

The Price of Anarchy for Polynomial Social Cost*

Martin Gairing[†] Thomas Lücking[†] Marios Mavronicolas[‡] Burkhard Monien[†]

*A preliminary version of this work appeared in the *Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science*, J. Fiala, V. Koubek and J. Kratochvíl eds., pp. 574–585, Vol. 3153, Lecture Notes in Computer Science, Springer-Verlag, Prague, Czech Republic, August 2004. This work has been partially supported by the IST Program of the European Union under contract numbers IST-2001-33116 (FLAGS), 001907 (DELIS) and 015964 (AEOLUS), by the DFG-Sonderforschungsbereich 376 (Massive Parallelität: Algorithmen, Entwurfsmethoden, Anwendungen), by the VEGA grant No. 2/3164/23, and by research funds at University of Cyprus.

[†]Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany. Email: {gairing,luck,bm}@uni-paderborn.de

[‡]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 for *self-ish routing* on *parallel links*, which reflects some influence from the much older *Wardrop model* [29]. In the new model, *user traffics* are still *unsplittable* and links are *identical*. *Social Cost* is now the expectation of the sum, over all links, of *Latency Costs*; each Latency Cost is modeled as a certain polynomial *latency cost function* evaluated at the *latency* incurred by all users choosing the link. The resulting Social Cost is called *Polynomial Social Cost*, or *Monomial Social Cost* when the latency cost function is a monomial. All considered polynomials are of degree d , where $d \geq 2$, and have non-negative coefficients.

We are interested in evaluating *Nash equilibria* in this model, and we use the *Monomial Price of Anarchy* and the *Polynomial Price of Anarchy* as our evaluation measures. Through establishing some remarkable relations of these costs and measures to some classical combinatorial numbers such as the *Stirling numbers of the second kind* and the *Bell numbers*, we obtain a multitude of results:

- For the special case of *identical* users:
 - The *fully mixed Nash equilibrium*, where all probabilities are strictly positive, maximizes Polynomial Social Cost.
 - The Monomial Price of Anarchy is no more than B_d , the Bell number of order d . This immediately implies that the Polynomial Price of Anarchy is no more than $\sum_{1 \leq t \leq d} B_t$.
For the special case of two links, the Monomial Price of Anarchy is no more than $2^{d-2} \left(1 + \left(\frac{1}{n} \right)^{d-1} \right)$, and this bound is *tight* for $n = 2$.
- The Monomial Price of Anarchy is *exactly* $\frac{(2^d - 1)^d}{(d-1)(2^d - 2)^{d-1}} \left(\frac{d-1}{d} \right)^d$ for *pure* Nash equilibria. This immediately implies that the Polynomial Price of Anarchy is no more than $\sum_{2 \leq t \leq d} \frac{(2^t - 1)^t}{(t-1)(2^t - 2)^{t-1}} \left(\frac{t-1}{t} \right)^t$.

1 Introduction

1.1 Framework

The *Price of Anarchy* [19, 25], also known as *Coordination Ratio*, is a widely adopted measure of the extent to which competition approximates cooperation. In the most general setting, the Price of Anarchy is the *worst-case* ratio between the value of a *global* objective function called *Social Cost* [19] over its optimal value, called *Optimum*. The Social Cost is evaluated at a *Nash equilibrium* [23, 24]; here, no *user* could unilaterally switch from its own (*mixed*) *strategy* in order to improve the value of its *local* objective function, called (*Expected*) *Individual Cost*. Yet, *Optimum* is the solution to some, usually hard, combinatorial optimization problem. So, the Price of Anarchy represents a rendezvous of Nash equilibrium, a fundamental concept from Game Theory, with *approximation*, an ubiquitous concept from Theoretical Computer Science.

Koutsoupias and Papadimitriou [19] introduced the Price of Anarchy in the context of some specific setting, widely known as the KP model and extensively studied in the last few years as a prevailing model for *selfish routing* (see, e.g., [10, 11, 13, 15, 16, 18, 21, 22]). In the KP model, there are m parallel *links* and n selfish *users* with (*unsplittable*) *traffics*. The (*expected*) *latency* incurred on a link is the (expected) total traffic of users choosing it. The (Expected) Individual Cost of a user is the (expected) latency on the link it chooses. In a Nash equilibrium, each user alone is minimizing its (Expected) Individual Cost. The Social Cost is the expectation of maximum latency; the *Optimum* is the least possible maximum latency.

The *Wardrop* model [29] is another prevailing model for selfish routing that dates back to the 1950s, when it was considered in the context of road traffic networks. In the Wardrop model, the network can be arbitrary; user traffics are infinitesimally *splittable*, and this rules out mixed strategies from consideration. In addition, Social Cost is defined here as the sum of all Individual Costs; each Individual Cost is a certain sum of *Latency Costs*. More specifically, the Latency Cost on a link is determined by a convex function, called *latency cost function*, of the latency on the link; the Individual Cost of a user is the sum of Latency Costs on links in the paths chosen by the user. Inspired by the vivid interest in the Price of Anarchy, Roughgarden and Tardos [27] initiated recently a reinvestigation of the Wardrop model.

1.2 The Model and its Relatives

A natural goal is to understand the dependence of the Price of Anarchy on the particular way of formulating Individual Cost and Social Cost. Towards this goal, some recent works [14, 20] have considered bridging the KP model with the Wardrop model and analyzing the bridged model. In this paper, we further pursue this goal by introducing and analyzing a new, interesting variant of the KP model that reflects some influence from the Wardrop model.

In our proposed model, we follow the KP model to consider the parallel links network, un-splittable traffics and mixed strategies. However, inspired by the Wardrop model, we introduce *Polynomial Social Cost* as the expectation of the sum of Latency Costs on links. We also assume that latency cost functions are *polynomial*; all polynomials we consider are of degree d and have non-negative coefficients. We assume that links are *identical*; so, all polynomials are identical. Polynomial Social Cost gives rise to *Polynomial Optimum* and *Polynomial Price of Anarchy* in the natural way. In some cases, we also consider *monomials* (that is, polynomials consisting of a single power with unit coefficient). We then talk about *Monomial Social Cost*, *Monomial Optimum* and *Monomial Price of Anarchy*.

Our model is closely related to two previously studied models:

- Relaxed to allow *arbitrary* links with (not necessarily identical) linear latencies, but restricted to *quadratic* latency cost functions, our model has been already studied by Lücking *et al.* [20]. The model in [20] adopted *Quadratic Social Cost*, which it defined (in an equivalent way) as the sum of *weighted* (Expected) Individual Costs. Quadratic Social Cost is the special case of our Monomial Social Cost with monomials of degree 2.
- Relaxed to allow arbitrary links with (not necessarily identical) convex latencies, our model was already studied by Gairing *et al.* [14]. However, Gairing *et al.* [14] modeled Social Cost as the sum of Individual Costs. Assuming identical users and linear latencies, the model of Gairing *et al.* [14] and the model of Lücking *et al.* [20] become identical.

The relation of our model to the KP model and the models of Lücking *et al.* [20] and Gairing *et al.* [14] is summarized in Figure 1. Restricted to *pure* Nash equilibria, where each user chooses a single link (with probability 1), our model was already studied for monotone latency cost functions in [9]. Besides pure Nash equilibria, we shall pay, in our study, some particular attention to the *fully mixed Nash equilibrium* [22], where each user chooses each link with non-zero probability.

Admittedly, our proposed model represents a significant departure from a long line of previous work (including [10, 14, 18, 20, 22, 27]). The reason is that, for the first time, the Social Cost is not a simple and natural function (for example, sum or maximum) of either the Individual Costs of the users or the latencies on the links. Thus, while the (Expected) Individual Cost is still the conditional expectation of the link latency, the Latency Cost on a link (which is what contributes to Social Cost) is now an (almost) arbitrary polynomial of the link latency. We argue, however, that our proposed model is significant and has potential applications in some economic scenaria:

- First, arbitrary polynomials have been long used for modeling Latency Costs in the context of studying communication and transportation networks – see, for example, [2, 6].

Model:	KP [19]	Lücking <i>et al.</i> [20]	Gairing <i>et al.</i> [14]	Present
Latency (as function of <i>total traffic</i>)	Linear	Linear	Convex	Identity
Latency Cost Function (as function of latency)	—	Quadratic	—	Polynomial
Individual Cost is expectation of:	Latency			
Social Cost is expectation of:	Maximum Latency	Sum of Latency Costs (= Weighted Sum of Individual Costs)	Sum of Individual Costs	Sum of Latency Costs

Figure 1: Comparison of present model to three relatives (namely, the KP model [19] and the models in [14, 20]). Note that all four models formulate Individual Cost (as a function of latency) in the same way, while they do not (in general) do so for latency. In the special case where the linear function for latency in the KP model and the model of Lücking *et al.* [20] becomes the identity function, the Individual Costs for these models become identical with the one for the present model, as also do their Nash equilibria.

- Second and at a more abstract level, recall that Latency Costs represent costs to the society, while Individual Costs represent costs to the individuals (making up the society). It is often the case that individuals receive support from some *authority* – for example, individuals count on *refunds* from the tax authority. This is best modeled by assuming that the actual Individual Costs are significantly lower than those corresponding to the Social Cost.
- Third, Social Cost is often overestimated in order to allow claims for higher support from funding agencies. Using higher degree polynomials provides a paradigm for such overestimation.

1.3 Contribution and Significance

We are primarily interested in analyzing the Polynomial Price of Anarchy for our new model. To do so, we introduce and study a natural conjecture, called the *Polynomial Fully Mixed Nash Equilibrium Conjecture* and abbreviated as the PFMNE Conjecture; it asserts that the fully mixed Nash equilibrium maximizes Polynomial Social Cost. Although the PFMNE conjecture is interesting in its own right, a resolution of it to the positive would also enable the derivation of upper bounds on the Polynomial Price of Anarchy via deriving upper bounds on the Polynomial

Social Cost of the fully mixed Nash equilibrium.

We address two important special settings of the problem:

1.3.1 Identical Users

For the case of identical users, we rely on a very thorough analysis of the fully mixed Nash equilibrium. For the analysis, we employ, as our chief combinatorial instrument, the *binomial function* originally introduced in [14]. We prove here that the binomial function can be expressed as a combinatorial sum of *Stirling numbers of the second kind* [28] (Proposition 2.2). We also observe that Polynomial Social Cost can be expressed as a sum of binomial functions (in the case of identical users). Together these two imply that the Polynomial Social Cost of the fully mixed Nash equilibrium is a combinatorial sum of Stirling numbers of the second kind (Corollary 3.2). Moreover, the Polynomial Social Cost of *any* (mixed) Nash equilibrium is upper bounded by a certain combinatorial sum of Stirling numbers of the second kind. Hence, comparison of these two Polynomial Social Costs reduces to comparing like terms in the two combinatorial sums. We obtain the following results:

- The PFMNE Conjecture is valid (Theorem 4.1). The proof follows a careful comparison of like terms in the combinatorial sums upper bounding and expressing the Polynomial Social Costs of an arbitrary and the fully mixed Nash equilibria, respectively.
- The Monomial Price of Anarchy is upper bounded by B_d , (Theorem 4.4); here, B_d is the *Bell number* of order d . This follows from the PFMNE Conjecture. From this bound, an upper bound of $\sum_{2 \leq t \leq d} B_t$ on Polynomial Price of Anarchy follows immediately (Corollary 4.5).

For the special case of two links, the Monomial Price of Anarchy is upper bounded by $2^{d-2} \left(1 + \left(\frac{1}{n} \right)^{d-1} \right)$ (Theorem 4.6). Furthermore, this upper bound is *tight* for the sub-case of two users. From this upper bound, an upper bound of $2^{d-1} - 1 + \frac{d-1}{n}$ on Polynomial Price of Anarchy follows immediately (Corollary 4.7).

1.3.2 Pure Nash Equilibria

The Monomial Price of Anarchy is *exactly* $\frac{(2^d - 1)^d}{(d-1)(2^d - 2)^{d-1}} \left(\frac{d-1}{d} \right)^d$ (Theorem 5.4). The asymptotic behavior of this exact bound is closely described by the simple function $\frac{(2^d)^d}{d(2^d)^{d-1}} = \frac{2^d}{d}$. From this exact bound, an upper bound of $\sum_{2 \leq t \leq d} \frac{(2^t - 1)^t}{(t-1)(2^t - 2)^{t-1}} \left(\frac{t-1}{t} \right)^t$ on Polynomial Price of Anarchy follows immediately (Corollary 5.5).

Bounds for Identical Users:		
Monomial Price of Anarchy		Polynomial Price of Anarchy
Lower	Upper	Upper
—	B_d , if $m \geq 2$	$\sum_{2 \leq t < d} B_t$, if $m \geq 2$
$2^{d-2} + \frac{1}{2}$, if $m = n = 2$	$2^{d-2} \left(1 + \left(\frac{1}{n}\right)^{d-1}\right)$, if $m = 2$	$2^{d-1} - 1 + \frac{d-1}{n}$, if $m = 2$
Bounds for Pure Nash Equilibria:		
Monomial Price of Anarchy		Polynomial Price of Anarchy
Lower	Upper	Upper
$\frac{(2^d - 1)^d}{(d-1)(2^d - 2)^{d-1}} \left(\frac{d-1}{d}\right)^d$	$\frac{(2^d - 1)^d}{(d-1)(2^d - 2)^{d-1}} \left(\frac{d-1}{d}\right)^d$	$\sum_{2 \leq t \leq d} \frac{(2^t - 1)^t}{(t-1)(2^t - 2)^{t-1}} \left(\frac{t-1}{t}\right)^t$

Figure 2: Summary of shown bounds on Monomial and Polynomial Prices of Anarchy.

1.3.3 Summary and Remarks

All shown bounds are summarized in Figure 2. We remark that (almost) all shown bounds are independent of m and n , but depend on d . The lower bounds imply that this dependence is inherent. Finally, we remark that all upper bounds on Polynomial Price of Anarchy are obtained through naive reductions to corresponding upper bounds on Monomial Price of Anarchy. At present, we do not know if there are better bounds on Polynomial Price of Anarchy that bypass the naive reduction.

1.4 Related Work and Comparison

Gairing *et al.* [15, 16] were the *first* to explicitly state the related *Fully Mixed Nash Equilibrium Conjecture* that the fully mixed Nash equilibrium maximizes Social Cost for the KP model. Up to now, the conjecture has been proved for several particular cases of the KP model [13, 15, 21]. Recently, Fischer and Vöcking [12] presented a counterexample to the Fully Mixed Nash Equilibrium Conjecture for the case of identical links. The validity of the Fully Mixed Nash Equilibrium Conjecture for the case of identical users (but arbitrary links) remains open.

Lücking *et al.* [20] formulated the *Quadratic Fully Mixed Nash Equilibrium Conjecture*, which asserts that the fully mixed Nash equilibrium maximizes Quadratic Social Cost for their model. Lücking *et al.* [20, Theorem 4.8] proved the Quadratic Fully Mixed Nash Equilibrium Conjecture for the case of identical users and identical links. Our PFMNE Conjecture generalizes the Quadratic Fully Mixed Nash Equilibrium Conjecture of Lücking *et al.* [20] to polynomial latency cost functions of arbitrary degree. Gairing *et al.* [14] also formulated a corresponding conjecture for their model, stating that the fully mixed Nash equilibrium maximizes Social

Cost (sum of Individual Costs) for their model. Gairing *et al.* [14, Theorem 3.5] proved their conjecture for the case of identical users and arbitrary links with non-decreasing, non-constant and convex latencies.

Our exact bound on Monomial Price of Anarchy for pure Nash equilibria includes, as the special case with $d = 2$, the exact bound of $\frac{9}{8}$ on Quadratic Price of Anarchy for pure Nash equilibria shown in [20, Theorem 5.2]. Our proof generalizes the one for [20, Theorem 5.2].

Our upper bound of $2^{d-2} \left(1 + \left(\frac{1}{n}\right)^{d-1}\right)$ on Monomial Price of Anarchy for the case of identical users and two (identical) links implies, as the special case where $d = 2$, an upper bound of $1 + \frac{1}{n} \leq \frac{3}{2}$ on Quadratic Price of Anarchy. This complements the corresponding exact bound of $\frac{4}{3}$ shown in [20, Theorem 5.1] for the case of identical users and pure Nash equilibria (but for arbitrarily many *related* links).

Our upper bound of B_d for the case of identical users implies, as the special case where $d = 2$, an upper bound of $B_2 = 2$. The implied upper bound exceeds the corresponding upper bound of $1 + \min\left\{\frac{m-1}{n}, \frac{n-1}{m}\right\} < 2$ on Quadratic Price of Anarchy shown in [20, Theorem 5.4]. So, our upper bound of B_d for the case of identical users is *not* tight for the particular case where $d = 2$.

Other bounds on Price of Anarchy include tight asymptotic bounds (depending on m) for the KP model [10, 18] and exact constant bounds for the Wardrop model [27]. Some recent works [1, 4, 7, 8] have provided tight (and even *exact*) bounds on Price of Anarchy for *congestion games* [26] and their variants. (The KP model is itself a special case of congestion games.) These works have considered both Nash and *correlated equilibria* [3], linear and polynomial latencies and different Social Cost functions.

1.5 Road Map

Section 2 summarizes some mathematical preliminaries. Section 3 introduces our theoretical model. The case of identical users is considered in Section 4. Pure Nash equilibria are treated in Section 5. We conclude, in Section 6, with a discussion of our results and some open problems.

2 Mathematical Preliminaries

Throughout, denote for any integer $k \geq 1$, $[k] = \{1, \dots, k\}$. A *monomial* function $g : \mathbb{R} \rightarrow \mathbb{R}$ has the form $g(\lambda) = \lambda^d$ for some integer $d \geq 0$. A *polynomial* function is a linear combination of monomials. We shall only consider polynomial functions with non-negative coefficients. For a random variable X with associated probability distribution \mathbf{P} , denote $\mathbb{E}_{\mathbf{P}}(X)$ the *expectation* of X . In some later proof, we shall make use of the following simple mathematical fact that

follows directly from the convexity of the monomial function $\pi_d(\lambda) = \lambda^d$.

Lemma 2.1 *Let $x, y_1, y_2 \in \mathbb{R}$ with $0 < x \leq y_1 < y_2 + x$. Then, for each integer $d \geq 2$,*

$$(y_1 - x)^d + (y_2 + x)^d > y_1^d + y_2^d.$$

2.1 Falling Factorials, Stirling Numbers and Bell Numbers

For any pair of integers $k \geq 1$ and $t \geq 1$, the *falling factorial of k order t* , denoted as $k^{\underline{t}}$, is given by $k^{\underline{t}} = \underbrace{k \cdot (k-1) \cdot \dots \cdot (k-(t-1))}_{t \text{ factors}}$, when $k \geq t$. Otherwise ($t \geq k+1$), $k^{\underline{t}} = 0$.

For any pair of integers $d \geq 1$ and $t \in [d] \cup \{0\}$, the *Stirling number of the second kind*, denoted as $S(d, t)$, counts the number of partitions of a set with d elements into exactly t blocks (non-empty subsets). In particular, $S(d, 1) = 1$. Also, for all integers $d \geq 2$, $S(d, 2) = 2^{d-1} - 1$. Stirling numbers of the second kind satisfy the recurrence relation

$$S(d, t) = \sum_{q: t \leq q \leq d-1} \binom{d-1}{q-1} \cdot S(q-1, t-1)$$

for all integers $d \geq 2$ and $t \in [d]$ (see, e.g., [17, Table 265, Identity (6.15)]). It is also known that for all integers $d \geq 2$ and $k \geq 1$, $k^d = \sum_{t \in [d]} S(d, t) \cdot k^{\underline{t}}$.

For any integer $d \geq 1$, the *Bell number* of order d [5], denoted as B_d , counts the number of partitions of a set with d elements into blocks. So, clearly, $B_0 = 1$ and $B_d = \sum_{t \in [d]} S(d, t)$.

2.2 Binomial Function

We start with the definition of a binomial function [14, Definition 1].

Definition 2.1 *For any integer $r \geq 1$, consider a vector of probabilities $\mathbf{p} = \langle p_1, \dots, p_r \rangle$. Fix a function $g(\lambda) : \mathbb{R} \rightarrow \mathbb{R}$. Then, the binomial function $\text{BF}(\mathbf{p}, g)$ is given by*

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

Strictly speaking, Definition 2.1 defines a *functional*. If all probabilities have the same value p , then we talk about a *constant* vector of probabilities, and we (abuse notation to) write $\text{BF}(p, r, g)$. Clearly, in this case,

$$\text{BF}(p, r, g) = \sum_{0 \leq k \leq r} \binom{r}{k} p^k (1-p)^{r-k} g(k).$$

We show that when g is monomial, the binomial function takes a special form.

Proposition 2.2 For each integer $d \geq 1$,

$$\text{BF}(p, r, \lambda^d) = \sum_{t \in [d]} p^t \cdot S(d, t) \cdot r^t.$$

Proof: The proof will first establish a recurrence relation for $\text{BF}(p, r, \lambda^d)$. This relation will then become instrumental for carrying out an inductive proof of the claim. We continue with the details of the formal proof. Clearly,

$$\begin{aligned} & \text{BF}(p, r, \lambda^d) \\ = & \sum_{k \in [r]} \binom{r}{k} p^k (1-p)^{r-k} k^d && \text{(by definition of binomial function)} \\ = & \sum_{k \in [r]} \frac{r}{k} \binom{r-1}{k-1} p^k (1-p)^{r-k} k^d && \text{(since } \binom{r}{k} = \frac{r}{k} \binom{r-1}{k-1} \text{)} \\ = & p \cdot r \cdot \sum_{k \in [r]} \binom{r-1}{k-1} p^{k-1} (1-p)^{r-k} k^{d-1} \\ = & p \cdot r \cdot \sum_{0 \leq k \leq r-1} \binom{r-1}{k} p^k (1-p)^{r-1-k} (k+1)^{d-1} && \text{(by change of variable)} \\ = & p \cdot r \cdot \sum_{0 \leq k \leq r-1} \binom{r-1}{k} p^k (1-p)^{r-1-k} \left(\sum_{0 \leq q \leq d-1} \binom{d-1}{q} k^q \right) && \text{(by the binomial theorem)} \\ = & p \cdot r \cdot \sum_{0 \leq q \leq d-1} \binom{d-1}{q} \left(\sum_{0 \leq k \leq r-1} \binom{r-1}{k} p^k (1-p)^{r-1-k} k^q \right) && \text{(by changing order of summation)} \\ = & p \cdot r \cdot \sum_{0 \leq q \leq d-1} \binom{d-1}{q} \text{BF}(p, r-1, \lambda^q) && \text{(by definition of binomial function).} \end{aligned}$$

We now use the obtained recurrence relation for the binomial function to prove the claim by induction on r . For the basis case, let $r = 1$. Then, $\text{BF}(p, 1, \lambda^d) = \binom{1}{1} p^1 1^d = p$ and $\sum_{t \in [d]} p^t S(d, t) 1^t = p^1 S(d, 1) 1^1 = p$, so that the claim follows. Assume inductively that the claim holds for some integer $r-1 \geq 1$. For the induction step, note that

$$\begin{aligned} & \text{BF}(p, r, \lambda^d) \\ = & p \cdot r \cdot \binom{d-1}{0} \text{BF}(p, r-1, 1) \\ & + p \cdot r \cdot \sum_{q \in [d-1]} \binom{d-1}{q} \text{BF}(p, r-1, \lambda^q) && \text{(by recurrence relation)} \\ = & p \cdot r + p \cdot r \cdot \sum_{q \in [d-1]} \binom{d-1}{q} \text{BF}(p, r-1, \lambda^q) && \text{(by definition of binomial function)} \\ = & p \cdot r + p \cdot r \cdot \sum_{q \in [d-1]} \binom{d-1}{q} \left(\sum_{t \in [q]} p^t \cdot S(q, t) \cdot (r-1)^t \right) && \text{(by induction hypothesis)} \\ = & p \cdot r + \sum_{q \in [d-1]} \binom{d-1}{q} \left(\sum_{t \in [q]} p^{t+1} \cdot S(q, t) \cdot r^{t+1} \right) \\ = & p \cdot r + \sum_{t \in [d-1]} p^{t+1} \cdot r^{t+1} \cdot \left(\sum_{q: t \leq q \leq d-1} \binom{d-1}{q} \cdot S(q, t) \right) && \text{(by changing order of summation)} \\ = & p \cdot r + \sum_{2 \leq t \leq d} p^t \cdot r^t \cdot \left(\sum_{q: t \leq q \leq d} \binom{d-1}{q-1} \cdot S(q-1, t-1) \right) && \text{(by change of variables)} \\ = & p \cdot r + \sum_{2 \leq t \leq d} p^t \cdot r^t \cdot S(d, t) && \text{(by recurrence relation for } S(d, t) \text{)} \\ = & \sum_{t \in [d]} p^t \cdot r^t \cdot S(d, t) && \text{(since } S(d, 1) = 1 \text{).} \end{aligned}$$

as needed. ■

Proposition 2.2 implies that for a constant vector of probabilities and a monomial function, the binomial function is a combinatorial sum of Stirling numbers of the second kind.

It is known [14, Lemma 3] that in case g is *convex*, the binomial function does not decrease when replacing all probabilities in the vector of probabilities \mathbf{p} by the average probability $\tilde{p} = \frac{\sum_{i \in [r]} p_i}{r}$.

Lemma 2.3 ([14]) *For a convex function g , $\text{BF}(\mathbf{p}, g) \leq \text{BF}(\tilde{p}, r, g)$.*

3 Model and Preliminaries

Our model definitions are built on top of those for the KP model [19], which are extended to accommodate some features from the Wardrop model [29].

3.1 General

We consider a simple *network* consisting of a set of m parallel *links* $1, 2, \dots, m$ from a *source* node to a *destination* node. Each of n *users* $1, 2, \dots, n$ wishes to route a traffic along a (non-fixed) link from source to destination. Denote w_i the *traffic* of user $i \in [n]$; denote $W = \sum_{i \in [n]} w_i$. Define the $n \times 1$ *traffic vector* \mathbf{w} in the natural way. We assume that all links are *identical*. Thus, an *instance* is a tuple $\langle \mathbf{w}, m \rangle$. In the model of *identical users*, all traffics are equal to 1. In that case, an instance is a pair $\langle n, m \rangle$. Assume throughout that $m \geq 2$ and $n \geq 2$.

The *latency* λ on a link is the total traffic on it. Associated with each link is a *latency cost function*, which is a polynomial $\pi_d(\lambda) = \sum_{0 \leq t \leq d} a_t \lambda^t$ of degree $d \geq 2$ with non-negative coefficients. In the special case of a monomial, $\pi_d(\lambda) = \lambda^d$. The *Latency Cost* for latency λ on the link is given by $\pi_d(\lambda)$.

3.2 Strategies and Assignments

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; so, it is a probability distribution over links.

A *pure assignment* is an n -tuple $\mathbf{L} = \langle \ell_1, \ell_2, \dots, \ell_n \rangle \in [m]^n$; a *mixed assignment* is an $n \times m$ *probability matrix* \mathbf{P} of nm probabilities p_{ij} , with $i \in [n]$ and $j \in [m]$, where p_{ij} is the probability that user i chooses link j . Note that a mixed assignment induces a probability distribution on the set of pure assignments. For each link $j \in [m]$, denote $r_j = |\{i \in [n] \mid p_{ij} > 0\}|$. Consider now a link $j \in [m]$ such that $r_j > 0$. Then, the *average probability* \hat{p}_j on link j is $\tilde{p}_j = \frac{\sum_{i \in [n]} p_{ij}}{r_j}$. A mixed assignment \mathbf{P} is *fully mixed* [22, Section 2.2] if for all users $i \in [n]$ and links $j \in [m]$, $p_{ij} > 0$.

Fix now a mixed assignment \mathbf{P} . The *latency* $\lambda_j(\mathbf{P})$ on link $j \in [m]$ induced by \mathbf{P} is the total traffic assigned to the link according to \mathbf{P} ; so, $\lambda_j(\mathbf{P})$ is a random variable. Denote $\Lambda_j(\mathbf{P})$ the *expected latency* on link $j \in [m]$; thus, $\Lambda_j(\mathbf{P}) = \mathbb{E}_{\mathbf{P}}(\lambda_j(\mathbf{P})) = \sum_{i \in [n]} p_{ij} w_i$.

3.3 Cost Measures

3.3.1 Individual Cost and Expected Individual Cost

For a pure assignment \mathbf{L} , the *Individual Cost* for user $i \in [n]$, denoted as $\text{IC}_i(\mathbf{L})$, is $\text{IC}_i(\mathbf{L}) = \Lambda_{\ell_i}(\mathbf{L})$; so, the Individual Cost for a user is the latency of the link it chooses. For a mixed assignment \mathbf{P} , the *Expected Individual Cost* for user $i \in [n]$, denoted again as $\text{IC}_i(\mathbf{P})$, is the expectation according to \mathbf{P} of the Individual Cost for the user.

The *Conditional Expected Individual Cost* $\text{IC}_{ij}(\mathbf{P})$ for user $i \in [n]$ on link $j \in [m]$ is the conditional expectation according to \mathbf{P} of the Individual Cost of user i had it been assigned to link j . So, $\text{IC}_{ij}(\mathbf{P}) = \Lambda_j(\mathbf{P}) + (1 - p_{ij})w_i$. Clearly, for each user $i \in [n]$, $\text{IC}_i(\mathbf{P}) = \sum_{j \in [m]} p_{ij} \text{IC}_{ij}(\mathbf{P})$.

3.3.2 Polynomial Social Cost

Associated with an instance $\langle \mathbf{w}, m \rangle$, a latency cost function $\pi_d(\lambda)$ and a mixed assignment \mathbf{P} is the *Polynomial Social Cost*, denoted $\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P})$, which is the expectation of the sum of Latency Costs; so, by linearity of expectation,

$$\begin{aligned} \text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P}) &= \mathbb{E}_{\mathbf{P}} \left(\sum_{j \in [m]} \pi_d \left(\sum_{k \in [n] | \ell_k = j} w_k \right) \right) \\ &= \sum_{j \in [m]} \mathbb{E}_{\mathbf{P}} \left(\pi_d \left(\sum_{k \in [n] | \ell_k = j} w_k \right) \right) \\ &= \sum_{j \in [m]} \sum_{A \subseteq [n]} \left(\prod_{i \in A} p_{ij} \right) \left(\prod_{i \notin A} (1 - p_{ij}) \right) \pi_d \left(\sum_{k \in A | \ell_k = j} w_k \right). \end{aligned}$$

The displayed formulas for Polynomial Social Cost refer to a pure assignment $\mathbf{L} = \langle \ell_1, \dots, \ell_n \rangle$ drawn according to the probability distribution (on the set of pure assignments) induced by the mixed assignment \mathbf{P} . Note that

$$\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P}) = \sum_{0 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P}).$$

So, Polynomial Social Cost is a linear combination (with non-negative coefficients) of *Monomial Social Costs*. This property will later reduce the comparison of the Polynomial Social Costs of two different assignments to the pairwise comparison of their Monomial Social Costs.

We remark that the Polynomial Social Cost is a generalization of the Quadratic Social Cost [20] to latency cost functions that are polynomials of arbitrary degree.

3.3.3 Polynomial Optimum

Associated with an instance $\langle \mathbf{w}, m \rangle$ and a latency cost function $\pi_d(\lambda)$ is the *Polynomial Optimum*, denoted $\text{POPT}_{\pi_d(\lambda)}(\mathbf{w}, m)$, which is the least possible, over all pure assignments, Polynomial Social Cost; thus,

$$\text{POPT}_{\pi_d(\lambda)}(\mathbf{w}, m) = \min_{\mathbf{L} \in [m]^n} \text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{L}).$$

A (pure) assignment \mathbf{L} such that $\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{L}) = \text{POPT}_{\pi_d(\lambda)}(\mathbf{w}, m)$ will be called *optimal* (for the instance $\langle \mathbf{w}, m \rangle$ and the latency cost function $\pi_d(\lambda)$). We remark that the Polynomial Optimum is a generalization of the Quadratic Optimum [20] to latency cost functions that are polynomials of arbitrary degree. *Monomial Optimum* is defined as the natural special case of Polynomial Optimum.

3.4 Nash Equilibria

Given an instance $\langle \mathbf{w}, m \rangle$, the mixed assignment \mathbf{P} is a *Nash equilibrium* [19, Section 2] if for each user $i \in [n]$, it minimizes the Expected Individual Cost $\text{IC}_i(\mathbf{P})$ over all mixed assignments \mathbf{Q} that differ from \mathbf{P} only with respect to the mixed strategy of user i ; that is, for all such mixed assignments \mathbf{Q} , $\text{IC}_i(\mathbf{P}) \leq \text{IC}_i(\mathbf{Q})$. Thus, in a Nash equilibrium, there is no incentive for a user to unilaterally deviate from its mixed strategy.

We remark that latency and (Expected) Individual Cost are defined for our model in the same way they are defined for the KP model (and for the model of Lücking *et al.* [20] as well) in the case of identical links. Thus, the sets of Nash equilibria for the two models coincide.

The particular definition of Expected Individual Cost implies that in a Nash equilibrium, for each user $i \in [n]$ and link $j \in [m]$ such that $p_{ij} > 0$, all Conditional Expected Individual Costs $\text{IC}_{ij}(\mathbf{P})$ are the same and no more than any Conditional Expected Individual Cost $\text{IC}_{il}(\mathbf{P})$ with $p_{il} = 0$.

3.5 The Fully Mixed Nash Equilibrium

For the KP model, it is known [22] that the fully mixed Nash equilibrium \mathbf{F} exists uniquely in the case of identical links (with $f_{ij} = \frac{1}{m}$ for all users $i \in [n]$ and links $j \in [m]$). As the set of Nash equilibria in the KP model (in the case of identical links) and the present model coincide, the same holds for the fully mixed Nash equilibria \mathbf{F} in our model.

We formulate a natural conjecture related to Polynomial Social Costs of Nash equilibria in our model, called the *Polynomial Fully Mixed Nash Equilibrium Conjecture* and abbreviated as the PFMNE Conjecture.

Conjecture 3.1 (Polynomial Fully Mixed Nash Equilibrium Conjecture) *For any instance $\langle \mathbf{w}, m \rangle$ and associated Nash equilibrium \mathbf{P} , $\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P}) \leq \text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{F})$.*

The PFMNE Conjecture generalizes the *Quadratic Fully Mixed Nash Equilibrium Conjecture* to latency cost functions that are polynomials of arbitrary degree. It is also a variant of the well known *Fully Mixed Nash Equilibrium Conjecture* [15, 16] for the original KP model.

3.6 Monomial and Polynomial Price of Anarchy

The *Polynomial Price of Anarchy*, denoted PPOA, is the *worst-case* ratio $\frac{\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P})}{\text{POPT}_{\pi_d(\lambda)}(\mathbf{w}, m)}$ over all instances $\langle \mathbf{w}, m \rangle$ and associated Nash equilibria \mathbf{P} . This generalizes the *Quadratic Price of Anarchy* [20] to latency cost functions that are polynomials of arbitrary degree. The *Monomial Price of Anarchy*, denoted MPOA, is the natural special case of the Polynomial Price of Anarchy.

The following simple fact will be instrumental for reducing the Polynomial Price of Anarchy for arbitrary polynomials (with non-negative coefficients) to the Monomial Price of Anarchy.

Lemma 3.1 (From Polynomials to Monomials) *Fix any instance $\langle \mathbf{w}, m \rangle$ with an associated Nash equilibrium \mathbf{P} . Then,*

$$\frac{\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P})}{\text{POPT}_{\pi_d(\lambda)}(\mathbf{w}, m)} \leq \sum_{2 \leq t \leq d} \frac{\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\text{POPT}_{\lambda^t}(\mathbf{w}, m)}.$$

Proof: Our proof will use the expression of Polynomial Social Cost as a linear combination of Monomial Social Costs (see Section 3.3.2). We will manipulate sums of fractions while relying on the non-negativeness of the coefficients in the latency cost function. We continue with the details of the formal proof. Let \mathbf{Q} be an optimal assignment for the instance $\langle \mathbf{w}, m \rangle$. Then,

$$\begin{aligned} \frac{\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P})}{\text{POPT}_{\pi_d(\lambda)}(\mathbf{w}, m)} &= \frac{\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P})}{\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{Q})} \\ &= \frac{a_0 + a_1 \cdot \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{P}) + \sum_{2 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{a_0 + a_1 \cdot \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{Q}) + \sum_{2 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})}. \end{aligned}$$

Note that $\text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{P}) = \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{P}) = W$, so that $a_0 + a_1 \cdot \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{P}) = a_0 + a_1 \cdot \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{P})$. Since \mathbf{Q} is an optimal assignment, $\text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{P}) \geq \text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{Q})$, which implies that $\sum_{2 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P}) \geq \sum_{2 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})$. Since $a_0 + a_1 \cdot \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{P}) = a_0 + a_1 \cdot \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{P}) \geq 0$, and we consider polynomials with non-negative coefficients, this implies that

$$\begin{aligned} \frac{a_0 + a_1 \cdot \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{P}) + \sum_{2 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{a_0 + a_1 \cdot \text{PSC}_{\lambda^1}(\mathbf{w}, m, \mathbf{Q}) + \sum_{2 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} &\leq \frac{\sum_{2 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\sum_{2 \leq t \leq d} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} \\ &= \frac{\sum_{2 \leq t \leq d | a_t > 0} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\sum_{2 \leq t \leq d | a_t > 0} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})}. \end{aligned}$$

Clearly, all terms $a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})$ and $a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})$ in the last fraction are strictly positive, and this implies that

$$\begin{aligned} \frac{\sum_{2 \leq t \leq d |a_t > 0} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\sum_{2 \leq t \leq d |a_t > 0} a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} &\leq \sum_{2 \leq t \leq d |a_t > 0} \frac{a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{a_t \cdot \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} \\ &= \sum_{2 \leq t \leq d |a_t > 0} \frac{\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} \\ &\leq \sum_{2 \leq t \leq d} \frac{\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})}. \end{aligned}$$

By definition of Monomial Optimum, $\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q}) \geq \text{POPT}_{\lambda^t}(\mathbf{w}, m)$. Hence,

$$\sum_{2 \leq t \leq d} \frac{\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})} \leq \sum_{2 \leq t \leq d} \frac{\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{P})}{\text{POPT}_{\lambda^t}(\mathbf{w}, m)}.$$

Combining now all inequalities yields the claim. \blacksquare

3.7 Identical Users

Restricted to identical users, Polynomial Social Cost reduces to

$$\begin{aligned} \text{PSC}_{\pi_d(\lambda)}(n, m, \mathbf{P}) &= \sum_{j \in [m]} \sum_{A \subseteq [n]} \left(\prod_{i \in A} p_{ij} \right) \left(\prod_{i \notin A} (1 - p_{ij}) \right) \pi_d(|A|) \\ &= \sum_{j \in [m]} \text{BF}(\langle p_{1j}, \dots, p_{nj} \rangle, \pi^d(\lambda)) \end{aligned}$$

So, Polynomial Social Cost is now a sum of binomial functions, one for each link. Recall that in the case of identical users, all probabilities are identical (and equal to $\frac{1}{m}$) for the fully mixed Nash equilibrium \mathbf{F} . Hence, Proposition 2.2 implies now that the Monomial Social Cost of the fully mixed Nash equilibrium \mathbf{F} is a combinatorial sum of Stirling numbers of the second kind.

Corollary 3.2 *Consider the case of identical users. Fix an instance $\langle n, m \rangle$. Then,*

$$\text{PSC}_{\lambda^d}(n, m, \mathbf{F}) = m \sum_{t \in [d]} \left(\frac{1}{m} \right)^t \cdot S(d, t) \cdot n^t.$$

A lower bound on Monomial Optimum for the case of identical users is $\text{POPT}_{\lambda^d}(\mathbf{w}, m) \geq m \left(\frac{n}{m} \right)^d$ if $n \geq m$, while $\text{POPT}_{\lambda^d}(n, m) = n$ if $n < m$.

4 Identical Users

The PFMNE Conjecture is considered in Section 4.1. Bounds on the Monomial and Polynomial Prices of Anarchy are presented in Section 4.2.

4.1 The PFMNE Conjecture

We prove the validity of the PFMNE Conjecture.

Theorem 4.1 *Consider the case of identical users. Then, the PFMNE Conjecture is valid.*

Proof: Fix an instance $\langle n, m \rangle$ with associated Nash equilibrium \mathbf{P} and fully mixed Nash equilibrium \mathbf{F} . Since Polynomial Social Cost is a linear combination (with non-negative coefficients) of Monomial Social Costs, it suffices to prove the PFMNE Conjecture for a monomial latency cost function $\pi_d(\lambda) = \lambda^d$.

Denote $\alpha = \frac{n}{m}$. Assume, without loss of generality, that for each link $j \in [m]$, $r_j = |\{i \in [n] : p_{ij} > 0\}| \geq 1$. Clearly, the average probability on link j is $\frac{\Lambda_j(\mathbf{P})}{r_j}$.

We start with an informal outline of our proof. We will separately calculate the Polynomial Social Costs of \mathbf{P} and \mathbf{F} ; we will express their difference as a linear combination (with non-negative coefficients) of terms, and we will use induction to prove that each term is non-negative. The inductive proof will establish and use an upper bound on the average probability for a link. We now continue with the details of the formal proof. On one hand,

$$\begin{aligned}
& \text{PSC}_{\lambda^d}(n, m, \mathbf{P}) \\
&= \sum_{j \in [m]} \text{BF}(\langle p_{1j}, \dots, p_{nj} \rangle, \lambda^d) \\
&\leq \sum_{j \in [m]} \text{BF}\left(\frac{\Lambda_j(\mathbf{P})}{r_j}, r_j, \lambda^d\right) \quad (\text{by Lemma 2.3}) \\
&= \sum_{j \in [m]} \sum_{t \in [d]} \left(\frac{\Lambda_j(\mathbf{P})}{r_j}\right)^t \cdot S(d, t) \cdot (r_j)^{\underline{t}} \quad (\text{by Proposition 2.2}) \\
&= \sum_{t \in [d]} S(d, t) \cdot \left(\sum_{j \in [m]} \left(\frac{\Lambda_j(\mathbf{P})}{r_j}\right)^t \cdot (r_j)^{\underline{t}}\right) \quad (\text{by changing order of summation}).
\end{aligned}$$

On the other hand,

$$\begin{aligned}
& \text{PSC}_{\lambda^d}(n, m, \mathbf{F}) \\
&= m \cdot \sum_{t \in [d]} \left(\frac{1}{m}\right)^t \cdot S(d, t) \cdot n^{\underline{t}} \quad (\text{by Corollary 3.2}) \\
&= \sum_{t \in [d]} S(d, t) \cdot m \alpha^t \cdot \frac{n^{\underline{t}}}{n^{\underline{t}}}.
\end{aligned}$$

So, clearly,

$$\text{PSC}_{\lambda^d}(n, m, \mathbf{F}) - \text{PSC}_{\lambda^d}(n, m, \mathbf{P}) \geq \sum_{t \in [d]} S(d, t) \cdot \Delta(t),$$

where for each integer $t \in [d]$,

$$\Delta(t) = m \alpha^t \cdot \frac{n^{\underline{t}}}{n^{\underline{t}}} - \sum_{j \in [m]} \left(\frac{\Lambda_j(\mathbf{P})}{r_j}\right)^t \cdot (r_j)^{\underline{t}}.$$

We prove:

Lemma 4.2 For each integer $t \geq 1$, $\Delta(t) \geq 0$.

Proof: Assume, without loss of generality, that for each integer $t \in [d]$, for each link $j \in [m]$, $r_j \geq t$ (since otherwise $r_j^t = 0$ and $\Delta(t)$ can only increase). The proof is by induction on t . For the basis case where $t = 1$, note that

$$\begin{aligned}\Delta(1) &= m\alpha - \sum_{j \in [m]} \frac{\Lambda_j(\mathbf{P})}{r_j} \cdot r_j \\ &= n - \sum_{j \in [m]} \Lambda_j(P) \\ &= 0,\end{aligned}$$

as needed.

Assume inductively that the claim holds for $(t-1)$, for some integer $t \geq 2$. For the induction step, we will prove the claim for t . We first prove a preliminary claim:

Lemma 4.3 For each integer $t \geq 2$, for each link $j \in [m]$ such that $r_j \geq t$,

$$\frac{\Lambda_j(\mathbf{P})}{r_j} (r_j - (t-1)) \leq \alpha \cdot \frac{n - (t-1)}{n}.$$

Proof: Fix a link $j \in [m]$ and a user $i \in [n]$ such that $0 < p_{ij} \leq \frac{\Lambda_j(\mathbf{P})}{r_j}$. (Clearly, such a user exists.) Since \mathbf{P} is a Nash equilibrium and $p_{ij} > 0$, it follows that for each link $\ell \in [m]$, $lC_{ij}(\mathbf{P}) \leq lC_{i\ell}(\mathbf{P})$, or $\Lambda_j(\mathbf{P}) - p_{ij} \leq \Lambda_\ell(\mathbf{P}) - p_{i\ell}$. Since $p_{ij} \leq \frac{\Lambda_j(\mathbf{P})}{r_j}$, it follows that $\Lambda_j(\mathbf{P}) - p_{ij} \geq \Lambda_j(\mathbf{P}) - \frac{\Lambda_j(\mathbf{P})}{r_j} = \frac{r_j - 1}{r_j} \Lambda_j(\mathbf{P})$. Hence, it follows that

$$\frac{r_j - 1}{r_j} \Lambda_j(\mathbf{P}) \leq \Lambda_\ell(\mathbf{P}) - p_{i\ell}.$$

Summing up over all links $\ell \in [m]$ yields that

$$\begin{aligned}\sum_{\ell \in [m]} \frac{r_j - 1}{r_j} \Lambda_j(\mathbf{P}) &\leq \sum_{\ell \in [m]} (\Lambda_\ell(\mathbf{P}) - p_{i\ell}) \\ &= \sum_{\ell \in [m]} \Lambda_\ell(\mathbf{P}) - \sum_{\ell \in [m]} p_{i\ell} \\ &= n - 1.\end{aligned}$$

Since $r_j \geq t$ and $t \geq 2$, $r_j \geq 2$. Hence, it follows that

$$\Lambda_j(\mathbf{P}) \leq \frac{n-1}{m} \cdot \frac{r_j}{r_j-1}.$$

Since $r_j \geq t$, $r_j - (t - 1) > 0$; hence, it follows that

$$\frac{\Lambda_j(\mathbf{P})}{r_j} (r_j - (t - 1)) \leq \frac{n - 1}{m} \cdot \frac{r_j - (t - 1)}{r_j - 1}.$$

Note that for each link $j \in [m]$, the fraction $\frac{r_j - (t - 1)}{r_j - 1}$ is monotonically increasing in r_j (since $t \geq 2$); since $r_j \leq n$, it follows that $\frac{r_j - (t - 1)}{r_j - 1} \leq \frac{n - (t - 1)}{n - 1}$. Hence, it follows that

$$\begin{aligned} \frac{\Lambda_j(\mathbf{P})}{r_j} (r_j - (t - 1)) &\leq \frac{n - 1}{m} \cdot \frac{n - (t - 1)}{n - 1} \\ &= \alpha \cdot \frac{n - (t - 1)}{n}, \end{aligned}$$

as needed. ■

We are now ready to prove that $\Delta(t) \geq 0$. Clearly,

$$\begin{aligned} &\sum_{j \in [m]} \left(\frac{\Lambda_j(\mathbf{P})}{r_j} \right)^t \cdot (r_j)^t \\ &= \sum_{j \in [m]} \frac{\Lambda_j(\mathbf{P})}{r_j} (r_j - (t - 1)) \left(\frac{\Lambda_j(\mathbf{P})}{r_j} \right)^{t-1} \cdot (r_j)^{t-1} \\ &\leq \sum_{j \in [m]} \alpha \cdot \frac{n - (t - 1)}{n} \cdot \left(\frac{\Lambda_j(\mathbf{P})}{r_j} \right)^{t-1} \cdot (r_j)^{t-1} \quad (\text{by Lemma 4.3}) \\ &= \alpha \cdot \frac{n - (t - 1)}{n} \sum_{j \in [m]} \left(\frac{\Lambda_j(\mathbf{P})}{r_j} \right)^{t-1} \cdot (r_j)^{t-1} \\ &\leq \alpha \cdot \frac{n - (t - 1)}{n} \cdot m \alpha^{t-1} \cdot \frac{n^{t-1}}{n^{t-1}} \quad (\text{by induction hypothesis}) \\ &= m \alpha^t \cdot \frac{n^t}{n^t}. \end{aligned}$$

This implies that $\Delta(t) \geq 0$, as needed. ■

Lemma 4.2 implies now the claim. ■

4.2 The Monomial and Polynomial Prices of Anarchy

We prove:

Theorem 4.4 *Consider the case of identical users. Then, $\text{MPoA} \leq B_d$.*

Proof: Fix any instance $\langle n, m \rangle$ with an associated fully mixed Nash equilibrium \mathbf{F} . By Corollary 3.2,

$$\begin{aligned} \text{PSC}_{\lambda^d}(n, m, \mathbf{F}) &= m \cdot \sum_{t \in [d]} \left(\frac{1}{m} \right)^t \cdot S(d, t) \cdot n^t \\ &\leq m \cdot \sum_{t \in [d]} \left(\frac{1}{m} \right)^t \cdot S(d, t) \cdot n^t. \end{aligned}$$

We now proceed by case analysis.

1. Assume first that $n \geq m$. Recall that in this case, $\text{POPT}_{\lambda^d}(\mathbf{w}, m) \geq m \cdot \left(\frac{n}{m}\right)^d$. Hence,

$$\begin{aligned}
& \frac{\text{PSC}_{\lambda^d}(n, m, \mathbf{F})}{\text{POPT}_{\lambda^d}(n, m)} \\
& \leq \frac{1}{m} \cdot \left(\frac{m}{n}\right)^d \cdot m \cdot \sum_{t \in [d]} \left(\frac{1}{m}\right)^t \cdot S(d, t) \cdot n^t \\
& = \sum_{t \in [d]} \left(\frac{m}{n}\right)^{d-t} \cdot S(d, t) \\
& \leq \sum_{t \in [d]} S(d, t) \quad (\text{since } m \leq n) \\
& = B_d.
\end{aligned}$$

2. Assume now that $n < m$. Recall that, in this case, $\text{POPT}_{\lambda^d}(n, m) = n$. Hence,

$$\begin{aligned}
& \frac{\text{PSC}_{\lambda^d}(n, n, \mathbf{F})}{\text{POPT}_{\lambda^d}(n, m)} \\
& \leq \frac{1}{n} m \cdot \sum_{t \in [d]} \left(\frac{1}{m}\right)^t \cdot S(d, t) \cdot n^t \\
& = \sum_{t \in [d]} \left(\frac{n}{m}\right)^{t-1} \cdot S(d, t) \\
& < \sum_{t \in [d]} S(d, t) \quad (\text{since } n < m) \\
& = B_d.
\end{aligned}$$

So, in all cases, $\frac{\text{PSC}_{\lambda^d}(n, m, \mathbf{F})}{\text{POPT}_{\lambda^d}(n, m)} \leq B_d$. Theorem 4.1 implies now the claim. ■

By Lemma 3.1, Theorem 4.4 immediately implies:

Corollary 4.5 *Consider the case of identical users. Then, $\text{PPoA} \leq \sum_{2 \leq t \leq d} B_t$.*

We next consider the special case of (identical users and) two links. We prove:

Theorem 4.6 *Consider the case of identical users and two links. Then,*

$$\text{MPoA} \leq 2^{d-2} \left(1 + \left(\frac{1}{n}\right)^{d-1} \right).$$

This bound is tight for $n = 2$.

Proof: We start with the upper bound. Fix any instance $\langle n, 2 \rangle$ with an associated Nash equilibrium \mathbf{P} . Clearly,

$$\begin{aligned}
& \text{PSC}_{\lambda^d}(n, 2, \mathbf{P}) \\
& \leq \text{PSC}_{\lambda^d}(n, 2, \mathbf{F}) && \text{(by Theorem 4.1)} \\
& = 2 \cdot \sum_{t \in [d]} \left(\frac{1}{2}\right)^t \cdot S(d, t) \cdot n^t && \text{(by Corollary 3.2)} \\
& \leq 2 \cdot \left(\frac{1}{2} \cdot S(d, 1) \cdot n + \frac{1}{4} \cdot \sum_{2 \leq t \leq d} S(d, t) n^t\right) && \text{(since } \left(\frac{1}{2}\right)^t \leq \frac{1}{4} \text{ for } t \geq 2) \\
& = 2 \cdot \left(\frac{1}{2} \cdot S(d, 1) \cdot n + \frac{1}{4} \cdot (n^d - S(d, 1) \cdot n)\right) && \text{(since } n^d = \sum_{t \in [d]} S(d, t) n^t) \\
& = 2 \cdot \left(\frac{n}{4} + \frac{n^d}{4}\right) && \text{(since } S(d, 1) = 1).
\end{aligned}$$

On the other hand, $\text{POPT}_{\lambda^d}(n, 2) \geq 2 \cdot \left(\frac{n}{2}\right)^d$. It follows that

$$\begin{aligned}
\frac{\text{PSC}_{\lambda^d}(\mathbf{w}, 2, \mathbf{F})}{\text{POPT}_{\lambda^d}(\mathbf{w}, 2)} & \leq \left(\frac{2}{n}\right)^d \cdot \left(\frac{n}{4} + \frac{n^d}{4}\right) \\
& = 2^{d-2} \left(1 + \left(\frac{1}{n}\right)^{d-1}\right),
\end{aligned}$$

as needed.

To prove that the upper bound is tight for $n = 2$, note that for $n = 2$ it becomes $2^{d-2} + \frac{1}{2}$. We continue to prove that this is also a lower bound for $n = 2$. Fix an instance $\langle 2, 2 \rangle$. Then, $\text{POPT}_{\lambda^d}(n, m) = 2$, while

$$\begin{aligned}
& \text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{F}) \\
& = 2 \cdot \sum_{t \in [d]} \left(\frac{1}{2}\right)^t \cdot S(d, t) \cdot 2^t && \text{(by Corollary 3.2)} \\
& = 2 \cdot \left(\frac{1}{2} \cdot S(d, 1) \cdot 2 + \frac{1}{4} \cdot S(d, 2) \cdot 2 \cdot 1\right) && \text{(since } 2^t = 0 \text{ for } t \geq 3) \\
& = 2 \cdot (S(d, 1) + \frac{1}{2} \cdot S(d, 2)) \\
& = 2 \cdot \left(1 + \frac{1}{2} \cdot (2^{d-1} - 1)\right) && \text{(since } S(d, 1) = 1 \text{ and } S(d, 2) = 2^{d-1} - 1) \\
& = 2 \cdot \left(2^{d-2} + \frac{1}{2}\right).
\end{aligned}$$

It follows that $\text{MPoA} \geq 2^{d-2} + \frac{1}{2}$, which establishes the claimed tightness. ■

By Lemma 3.1 and Theorem 4.6, we obtain:

Corollary 4.7 *Consider the case of identical users and two links. Then,*

$$\text{PPoA} \leq 2^{d-1} - 1 + \frac{d-1}{n}.$$

Proof: Clearly,

$$\begin{aligned}
& \text{PPoA} \\
& \leq \sum_{2 \leq t \leq d} 2^{t-2} \left(1 + \left(\frac{1}{n}\right)^{t-1}\right) \quad (\text{by Lemma 3.1 and Theorem 4.6}) \\
& = \sum_{0 \leq t \leq d-2} 2^t + \frac{1}{n} \sum_{0 \leq t \leq d-2} \left(\frac{2}{n}\right)^t \\
& \leq 2^{d-1} - 1 + \frac{d-1}{n} \quad (\text{since } \frac{2}{n} \leq 1 \text{ for } n \geq 2),
\end{aligned}$$

as needed. ■

5 Pure Nash Equilibria

We first recall some technical definition from [20]. For a given instance $\langle \mathbf{w}, m \rangle$, call a user $i \in [n]$ *bursty* [20, Section 3] if $w_i > \frac{W}{m}$. Intuitively, the traffic of a bursty user exceeds the fair share of traffic for a link. Say that an instance $\langle \mathbf{w}, m \rangle$ is *bursty* if some user $i \in [n]$ is bursty; else, $\langle \mathbf{w}, m \rangle$ is *non-bursty*. We first prove a simple property of bursty users.

Lemma 5.1 *A bursty user is solo in an optimal assignment.*

Proof: Fix an instance $\langle \mathbf{w}, m \rangle$ with bursty user $i \in [n]$; so, $w_i > \frac{W}{m}$. Consider an optimal assignment $\mathbf{Q} = \langle q_1, \dots, q_n \rangle$. Note that $\lambda_{q_i}(\mathbf{Q}) \geq w_i$. Since i is bursty, it follows that $\lambda_{q_i}(\mathbf{Q}) > \frac{W}{m}$. Since $\sum_{j \in [m]} \lambda_j(\mathbf{Q}) = W$, there is some other link $j \in [m]$ with $j \neq q_i$ such that $\lambda_j(\mathbf{Q}) < \frac{W}{m}$. Thus, $\lambda_j(\mathbf{Q}) < w_i$. Assume now, by way of contradiction, that some user $k \neq i$ is assigned to link q_i . Modify \mathbf{Q} to obtain \mathbf{Q}' by switching user k to link j . Then,

$$\begin{aligned}
& \text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{Q}') - \text{PSC}_{\pi_d(\lambda)}(\mathbf{w}, m, \mathbf{Q}) \\
& = \sum_{t \in [d]} a_t \left((\lambda_{q_i}(\mathbf{Q}'))^t + (\lambda_j(\mathbf{Q}'))^t - (\lambda_{q_i}(\mathbf{Q}))^t - (\lambda_j(\mathbf{Q}))^t \right) \\
& = \sum_{t \in [d]} a_t \left((\lambda_{q_i}(\mathbf{Q}'))^t + (\lambda_j(\mathbf{Q}) + w_k)^t - (\lambda_{q_i}(\mathbf{Q}') + w_k)^t - (\lambda_j(\mathbf{Q}))^t \right) \\
& < \sum_{t \in [d]} a_t \left((\lambda_{q_i}(\mathbf{Q}') + w_k)^t + (\lambda_j(\mathbf{Q}) + w_k - w_k)^t - (\lambda_{q_i}(\mathbf{Q}') + w_k)^t - (\lambda_j(\mathbf{Q}))^t \right) \quad (\text{by Lemma 2.1}) \\
& = 0.
\end{aligned}$$

Since \mathbf{Q} is optimum, $\text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q}') \geq \text{PSC}_{\lambda^t}(\mathbf{w}, m, \mathbf{Q})$. A contradiction. ■

The following two simple properties of bursty users in pure Nash equilibria were shown in [20, Section 3]; they carry over to our model since their sets of Nash equilibria coincide.

Lemma 5.2 *A bursty user is solo in a pure Nash equilibrium.*

Lemma 5.3 *Consider a pure Nash equilibrium \mathbf{P} for a non-bursty instance $\langle \mathbf{w}, m \rangle$. Then, for each link $j \in [m]$, $\lambda_j(\mathbf{P}) \leq 2 \min_{\ell \in [m]} \lambda_\ell(\mathbf{L})$.*

We are now ready to prove:

Theorem 5.4 *For pure Nash equilibria.*

$$\text{MPoA} = \frac{(2^d - 1)^d}{(d - 1)(2^d - 2)^{d-1}} \left(\frac{d - 1}{d} \right)^d.$$

In our proof, we will make use of the following notations. Consider an instance $\langle \mathbf{w}, m \rangle$ with an associated pure assignment \mathbf{L} . Fix a set of links \mathcal{L} , inducing a set of users \mathcal{U} that are assigned by \mathbf{L} to links in \mathcal{L} . Then, $\mathbf{w} \setminus \mathcal{U}$ and $m \setminus \mathcal{L}$ denote the traffic vector and the number of links resulting from \mathbf{w} and m , respectively, by respective eliminations of all entries corresponding to users in \mathcal{U} and links in \mathcal{L} . Also, $\mathbf{L} \setminus (\mathcal{U}, \mathcal{L})$ denotes the assignment induced by these eliminations. We are now ready for the proof.

Proof: We first prove the upper bound. Consider any arbitrary instance $\langle \mathbf{w}, m \rangle$ with associated pure Nash equilibrium $\mathbf{L} = \langle \ell_1, \dots, \ell_n \rangle$ and optimal assignment $\mathbf{Q} = \langle q_1, \dots, q_n \rangle$. Denote $\lambda(\mathbf{L}) = \min_{\ell \in [m]} \lambda_\ell(\mathbf{L})$. We distinguish between two cases:

1. The instance $\langle \mathbf{w}, m \rangle$ is non-bursty:

Recall that in this case, by Lemma 5.3, for each link $j \in [m]$, $\lambda_j(\mathbf{L}) \leq 2\lambda(\mathbf{L})$. So, transform the set of loads $\{\lambda_\ell(\mathbf{L}) \mid \ell \in [m]\}$ into a new set of loads $\{\widehat{\lambda}_\ell \mid \ell \in [m]\}$ as the output of the following repetitive procedure:

for each link $\ell \in [m]$ **do**
 $\widehat{\lambda}_\ell \leftarrow \lambda_\ell(\mathbf{L});$
while there are distinct links $j_1, j_2 \in [m]$ with $\lambda(\mathbf{L}) < \widehat{\lambda}_{j_1} \leq \widehat{\lambda}_{j_2} < 2\lambda(\mathbf{L})$ **do**

$$\begin{pmatrix} \widehat{\lambda}_{j_1} \\ \widehat{\lambda}_{j_2} \end{pmatrix} \leftarrow \begin{pmatrix} \widehat{\lambda}_{j_1} - \min\{\widehat{\lambda}_{j_1} - \lambda(\mathbf{L}), 2\lambda(\mathbf{L}) - \widehat{\lambda}_{j_2}\} \\ \widehat{\lambda}_{j_2} + \min\{\widehat{\lambda}_{j_1} - \lambda(\mathbf{L}), 2\lambda(\mathbf{L}) - \widehat{\lambda}_{j_2}\} \end{pmatrix}.$$

Intuitively, our transformation procedure chooses at each step two intermediate latencies $\widehat{\lambda}_{j_1}$ and $\widehat{\lambda}_{j_2}$ (that is, two latencies that are not yet pushed either to the upper or to the lower end of the interval of link loads). It transfers the (strictly) positive quantity $\min\{\widehat{\lambda}_{j_1} - \lambda(\mathbf{L}), 2\lambda(\mathbf{L}) - \widehat{\lambda}_{j_2}\}$ from the small latency $\widehat{\lambda}_{j_1}$ to the large latency $\widehat{\lambda}_{j_2}$. Clearly, each step of the procedure either pushes the small latency $\widehat{\lambda}_{j_1}$ to the lower end $\lambda(\mathbf{L})$ of the interval of link latencies, or pushes the large load $\widehat{\lambda}_{j_2}$ to the upper end $2\lambda(\mathbf{L})$ of the interval of link latencies (or both). So, clearly, when the procedure terminates, there is at most one intermediate latency. Hence, by reordering links, we obtain that for some

integer $k \in [m-1] \cup \{0\}$, for each link $j \in [m]$,

$$\widehat{\lambda}_j = \begin{cases} 2\lambda(\mathbf{L}) & \text{if } j \in [k] \\ (1+x)\lambda(\mathbf{L}) & \text{if } j = k+1 \\ \lambda(\mathbf{L}) & \text{if } j \in [m] \setminus [k+1], \end{cases}$$

where $0 \leq x < 1$. Intuitively, k is the number of overloaded links. Note that $\sum_{j \in [m]} \widehat{\lambda}_j = k \cdot 2\lambda(\mathbf{L}) + (1+x) \cdot \lambda(\mathbf{L}) + (m - (k+1)) \cdot \lambda(\mathbf{L}) = (m+k+x) \cdot \lambda(\mathbf{L})$.

Note that this transformation procedure maps a set of latencies to a new set of latencies, without explicitly mapping an instance to a new instance. However, for the sake of our analysis, we will also consider that the procedure maps an instance $\langle \mathbf{w}, m \rangle$ and a Nash equilibrium \mathbf{L} to a new instance $\langle \widehat{\mathbf{w}}, m \rangle$ and a new Nash equilibrium $\widehat{\mathbf{L}}$. Note also that this transformation preserves (at each of its steps) the sum of latencies. Hence, it also preserves the total latencies, so that $W = \widehat{W}$. Clearly, for each link $j \in [m]$, $\lambda_j(\widehat{\mathbf{L}}) = \widehat{\lambda}_j$. Hence, it follows that $\sum_{j \in [m]} \lambda_j(\widehat{\mathbf{L}}) = (m+k+x) \cdot \lambda(\mathbf{L})$.

For any individual step of our repetitive procedure, Lemma 2.1 implies that

$$\begin{aligned} & \text{PSC}_{\lambda^d}(\widehat{\mathbf{w}}, m, \widehat{\mathbf{L}}) - \text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{L}) \\ &= \left(\left(\widehat{\lambda}_{j_1} - \min \left\{ \widehat{\lambda}_{j_1} - \lambda(\mathbf{L}), 2\lambda(\mathbf{L}) - \widehat{\lambda}_{j_2} \right\} \right)^d + \left(\widehat{\lambda}_{j_2} + \min \left\{ \widehat{\lambda}_{j_1} - \lambda(\mathbf{L}), 2\lambda(\mathbf{L}) - \widehat{\lambda}_{j_2} \right\} \right)^d \right) \\ & \quad - \left(\left(\lambda_{j_1} \right)^d - \left(\lambda_{j_2} \right)^d \right) \\ &> 0. \end{aligned}$$

Hence, it follows that,

$$\begin{aligned} \text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{L}) &\leq \text{PSC}_{\lambda^d}(\widehat{\mathbf{w}}, m, \widehat{\mathbf{L}}) \\ &= \sum_{j \in [m]} \left(\lambda_j(\widehat{\mathbf{L}}) \right)^d \\ &= k(2\lambda(\mathbf{L}))^d + ((1+x)\lambda(\mathbf{L}))^d + (m-k-1)\lambda(\mathbf{L})^d \\ &= \left(m + (2^d - 1)k - 1 + (1+x)^d \right) \lambda(\mathbf{L})^d. \end{aligned}$$

On the other hand,

$$\begin{aligned} \text{POPT}_{\lambda^d}(\mathbf{w}, m) &\geq m \left(\frac{W}{m} \right)^d \\ &= \frac{\widehat{W}^d}{m^{d-1}} \\ &= \frac{\left(\sum_{j \in [m]} \lambda_j(\widehat{\mathbf{L}}) \right)^d}{m^{d-1}} \\ &= \frac{(m+k+x)^d \lambda(\mathbf{L})^d}{m^{d-1}}. \end{aligned}$$

It follows that

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

Define the real function

$$f(k) = \frac{(m + (2^d - 1)k - 1 + (1 + x)^d)m^{d-1}}{(m + k + x)^d}$$

of a real variable k . (The quantity x is taken as a parameter, while m is a fixed constant). Clearly, $\text{MPoA} \leq \sup_k f(k)$. So, we will determine $\sup_k f(k)$.

To gain some intuition about the function $f(k)$ and its supremum, observe that the value of x is not really important for the upper bound; so, consider that $x = 0$. Setting then $y = \frac{k}{m}$, so that $0 < y \leq 1$, the resulting function $g(y)$ may be written as $g(y) = \frac{1 + (2^d - 1)y}{(y + 1)^d}$. It may be easily verified that $g(y)$ has a unique maximum y_0 in $(0, 1]$, where $y_0 = \frac{2^d - 1 - d}{(2^d - 1)(d - 1)}$. The presence of x in $g(k)$ will result in a bit more complicated calculations. We now continue with the details of the formal proof.

To maximize the function $f(k)$, observe that the first and second derivatives of $f(k)$ are

$$f'(k) = \frac{(2^d - 1)m^{d-1}}{(m + k + x)^d} - \frac{(m + (2^d - 1)k - 1 + (1 + x)^d)m^{d-1}d}{(m + k + x)^{d+1}}$$

and

$$f''(k) = \frac{m^{d-1}d((2^d - 1)(d - 1)k - 2(2^d - 1)(m + x) + (m - 1 + (1 + x)^d)(d + 1))}{(m + k + x)^{d+2}},$$

respectively. The only root of $f'(k)$ is

$$k_0 = \frac{(2^d - 1)(m + x) + d(-m + 1 - (1 + x)^d)}{(2^d - 1)(d - 1)}.$$

For $k = k_0$, the second derivative evaluates to

$$f''(k_0) = \frac{m^{d-1}d(-m(2^d - 2) - (2^d - 1)x + (1 + x)^d - 1)}{(m + k_0 + x)^{d+2}}.$$

Since $-(2^d - 1)x + (1 + x)^d \leq 2^d$ holds for all $x \in [0, 1]$, it follows that $f''(k_0) < 0$. Thus, k_0 is a local maximum of the function $f(k)$. Since $f(k)$ is a continuous function with a single extreme point that is a local maximum, it follows that

$$\begin{aligned} f(k) &\leq f(k_0) \\ &= \frac{(2^d - 1)^d}{d - 1} \left(\frac{d - 1}{d}\right)^d \cdot \frac{m^{d-1}}{(m(2^d - 2) + x(2^d - 1) - (1 + x)^d + 1)^{d-1}}. \end{aligned}$$

Note that the minimum value of the function $h(x) = x(2^d - 1) - (1 + x)^d + 1$ for $x \in [0, 1]$ is $h(0) = h(1) = 0$. Thus

$$\begin{aligned} f(k) &\leq \frac{(2^d - 1)^d}{d - 1} \cdot \left(\frac{d - 1}{d}\right)^d \cdot \frac{m^{d-1}}{(m(2^d - 2))^{d-1}} \\ &= \frac{(2^d - 1)^d}{(d - 1)(2^d - 2)^{d-1}} \left(\frac{d - 1}{d}\right)^d, \end{aligned}$$

as needed.

2. The instance $\langle \mathbf{w}, m \rangle$ is bursty:

We remark that Lemmas 5.1 and 5.2 imply that the existence of bursty users cannot increase the Polynomial Price of Anarchy since their assignment in a Nash equilibrium coincides with that in an optimal assignment.

Denote \mathcal{U} the (non-empty) set of bursty users. Recall that, by Lemmas 5.1 and 5.2, \mathcal{U} induces sets of solo links $\mathcal{L}_{\mathbf{L}}$ and $\mathcal{L}_{\mathbf{Q}}$ for the Nash equilibrium \mathbf{L} and the optimal assignment \mathbf{Q} , respectively, so that $|\mathcal{L}_{\mathbf{L}}| = |\mathcal{U}|$ and $|\mathcal{L}_{\mathbf{Q}}| = |\mathcal{U}|$. Since links are identical, we assume that $\mathcal{L}_{\mathbf{L}} = \mathcal{L}_{\mathbf{Q}} = \mathcal{L}$, with $|\mathcal{L}| \geq 1$. So,

$$\begin{aligned} \text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{L}) &= \sum_{j \in \mathcal{L}} (\lambda_j(\mathbf{L}))^d + \text{PSC}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}, \mathbf{L} \setminus (\mathcal{U}, \mathcal{L})) \\ &= \sum_{i \in \mathcal{U}} w_i^d + \text{PSC}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}, \mathbf{L} \setminus (\mathcal{U}, \mathcal{L})) \end{aligned}$$

and

$$\begin{aligned} \text{POPT}_{\lambda^d}(\mathbf{w}, m) &= \text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{Q}) \\ &= \sum_{j \in \mathcal{L}} (\lambda_j(\mathbf{L}))^d + \text{PSC}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}, \mathbf{Q} \setminus (\mathcal{U}, \mathcal{L})) \\ &= \sum_{i \in \mathcal{U}} w_i^d + \text{PSC}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}, \mathbf{Q} \setminus (\mathcal{U}, \mathcal{L})). \end{aligned}$$

Note first that the assignment $\mathbf{L} \setminus (\mathcal{U}, \mathcal{L})$ is a Nash equilibrium for the instance $\langle \mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L} \rangle$. Moreover, since \mathbf{Q} is an optimal assignment for the instance $\langle \mathbf{w}, m \rangle$, it follows that $\mathbf{Q} \setminus (\mathcal{U}, \mathcal{L})$ is an optimal assignment for the instance $\langle \mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L} \rangle$, so that

$$\text{PSC}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}, \mathbf{Q} \setminus (\mathcal{U}, \mathcal{L})) = \text{POPT}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}).$$

Thus,

$$\text{POPT}_{\lambda^d}(\mathbf{w}, m) = \sum_{i \in \mathcal{U}} w_i^d + \text{POPT}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}).$$

It follows that

$$\begin{aligned} \frac{\text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{L})}{\text{POPT}_{\lambda^d}(\mathbf{w}, m)} &= \frac{\sum_{i \in \mathcal{U}} w_i^d + \text{PSC}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}, \mathbf{L} \setminus (\mathcal{U}, \mathcal{L}))}{\sum_{i \in \mathcal{U}} w_i^d + \text{POPT}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L})} \\ &\leq \frac{\text{PSC}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L}, \mathbf{L} \setminus (\mathcal{U}, \mathcal{L}))}{\text{POPT}_{\lambda^d}(\mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L})}. \end{aligned}$$

Consider the smaller instance $\langle \mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L} \rangle$ and the associated Nash equilibrium $\mathbf{L} \setminus (\mathcal{U}, \mathcal{L})$. There are two possibilities depending on whether $\langle \mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L} \rangle$ is bursty or not.

- Assume first that the instance $\langle \mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L} \rangle$ is non-bursty. Then, we are reduced to the previous case of non-bursty instances, and the upper bound follows.
- Assume now that the smaller instance $\langle \mathbf{w} \setminus \mathcal{U}, [m] \setminus \mathcal{L} \rangle$ is bursty. We repeatedly identify the set of bursty users for the smaller instance, and we reduce this smaller instance to an even smaller instance that may be bursty or non-bursty. This procedure eventually yields a non-bursty instance (even the trivial one with one user), and the claim for the original bursty instance follows inductively.

The proof of the upper bound is now complete.

We continue to prove the lower bound. Construct an instance $\langle \mathbf{w}, m \rangle$ as follows. There are $m = (2^d - 1)(d - 1)$ links. There are $2(2^d - d - 1)$ *heavy* users with traffic 1; there are $m \cdot (m - (2^d - d - 1))$ *light* users with traffic $\frac{1}{m}$. Consider now the following assignments:

- In the pure assignment \mathbf{L} , heavy users are evenly distributed to $2^d - d - 1$ links; light users are evenly distributed to the remaining $m - (2^d - d - 1)$ links. Clearly, \mathbf{L} is a Nash equilibrium with

$$\begin{aligned} \text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{L}) &= 2^d \cdot (2^d - d - 1) + 1^d \cdot ((2^d - 1)(d - 1) - (2^d - d - 1)) \\ &= (2^d - 1)(2^d - 2). \end{aligned}$$

- In the pure assignment \mathbf{Q} , each (of $2(2^d - d - 1)$) heavy user is assigned solo to each of $2(2^d - d - 1)$ links; $m(m - 2(2^d - d - 1))$ light users are evenly assigned to the remaining $m - 2(2^d - d - 1)$ links, while the remaining $m(2^d - d - 1)$ light users are evenly assigned to all m links. It is easy to see that the latency on each link induced by \mathbf{Q} is $1 + \frac{2^d - d - 1}{m} = \frac{m + 2^d - d - 1}{m}$. Thus,

$$\begin{aligned} \text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{Q}) &= m \cdot \left(\frac{m + 2^d - d - 1}{m} \right)^d \\ &= \frac{(d - 1)(2^d - 2)^d}{(2^d - 1)^{d-1}} \cdot \left(\frac{d}{d - 1} \right)^d. \end{aligned}$$

Thus,

$$\begin{aligned} \text{MPoA} &\geq \frac{\text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{L})}{\text{PSC}_{\lambda^d}(\mathbf{w}, m, \mathbf{Q})} \\ &= \frac{(2^d - 1)^d}{(d - 1)(2^d - 2)^{d-1}} \cdot \left(\frac{d - 1}{d}\right)^d, \end{aligned}$$

as needed. ■

By Lemma 3.1, Theorem 5.4 immediately implies:

Corollary 5.5 *For pure Nash equilibria.*

$$\text{PPoA} \leq \sum_{2 \leq t \leq d} \frac{(2^t - 1)^t}{(t - 1)(2^t - 2)^{t-1}} \left(\frac{t - 1}{t}\right)^t.$$

6 Epilogue

We introduced and analyzed an interesting variant of the well studied KP model [19] for selfish routing that reflects some influence from the much older Wardrop model [29]. Our analysis highlights some interesting connections to classical combinatorial numbers such as the Stirling numbers of the second kind [28] and the Bell numbers [5]. In particular, we formulated and proved the validity of the PFMNE Conjecture. In turn, this validity was instrumental for proving (sometimes tight) bounds on Monomial Price of Anarchy; these immediately implied upper bounds on Polynomial Price of Anarchy.

Several interesting problems remain open. On the most concrete level, we do not yet know any lower bounds on Polynomial Price of Anarchy. Are our upper bounds tight? We are also missing *general* bounds on Monomial and Polynomial Prices of Anarchy (ones that hold for arbitrary users, for an arbitrary number of links and for all (mixed) Nash equilibria). Proving such bounds remains a very challenging open problem.

Acknowledgements:

We would like to thank the two anonymous referees for their many helpful comments that significantly improved the presentation of this work. We especially thank one of the referees for contributing to this work a substantial improvement (together with a proof) to an earlier version of Theorem 4.1.

References

- [1] S. Aland, D. Dumrauf, M. Gairing, B. Monien and F. Schoppmann, "Exact Price of Anarchy for Polynomial Congestion Games," *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science*, B. Duran and W. Thomas eds., pp. 218–229, Vol. 3884, Lecture Notes in Computer Science, Springer-Verlag, February 2006.
- [2] E. Altman, T. Basar, T. Jiménez and N. Shimkin, "Competitive Routing in Networks with Polynomial Costs," *IEEE Transactions on Automatic Control*, Vol. 47, No. 1, pp. 92–96, January 2002.
- [3] R. J. Aumann, "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, Vol. 1, pp. 67–95, 1974.
- [4] B. Awerbuch, Y. Azar and A. Epstein, "The Price of Routing Unsplittable Flow," *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pp. 57–66, May 2005.
- [5] E. T. Bell, "Exponential Polynomials," *Annals of Mathematics*, Vol. 35, pp. 258–277, 1934.
- [6] D. E. Boyce, B. N. Janson and R. W. Eash, "The Effect on Equilibrium Trip Assignment of Different Congestion Functions," *Transport Research A*, Vol. 15, pp. 223–232, 1981.
- [7] G. Christodoulou and E. Koutsoupias, "The Price of Anarchy of Finite Congestion Games," *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pp. 67–73, May 2005.
- [8] G. Christodoulou and E. Koutsoupias, "On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games," *Proceedings of the 13th Annual European Symposium on Algorithms*, G. Brodal and S. Leonardi eds., pp. 59–70, Vol. 3669, Lecture Notes in Computer Science, Springer-Verlag, October 2005.
- [9] A. Czumaj, P. Krysta and B. Vöcking, "Selfish Traffic Allocation for Server Farms," *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pp. 287–296, May 2002.
- [10] A. Czumaj and B. Vöcking, "Tight Bounds for Worst-Case Equilibria," *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 413–420, January 2002.
- [11] R. Feldmann, M. Gairing, T. Lücking, B. Monien and M. Rode, "Nashification and the Coordination Ratio for a Selfish Routing Game," *Proceedings of the 30th International*

- Colloquium on Automata, Languages and Programming*, J. C. M. Baeten, J. K. Lenstra, J. Parrow and G. J. Woeginger eds., pp. 514–526, Vol. 2719, Lecture Notes in Computer Science, Springer-Verlag, June/July 2003.
- [12] S. Fischer and B. Vöcking, "On the Structure and Complexity of Worst-Case Equilibria," *Proceedings of the 1st International Workshop on Internet and Network Economics*, X. Deng and Y. Ye eds., pp. 151–160, Vol. 3828, Lecture Notes in Computer Science, Springer-Verlag, December 2005.
- [13] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas and P. Spirakis, "The Structure and Complexity of Nash Equilibria for a Selfish Routing Game," *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, P. Widmayer, F. Triguero, R. Morales, M. Hennessy, S. Eidenbenz and R. Conejo eds., pp. 123–134, Vol. 2380, Lecture Notes in Computer Science, Springer-Verlag, July 2002.
- [14] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and M. Rode, "Nash Equilibria in Discrete Routing Games with Convex Latency Functions," *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, J. Diaz, J. Karhumäki, A. Lepistö and D. Sannella eds., pp. 645–657, Vol. 3142, Lecture Notes in Computer Science, Springer-Verlag, July 2004.
- [15] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis, "Extreme Nash Equilibria," *Proceedings of the 8th Italian Conference on Theoretical Computer Science*, C. Blundo and C. Laneve eds., pp. 1–20, Vol. 2841, Lecture Notes in Computer Science, Springer-Verlag, October 2003.
- [16] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis, "Structure and Complexity of Extreme Nash Equilibria," *Theoretical Computer Science*, (Special Issue "Game Theory Meets Theoretical Computer Science", M. Mavronicolas and S. Abramsky guest eds.), Vol. 343, Nos. 1-2, pp. 133-157, October 2005.
- [17] R. L. Graham, D. E. Knuth and O. Patashnik, *Concrete Mathematics – A Foundation for Computer Science*, Second Edition, Addison-Wesley, 1994.
- [18] E. Koutsoupias, M. Mavronicolas and P. Spirakis, "Approximate Equilibria and Ball Fusion," *Theory of Computing Systems*, Vol. 36, No. 6, pp. 683–693, November 2003.
- [19] E. Koutsoupias and C. H. Papadimitriou, "Worst-Case Equilibria," *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science*, G. Meinel and S. Tison eds., pp. 404–413, Vol. 1563, Lecture Notes in Computer Science, Springer-Verlag, March 1999.

- [20] T. Lücking, M. Mavronicolas, B. Monien and M. Rode, "A New Model for Selfish Routing," *Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science*, V. Diekert and M. Habib eds., pp. 547–558, Vol. 2996, Lecture Notes in Computer Science, Springer-Verlag, March 2004.
- [21] T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis and I. Vrto, "Which is the Worst-Case Nash Equilibrium?," *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science*, B. Rovan and P. Vojtas eds., pp. 551–561, Vol. 2747, Lecture Notes in Computer Science, Springer-Verlag, August 2003.
- [22] M. Mavronicolas and P. Spirakis, "The Price of Selfish Routing," *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 510–519, July 2001.
Full version accepted to *Algorithmica*.
- [23] J. F. Nash, "Equilibrium Points in N -Person Games," *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 36, pp. 48–49, 1950.
- [24] J. F. Nash, "Non-Cooperative Games," *Annals of Mathematics*, Vol. 54, pp. 286–295, 1951.
- [25] C. H. Papadimitriou, "Algorithms, Games and the Internet," *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 510–519, July 2001.
- [26] R. W. Rosenthal, "A Class of Games Possessing Pure-Strategy Nash Equilibria," *International Journal of Game Theory*, Vol. 2, pp. 65–67, 1973.
- [27] T. Roughgarden and É. Tardos, "How Bad is Selfish Routing?," *Journal of the ACM*, Vol. 49, No. 2, pp. 236–259, May 2002.
- [28] J. Stirling, *Methodus Differentialis*, London, 1730.
- [29] J. G. Wardrop, "Some Theoretical Aspects of Road Traffic Research," *Proceedings of the Institute of Civil Engineers*, Pt. II, Vol. 1, pp. 325–378, 1956.