq \left(1 + \frac{1}{n-1}\right)^d \cdot B_d.$$

The upper bound on the Price of Anarchy in Theorem 4 consists of two factors. We believe that the factor $(1 + \frac{1}{n-1})^d$ is not necessary. Note also, that the second factor B_d exceeds (asymptotically) the tight bound $\frac{1}{2}(2^{d-1} + 1)$ on the Price of Anarchy established in Theorem 2 for the case of two links.

Corollary 2. *Consider the model of identical users. Then*

$$\max_{\mathbf{w}, \mathbf{P}} \frac{\text{SC}_{\pi^d(\delta)}(\mathbf{w}, \mathbf{P})}{\text{OPT}_{\pi^d(\delta)}(\mathbf{w})} \leq \sum_{0 \leq t \leq d} \left(1 + \frac{1}{n-1}\right)^t \cdot B_t.$$

6 Pure Nash Equilibria

We now prove a tight bound on the Price of Anarchy for pure Nash equilibria for the special case that $\pi^d(\delta) = \delta^d$.

Theorem 5. *Consider the model of arbitrary users, restricted to pure Nash equilibria \mathbf{P} , and let $d \geq 2$. Then,*

$$\max_{\mathbf{w}, \mathbf{P}} \frac{\text{SC}_{\delta^d}(\mathbf{w}, \mathbf{P})}{\text{OPT}_{\delta^d}(\mathbf{w})} = \frac{(2^d - 1)^d}{(d - 1)(2^d - 2)^{d-1}} \left(\frac{d - 1}{d} \right)^d.$$

References

1. M. Aigner. *Combinatorial Theory*. Springer-Verlag, 1979.
2. E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proc. of the 35th Annual ACM Symp. on Theory of Computing*, pp. 511–520, 2003.
3. B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in worst-case equilibria. In *Proc. of the 1st Int. Workshop on Approximation and Online Algorithms*, LNCS 2909, pp. 41–52, 2003.
4. M.J. Beckmann. On the theory of traffic flow in networks. *Traffic Quart*, 21:109–116, 1967.
5. M. Beckmann, C.B. McGuire and, C.B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
6. D. Braess. Über ein Paradoxon der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.
7. A. Czumaj. Selfish routing on the internet. *To appear as a chapter in Handbook of Scheduling: Algorithms, Models, and Performance Analysis*, edited by J. Leung, CRC Press, 2004.
8. A. Czumaj, P. Krysta, and B. Vöcking. Selfish traffic allocation for server farms. In *Proc. of the 34th Annual ACM Symp. on Theory of Computing*, pp. 287–296, 2002.
9. A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *Proc. of the 13th Annual ACM-SIAM Symp. on Discrete Algorithms*, pp. 413–420, 2002.
10. A. Fabrikant, A. Luthra, E. Maneva, C.H. Papadimitriou, and S. Shenker. On a network creation game. In *Proc. of the 22nd Annual Symp. on Principles of Distributed Computing*, pp. 347–351, 2003.
11. R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Nashification and the coordination ratio for a selfish routing game. In *Proc. of the 30th Int. Colloq. on Automata, Languages, and Programming*, LNCS 2719, pp. 514–526, 2003.
12. R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish routing in non-cooperative networks: A survey. In *Proc. of the 28th Int. Symp. on Mathematical Foundations of Computer Science*, LNCS 2747, pp. 21–45, 2003.
13. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *Proc. of the 29th Int. Colloq. on Automata, Languages, and Programming*, LNCS 2380, pp. 123–134, 2002.
14. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash equilibria in discrete routing games with convex latency functions. In *Proc. of the 31st Int. Colloq. on Automata, Languages, and Programming*, 2004, to appear.

15. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Extreme nash equilibria. In *Proc. of the 8th Italian Conf. on Theoretical Computer Science*, LNCS 2841, pp. 1–20, 2003. Accepted to *Theoretical Computer Science*.
16. E. Koutsoupias. Selfish task allocation. *Bulletin of the EATCS*, 81:79–88, 2003.
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. Papadimitriou. Worst-case equilibria. In *Proc. of the 16th Int. Symp. on Theoretical Aspects of Computer Science*, LNCS 1563, pp. 404–413, 1999.
19. V.S.A. Kumar and M.V. Marathe. Improved results for stackelberg scheduling strategies. In *Proc. of the 29th Int. Colloq. on Automata, Languages, and Programming*, LNCS 2380, pp. 776–787, 2002.
20. T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. In *Proc. of the 21st Int. Symp. on Theoretical Aspects of Computer Science*, LNCS 2996, pp. 547–558, 2004.
21. T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the Worst-case Nash Equilibrium? In *Proc. of the 28th Int. Symp. on Mathematical Foundations of Computer Science*, LNCS 2747, pp. 551–561, 2003.
22. M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proc. of the 33rd Annual ACM Symp. on Theory of Computing*, pp. 510–519, 2001.
23. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13:111–124, 1996.
24. J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
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. Designing networks for selfish users is hard. In *Proc. of the 42nd Annual Symp. on Foundations of Computer Science*, pp. 472–481, 2001.
27. T. Roughgarden. Stackelberg scheduling strategies. In *Proc. of the 33rd Annual ACM Symp. on the Theory of Computing*, pp. 104–113, 2001.
28. T. Roughgarden. The price of anarchy is independent of the network topology. In *Proc. of the 34th Annual ACM Symp. on the Theory of Computing*, pp. 428–437, 2002.
29. T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
30. V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.
31. J.G. Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II, Vol. 1*, pp. 325–378, 1956.