Nash Equilibria in Discrete Routing Games with Convex Latency Functions [⋆]

Martin Gairing [a,∗], Thomas Lücking [b,1], Marios Mavronicolas [c,2], Burkhard Monien [a], Manuel Rode [a,3]

[a] *Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany.*

[b] *SAP AG, Dietmar-Hopp-Allee 16, 69190 Walldorf, Germany.*

[c] *Department of Computer Science, University of Cyprus, P. O. Box 20537, Nicosia CY-1678, Cyprus.*

Abstract

In a *discrete routing game,* each of n selfish *users* employs a *mixed strategy* to ship her (*unsplittable*) *traffic* over m parallel *links*. The *(expected) latency* on a link is determined by an arbitrary *non-decreasing, non-constant* and *convex latency function* ϕ. In a *Nash equilibrium,* each user alone is minimizing her *(Expected) Individual Cost,* which is the (expected) latency on the link she chooses. To evaluate Nash equilibria, we formulate *Social Cost* as the sum of the users' (Expected) Individual Costs. The *Price of Anarchy* is the *worst-case* ratio of Social Cost for a Nash equilibrium over the least possible Social Cost. A Nash equilibrium is *pure* if each user deterministically chooses a single link; a Nash equilibrium is *fully mixed* if each user chooses each link with non-zero probability. We obtain:

For the case of *identical users,* the Social Cost of *any* Nash equilibrium is no more than the Social Cost of the fully mixed Nash equilibrium, which may exist only uniquely. Moreover, instances admitting a fully mixed Nash equilibrium enjoy an efficient characterization.

For the case of identical users, we derive two upper bounds on the Price of Anarchy: For the case of *identical links* with a *monomial* latency function $\phi(x) = x^d$, the Price of Anarchy is the *Bell number of order* $d + 1$. For pure Nash equilibria, a generic upper bound from the *Wardrop* model can be transfered to discrete routing games. For polynomial latency functions with non-negative coefficients and degree d, this yields an upper bound of $d + 1$.

For the case of identical users, a pure Nash equilibrium (and thereby an optimum pure assignment) can be computed in time $O(m \log m \log n)$. For the general case, computing the *best* or the *worst* pure Nash equilibrium is $\mathcal{NP}$-complete, even for identical links with an identity latency function.

Key words: Discrete Routing Games, Convex Latency Functions, Price of Anarchy, Fully Mixed Nash Equilibria

⋆ A preliminary version of this work appeared in the *Proceedings of the 31st International Colloquium on Automata, Languages and Programming,* pp. 645–657, Vol. 3142, Lecture Notes in Computer Science, Springer-Verlag, July 2004. This work has been partially supported by the IST Program of the European Union under contracts IST-2001-33116 (FLAGS), 001907 (DELIS) and 015964 (AEOLUS), by research funds at University of Cyprus, and by the VEGA grant No. 2/3164/23.

∗ Corresponding author at: International Computer Science Institute, 1947 Center Street, Berkeley, CA 94704, USA.

Email addresses: gairing@upb.de (Martin Gairing), luck@upb.de (Thomas Lücking), mavronic@ucy.ac.cy (Marios Mavronicolas), bm@upb.de (Burkhard Monien), rode@upb.de (Manuel Rode).

1 The work of this author was done while at the University of Paderborn.

2 Part of the work of this author was done while visiting the University of Paderborn.

3 International Graduate School of *Dynamic Intelligent Systems,* University of Paderborn, Germany.

1. Introduction

1.1. *Background*

Nash equilibrium [39,40] is one of the most significant concepts in *Non-Cooperative Game Theory.* For a given *strategic game,* a Nash equilibrium is a state where no player can improve her individual objective by unilaterally changing its *strategy.* A Nash equilibrium is called *pure* if each player chooses exactly one strategy; it is called *mixed* if each player makes her choice using a probability distribution over strategies. In a *fully mixed Nash equilibrium* [36], each player chooses each strategy with non-zero probability. The *Price of Anarchy* [31,41] is the *worst-case* ratio of the *Social Cost* in a Nash equilibrium and the least possible Social Cost.

Much of the recent algorithmic work on Non-Cooperative Game Theory considered *selfish routing,* where it focused on the KP model due to Koutsoupias and Papadimitriou [31] and the Wardrop model [48]. The KP model was proposed only recently in the context of studying the effects of selfish traffic over the Internet; in contrast, the Wardrop model dates back to the 1950s, when it was used for studying the economics of transportation networks (see, e.g., [4,5,12]).

- In the KP model, each of n selfish *users* employs a *mixed strategy,* which is a probability distribution over m parallel *links,* to ship its (*unsplittable*) *traffic*; so, the traffic of each user is shipped all together and with no splitting. The *(expected) latency* on a link is linear in 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 (expected) maximum latency; the *Optimum* is the least possible maximum latency.
- In the Wardrop model, the network can be arbitrary. Modeled as a *network flow* from *source* to *destination,* selfish traffic is infinitesimally *splittable*; this modeling rules out mixed strategies from consideration. Associated with each link is a *convex latency function,* which determines the *latency* on the link for a given traffic. In a *Wardrop equilibrium* [48], all used paths incur the same (total) latency. So, a Wardrop equilibrium can be interpreted as a Nash equilibrium for a strategic game with infinitely many users, each carrying an infinitesimal amount of traffic. The *Individual Cost* of each such user is the sum of link latencies on the path it chooses; the *Social Cost* is the integral of the Individual Costs; so, it is the overall cost incurred to the users.

1.2. *Discrete Routing Games*

In this work, we introduce the model of *discrete routing games* as a hybridization of the KP model and the Wardrop model.

We follow the KP model to consider the parallel links network with m links and n users with unsplittable traffics and mixed strategies. However, we allow *arbitrary* non-decreasing, non-constant and convex latency functions, whereas latency functions for the KP model are linear. So, the *latency function* for a link is a *convex* function of the total traffic of users choosing the link. The *Social Cost* is the sum of (Expected) Individual Costs; the *(Expected) Individual Cost* of a user is the (expected) latency on the link it chooses. So, as far as the generality of latency functions and the Social Cost are concerned, discrete routing games lean towards the Wardrop model; however, the network structure and the unsplittability of traffics come from the KP model.

The assumption of convex latency functions determines a very broad class of discrete routing games. Restricted to *monotone* latency functions and pure Nash equilibria, discrete routing games were already studied in [10]. Restricted to *linear* latency functions, they have been studied by Lücking *et al.* in [33], where Social Cost was formulated as the sum of *weighted* Expected Individual Costs and called *Quadratic Social Cost*; so, the model in [33] is the special case of discrete routing games where latency functions are linear and users are identical. To the best of our knowledge, discrete routing games represent the *first* model to *simultaneously* consider mixed Nash equilibria and arbitrary (convex) latency functions.

Discrete routing games are a particular instance of *weighted congestion games* [37,42], where each pure strategy is an *arbitrary* (not necessarily singleton) set of links. It is known that all *unweighted* congestion games admit a pure Nash equilibrium [42]. Hence, so do discrete routing games in the special case of identical (*unweighted*) users. However, it is straightforward to verify through a lexicographic argument (cf. [18, Theorem 1]) that discrete routing games admit a pure Nash equilibrium in the general case of arbitrary (*weighted*) users.

1.3. *Contribution*

Our results for discrete routing games are partitioned into four major groups:

1.3.1. *Fully Mixed Nash Equilibria*

Which is the *worst-case* Nash equilibrium for discrete routing games with respect to Social Cost? This is a very natural question, which we address for the special case of identical users. As our main result, we prove that for any discrete routing game with convex latency functions, whenever a fully mixed Nash equilibrium exists, it is a worst-case Nash equilibrium (Theorem 4.3). Therewith, we prove the *Fully Mixed Nash Equilibrium Conjecture* for discrete routing games (but only for the special case of identical users) [4]. The proof relies critically on the convexity assumption for the latency functions; we provide a simple counterexample to show that this assumption is essential (Proposition 4.4).

Furthermore, we prove that a fully mixed Nash equilibrium may exist only uniquely (Theorem 4.6). The proof utilizes the assumption that the latency functions are non-decreasing and non-constant, but it does *not* need the convexity assumption on them.

Finally, we provide a combinatorial characterization of instances admitting a fully mixed Nash equilibrium. Specifically, we identify the classes of *dead* and *special* links, and we prove some combinatorial properties for them (Lemmas 4.8 and 4.9). In turn, these properties are used for characterizing instances admitting a fully mixed Nash equilibrium (Theorem 4.10). Furthermore, we prove a generalization of the Fully Mixed Nash Equilibrium Conjecture for instances that do *not* admit a fully mixed Nash equilibrium (Theorem 4.11).

As our chief combinatorial instrument for the study of fully mixed Nash equilibria in discrete routing games, we introduce and study a novel combinatorial function, called the *binomial function* (Section 2).

1.3.2. *Price of Anarchy*

We focus on the special case of identical users, for which we present two upper bounds on the Price of Anarchy for mixed and pure Nash equilibria, respectively. We remind the reader that n and m are the numbers of users and links, respectively.

- We first treat mixed Nash equilibria, where we consider the special case of identical links with a *monomial* (convex) latency function $\phi(\lambda) = \lambda^d$. We prove that the Price of Anarchy is less than the *Bell number of order* $d+1$ (Theorem 5.1). When $m = n$ and in the limit, this bound can be attained arbitrarily close but *not* exactly.
- For pure Nash equilibria, we consider the case of arbitrary links. We revisit a generic upper bound on the Price of Anarchy for the Wardrop model, which was shown by Roughgarden and Tardos [44]. We utilize the assumption that latency functions are non-decreasing and non-constant to transfer this bound to discrete routing games (Proposition 5.3).

  For polynomial latency functions with non-negative coefficients and degree $d \geq 1$, the transfered bound immediately yields an upper bound of $d+1$ (Corollary 5.4).

Interestingly, both shown upper bounds on the Price of Anarchy are *constant* — independent of m and n.

[4] The validity of the Fully Mixed Nash Equilibrium Conjecture for the general case of arbitrary users in discrete routing games is left as an open problem. Apparently, the couterexample for the case of arbitrary users in the KP-model shown in [17] does not apply to discrete routing games where Social Cost is defined in a different way.

1.3.3. *Algorithmic Results*

We present algorithmic results for the case of identical users and arbitrary links. These results utilize the assumption that latency functions are non-decreasing.

– We show that a pure Nash equilibrium can be computed in $O(m \log m \log n)$ time (Theorem 6.1). This improves on the simple approach of assigning users one-by-one to their respective best links which is motivated by Graham's classical LPT scheduling algorithm [27] (cf. [18, Section 3]); the approach applies directly to the general case of arbitrary users, yielding a polynomial time algorithm. The improvement is achieved by an algorithm running in $\log n$ *phases*; in each phase, user chunks of halving size are switched together to a different link in order to improve their (common) Individual Cost. We show that the number of such switches per phase is $O(m)$, We use a suitable data structure for implementing each switch in $\Theta(\log m)$ time, for a total of $O(m \log m)$ time per phase; this implies that the total time is $O(m \log m \log n)$.
– Under a certain convexity assumption on the latency functions, we exhibit a very simple, polynomial time reduction from the problem of computing an optimum (pure) assignment to the problem of computing a pure Nash equilibrium (Proposition 6.4). Together with our efficient algorithm for computing a pure Nash equilibrium, this implies a corresponding efficient algorithm to compute an optimum assignment (Corollary 6.5).

1.3.4. *Complexity Results*

We present some complexity results for the problems of computing the *best-case* and *worst-case* Nash equilibrium (with respect to the particular Social Cost adopted for discrete routing games). Specifically, we prove that in the general case of arbitrary users, both problems are $\mathcal{NP}$-complete (Theorems 7.1 and 7.3, respectively). Both $\mathcal{NP}$-completeness results hold even for the case of identical links with an identity latency function.

Both proofs use polynomial time transformations from the $\mathcal{NP}$-complete PARTITION problem [29], whose counting version is known to be $\#\mathcal{P}$-complete [45]. The employed transformations are *parsimonious* — roughly speaking, they preserve the number of solutions (cf. [32, Definition 26.6]). This implies that the problems of counting the best-case and worst-case pure Nash equilibria are both $\#\mathcal{P}$-complete as well (Corollaries 7.2 and 7.4).

1.4. *Related Work and Comparison*

The KP model has received a lot of interest and attention – see, e.g., [10,11,13–15,18,21,24,30,34,36]. For a survey of early work on the Wardrop model, see [5]. Inspired by the new interest in the Price of Anarchy, Roughgarden and Tardos [44] initiated recently a reinvestigation of the Wardrop model; for recent results, we refer the reader to the book [43] (and the references therein).

The fully mixed Nash equilibrium was originally introduced and analyzed by Mavronicolas and Spirakis [36] for the KP model. For the KP model, it was shown that existence of a fully mixed Nash equilibrium implies its uniqueness [36]. This result applies to the special case of discrete routing games where latency functions are linear; hence, it is broadened by Theorem 4.6.

The original *Fully Mixed Nash Equilibrium Conjecture* for the KP model states that the worst-case Nash equilibrium is the fully mixed Nash equilibrium for instances where the fully mixed Nash equilibrium exists. This conjecture was originally motivated by some results in [18]; it was explicitly formulated in [24] and further studied and extended to other related models in [16,17,22,26,33–35]. In particular, Lücking *et al.* [33] proved the Fully Mixed Nash Equilibrium Conjecture for the special case of identical users and identical links in their model (which is itself a special case of discrete routing games). Recently, Fischer and Vöcking [17] provided a counterexample to the Fully Mixed Nash Equilibrium Conjecture for the special case of arbitrary users and identical links in the KP model.

Bounds on the Price of Anarchy for the KP model are given in [11,14,30,31,36]. These include (*tight*) bounds of $\Theta\left(\frac{\log m}{\log \log m}\right)$ for the case of identical links [11,30,31,36] and of $\Theta\left(\frac{\log m}{\log \log \log m}\right)$ for the case of arbitrary links [11]. Bounds on the Price of Anarchy for several variants and generalizations of the KP model were proved in [2,3,7,19–21,23,35,47].

Bounds on the Price of Anarchy for congestion games were proved in [1,2,8]. Christodoulou and Koutsoupias [8] consider congestion games with unweighted players and under linear and polynomial latency functions (of degree d and with non-negative coefficients); they define Social Cost as either the maximum (Expected) Individual Cost or the sum of (Expected) Individual Costs. The upper bounds obtained for their latter definition apply to discrete routing games as well. In particular, Christodoulou and Koutsoupias [8] prove that the Price of Anarchy for Social Cost as sum of (Expected) Individual Costs is $\Theta(d^{d(1-o(1))})$. (Exact values were later obtained in [1] for both cases of unweighted and weighted players.) Corollary 5.4 provides a much smaller upper bound of $d+1$ for the special case of discrete routing games where each pure strategy is a singleton set (but restricted to pure Nash equilibria).

For the KP model, Fotakis *et al.* [18] showed that a pure Nash equilibrium can be computed in polynomial time using Graham's LPT scheduling algorithm [27]. This result applies to the special case of discrete routing games where latency functions are linear; hence, it is broadened by Theorem 6.1.

Fotakis *et al.* [18] showed that computing the best-case and the worst-case (pure) Nash equilibria are both $\mathcal{NP}$-complete for the KP model where Social Cost is defined as the (expected) maximum latency; in contrast, Theorems 7.1 and 7.3 apply to discrete routing games where Social Cost is defined as the sum of (Expected) Individual Costs.

A result of Conitzer and Sandholm [9, Theorem 1] directly implies that it is $\mathcal{NP}$-complete to decide if an *arbitrary* strategic game has a (mixed) Nash equilibrium for which the sum of the (Expected) Individual Costs is no more than some *arbitrary* number; this holds even if the game is *symmetric* and has only two players. Theorem 7.1 is comparable to this result. On one hand, it restricts to a particular kind of strategic games (namely, discrete routing games) and it applies even to the case with only three links; on the other hand, it assumes an arbitrary number of players. Furthermore, Corollary 7.2 is comparable to [9, Corollary 12], which established the $\#\mathcal{P}$-completeness of the corresponding counting problem.

Subsequent to this work, Gairing *et al.* [22] introduced yet another hybridization of the KP model and the Wardrop model, where Social Cost is the expectation of the sum (over links) of *Latency Costs*; each Latency Cost is defined as a polynomial function of the total traffic of users choosing a link. For their model, Gairing *et al.* [22] established the Fully Mixed Nash Equilibrium Conjecture for the special case of identical users and identical links; they also proved several upper bounds on the Price of Anarchy. To do so, Gairing *et al.* [22] showed further interesting properties of the binomial function. (In fact, we used one such property, restated here as Lemma 2.3, to simplify some proofs in the preliminary version of this work.)

1.5. *Notation and Preliminaries*

Throughout, denote for any integer $m \geq 1$, $[m] = \{1, \ldots, m\}$. Denote as $\mathbf{0}$ and $\mathbf{1}$ the vectors (of any appropriate dimension) with all zeros and all ones, respectively. Denote as $\mathbb{R}_0^+$ and $\mathbb{N}_0$ the sets of non-negative real and natural numbers, respectively. For a function $\phi : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$, denote as $\widehat{\phi}$ the function defined by $\widehat{\phi}(\lambda) = \phi(\lambda + 1)$. For a random variable X with associated probability distribution $\mathbf{P}$, denote as $\mathbb{E}_{\mathbf{P}}(X)$ the *expectation* of X (according to $\mathbf{P}$). For an integer $m \geq 2$ and a *dimension* $j \in [m]$, the j*-characteristic* m-dimensional vector $\boldsymbol{\chi}_j$ has entry j equal to 1 and all other entries equal to 0.

Some of our analysis will bring into play some special numbers from classical combinatorics. Recall first the *Bell number of order* d [6], denoted as B_d, which counts the number of partitions of a set with d elements into *blocks* (non-empty subsets). It is known that $B_d = \sum_{k \in [d]} S(d, k)$, where for each $k \in [d]$, $S(d, k)$, the *Stirling number of the second kind* [46], counts the number of partitions of a set with d elements into exactly k blocks. For any pair of integers $r \geq 1$ and $k \geq 1$, the *falling factorial of* r *of order* k, denoted as $r^{\underline{k}}$, is given by $r^{\underline{k}} = r(r-1) \cdot \ldots \cdot (r-(k-1))$, when $r \geq k$. Otherwise ($k \geq r+1$), $r^{\underline{k}} = 0$.

1.6. *Organization*

In Section 2, we introduce the binomial function. Section 3 presents discrete routing games. The fully mixed Nash equilibrium is studied in Section 4. Section 5 contains the bounds on the Price of Anarchy. Our algorithms for computing pure Nash equilibria and optimum pure assignments appear in Section 6. Section 7 contains the complexity results for best-case and worst-case pure Nash equilibria. We conclude, in Section 8, with a discussion of our results and some open problems.

2. The Binomial Function

For a vector of probabilities $\mathbf{p} \in [0,1]^r$, denote as $\widetilde{\mathbf{p}}$ the vector of probabilities with all entries equal to $\frac{\sum_{i\in[r]} p_i}{r}$. We now define a combinatorial function:

Definition 2.1 (Binomial Function) *For any integer r, consider a triple of a vector of probabilities $\mathbf{p} = \langle p_1, \ldots, p_r\rangle$, a vector $\mathbf{w} = \langle w_1, \ldots, w_r\rangle \in (\mathbb{R}_0^+)^r$, and a function $\phi : \mathbb{R}_0^+ \to \mathbb{R}_0^+$. For a subset $\mathcal{U} \subseteq [r]$, denote $w_{\mathcal{U}} = \sum_{k\in\mathcal{U}} w_k$. The* binomial function $\mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi)$ *is defined by*

$$\mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi) = \sum_{\mathcal{U}\subseteq[r]} \prod_{k\in\mathcal{U}} p_k \prod_{k\notin\mathcal{U}} (1-p_k) \cdot \phi(w_{\mathcal{U}}) \,.$$

Roughly speaking, the binomial function BF represents the expectation of a function ϕ of a random variable that follows some kind of a binomial distribution – hence, its name. Clearly, the binomial function is a symmetric function in the probabilities $p_1, \ldots, p_r$ and the weights $w_1, \ldots, w_r$ – for each permutation π on $[r]$ that maps $\mathbf{p}$ to $\pi(\mathbf{p})$ and $\mathbf{w}$ to $\pi(\mathbf{w})$, $\mathsf{BF}(\pi(\mathbf{p}), \pi(\mathbf{w}), \phi) = \mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi)$. Moreover, the binomial function BF is a continuous function in the probabilities $p_1, \ldots, p_r$.

Say that the function ϕ is *non-constant on the vector* $\mathbf{w}$ if $\phi\left(\min_{i\in[r]} w_i\right) \neq \phi\left(\sum_{i\in[r]} w_i\right)$. If ϕ is both non-decreasing and non-constant on $\mathbf{w}$, then $\phi\left(\min_{i\in[r]} w_i\right) < \phi\left(\sum_{i\in[r]} w_i\right)$. We now prove a significant monotonicity property of the binomial function:

Lemma 2.1 (Monotonicity of Binomial Function) *Assume that ϕ is a non-decreasing and non-constant function on the vector* $\mathbf{w}$*. Then, $\mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi)$ is strictly increasing in each probability p_i, where $i \in [r]$.*

PROOF. Since BF is non-decreasing and non-constant, there is some index $r' \in [r]$ and some set $\mathcal{U} \subseteq [r] \setminus \{r'\}$ such that $\phi(w_{r'} + w_{\mathcal{U}}) > \phi(w_{\mathcal{U}})$. Assume, without loss of generality, that $r' = r$ so that $\mathcal{U} \subseteq [r-1]$. Since $\mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi)$ is a symmetric function in the probabilities p_i, where $i \in [r]$, it suffices to prove that $\mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi)$ is strictly increasing in the probability p_r. Clearly,

$$\begin{aligned}
\mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi) &= \sum_{\mathcal{U}\subseteq[r]} \prod_{k\in\mathcal{U}} p_k \prod_{k\notin\mathcal{U}} (1-p_k)\, \phi(w_{\mathcal{U}}) \\
&= \sum_{\mathcal{U}\subseteq[r-1]} \prod_{k\in\mathcal{U}} p_k \prod_{k\notin\mathcal{U}\cup\{r\}} (1-p_k) \cdot [p_r \cdot \phi(w_{\mathcal{U}} + w_r) + (1-p_r)\cdot \phi(w_{\mathcal{U}})] \\
&= \sum_{\mathcal{U}\subseteq[r-1]} \prod_{k\in\mathcal{U}} p_k \prod_{k\notin\mathcal{U}\cup\{r\}} (1-p_k) \cdot (\phi(w_{\mathcal{U}}) + p_r \cdot [\phi(w_{\mathcal{U}} + w_r) - \phi(w_{\mathcal{U}})]) \,.
\end{aligned}$$

Since ϕ is non-decreasing, $\phi(w_{\mathcal{U}} + w_r) - \phi(w_{\mathcal{U}}) \geq 0$ for all $\mathcal{U} \subseteq [r-1]$. Moreover, by assumption, $\phi(w_{\mathcal{U}} + w_r) - \phi(w_{\mathcal{U}}) > 0$ for some $\mathcal{U} \subseteq [r-1]$. It follows that $\mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi)$ is strictly increasing in p_r, as needed. ■

A significant special case occurs when $\mathbf{w}$ is a *constant* vector with all entries equal to 1. In this case, we abuse notation to write $\mathsf{BF}(\mathbf{p}, r, \phi)$ for $\mathsf{BF}(\mathbf{p}, \mathbf{w}, \phi)$. For this special case, we prove a monotonicity property of the binomial function with respect to averaging the probabilities:

Lemma 2.2 (Averaging Monotonicity of Binomial Function) *Assume that the function ϕ is convex. Then, $\mathsf{BF}(\mathbf{p}, r, \phi) \leq \mathsf{BF}(\widetilde{\mathbf{p}}, r, \phi)$.*

PROOF. Clearly, it suffices to prove that $\mathsf{BF}(\mathbf{p}, r, \phi)$ does not decrease when any two arbitrary probabilities in the vector $\mathbf{p}$ are replaced by their average. Since $\mathsf{BF}(\mathbf{p}, r, \phi)$ is symmetric in the probabilities p_i, with $i \in [r]$, it suffices to prove that $\mathsf{BF}(\mathbf{p}, r, \phi)$ does not decrease when p_1 and p_2 are replaced by their average

$\frac{p_1+p_2}{2}$. So, consider the vector of probabilities $\mathbf{q} = \langle q_1, \ldots, q_r \rangle$ with $q_1 = q_2 = \frac{p_1+p_2}{2}$, and $q_i = p_i$ for all $i \in [r] \setminus [2]$. Write

$$
\begin{aligned}
&\mathsf{BF}(\mathbf{p}, r, \phi) \\
&= \sum_{\mathcal{U} \subseteq [r]} \prod_{k \in \mathcal{U}} p_k \prod_{k \notin \mathcal{U}} (1 - p_k) \cdot \phi(|\mathcal{U}|) \\
&= \sum_{\mathcal{U} \subseteq [r] \setminus [2]} \prod_{k \in \mathcal{U}} p_k \prod_{k \notin \mathcal{U} \cup [2]} (1 - p_k) \cdot \big[(1 - p_1)(1 - p_2) \phi(|\mathcal{U}|) + \\
&\qquad\qquad + (p_1(1 - p_2) + p_2(1 - p_1)) \phi(|\mathcal{U}| + 1) + p_1 p_2 \phi(|\mathcal{U}| + 2)) \big] \\
&= \sum_{\mathcal{U} \subseteq [r] \setminus [2]} \prod_{k \in \mathcal{U}} p_k \prod_{k \notin \mathcal{U} \cup [2]} (1 - p_k) \cdot \big[p_1 p_2 \big(\phi(|\mathcal{U}| + 2) - 2\phi(|\mathcal{U}| + 1) + \phi(|\mathcal{U}|) \big) + \\
&\qquad\qquad + (p_1 + p_2) \left(\phi(|\mathcal{U}| + 1) - \phi(|\mathcal{U}|) \right) + \phi(|\mathcal{U}|) \big] ,
\end{aligned}
$$

so that also

$$
\begin{aligned}
&\mathsf{BF}(\mathbf{q}, r, \phi) \\
&= \sum_{\mathcal{U} \subseteq [r] \setminus [2]} \prod_{k \in \mathcal{U}} q_k \prod_{k \notin \mathcal{U} \cup [2]} (1 - q_k) \cdot \big[q_1 q_2 \big(\phi(|\mathcal{U}| + 2) - 2\phi(|\mathcal{U}| + 1) + \phi(|\mathcal{U}|) \big) + \\
&\qquad\qquad + (q_1 + q_2) \left(\phi(|\mathcal{U}| + 1) - \phi(|\mathcal{U}|) \right) + \phi(|\mathcal{U}|) \big] .
\end{aligned}
$$

Since $p_i = q_i$ for all users $i \in [r] \setminus [2]$, while $q_1 + q_2 = p_1 + p_2$, it follows that

$$
\begin{aligned}
&\mathsf{BF}(\mathbf{q}, r, \phi) - \mathsf{BF}(\mathbf{p}, r, \phi) \\
&= \sum_{\mathcal{U} \subseteq [r] \setminus [2]} \prod_{k \in \mathcal{U}} p_k \prod_{k \notin \mathcal{U} \cup [2]} (1 - p_k) \cdot \big[(q_1 q_2 - p_1 p_2) \left(\phi(|\mathcal{U}| + 2) - 2\phi(|\mathcal{U}| + 1) + \phi(|\mathcal{U}|) \right) \big] .
\end{aligned}
$$

Since $q_1 = q_2$ is the arithmetic mean of p_1 and p_2, it holds that $q_1 q_2 \geq p_1 p_2$. Since the function ϕ is convex, $\phi(|\mathcal{U}|+1) - \phi(|\mathcal{U}|) \leq \phi(|\mathcal{U}|+2) - \phi(|\mathcal{U}|+1)$; rearranging terms yields that $\phi(|\mathcal{U}|+2) - 2\phi(|\mathcal{U}|+1) + \phi(|\mathcal{U}|) \geq 0$. Thus, $\mathsf{BF}(\mathbf{q}, r, \phi) - \mathsf{BF}(\mathbf{p}, r, \phi) \geq 0$, as needed. ■

Another significant special case of the binomial function occurs when not only $\mathbf{w}$ is a constant vector with all entries equal to 1, but also $\mathbf{p}$ is a constant vector with all entries equal to $p > 0$. We shall analyze this special case under the additional assumption that ϕ is the *monomial* function $\phi(\lambda) = \lambda^d$. We will then abuse notation to write $\mathsf{BF}(p, r, d)$ for $\mathsf{BF}(\mathbf{p}, \mathbf{w}, \lambda^d)$. Clearly,

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

We shall later use a known fact about $\mathsf{BF}(p, r, d)$:

Lemma 2.3 (Gairing *et al.* [22]) *For any integer* $d \geq 1$, $\mathsf{BF}(p, r, d) = \sum_{k \in [d]} S(d, k) \cdot r^{\underline{k}} \cdot p^k$.

3. Discrete Routing Games

We extend definitions for the KP model to accommodate features from the Wardrop model.

3.1. *General*

We consider a simple *network* consisting of $m \geq 2$ parallel *links* $1, 2, \ldots, m$ from a *source* node to a *destination* node. Each of $n \geq 2$ *users* $1, 2, \ldots, n$ wishes to route a particular amount of (*unsplittable*) traffic along a (non-fixed) link from source to destination.

Denote as $w_i > 0$ the *traffic* of user $i \in [n]$. Define the $n \times 1$ *traffic vector* $\mathbf{w}$ in the natural way. For a user $i \in [n]$, eliminating w_i from the vector $\mathbf{w}$ yields the $(n-1)$-dimensional vector $\mathbf{w}_{-i}$. Associated with each link $j \in [m]$ is a *latency function* $\phi_j : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$, which is a *non-decreasing* and *non-constant* function with $\phi_j(0) = 0$. For each user $i \in [n]$, define the function $\phi_{ij} : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ by $\phi_{ij}(\lambda) = \phi_j(w_i + \lambda)$. We assume that $\phi_1(1) \leq \ldots \leq \phi_m(1)$; call link 1 the *smallest* link and say that link j is *smaller* than link k whenever $j < k$. Define the $m \times 1$ *latency function vector* Φ in the natural way. An *instance* is a pair $\langle \mathbf{w}, \Phi \rangle$.

In the case of *identical users*, all user traffics are 1 and an instance is a pair $\langle n, \Phi \rangle$. In the case of *identical links,* $\phi_j = \phi$ for all links $j \in [m]$, where ϕ is a non-decreasing and non-constant function (with $\phi(0) = 0$); in this case, an instance is a pair $\langle \mathbf{w}, \langle m, \phi \rangle \rangle$. For the case of identical users and identical links, an instance is a pair $\langle n, \langle m, \phi \rangle \rangle$. In the general case, we talk about *arbitrary users* and *arbitrary links.*

In the case of identical users, each latency function is a discrete function $\phi_j : [n] \cup \{0\} \rightarrow \mathbb{R}_0^+$ $\phi_j(0) = 0$; clearly, $\widehat{\phi_j} = \phi_{ij}$ for each (identical) user $i \in [n]$. In this case, say that the latency function ϕ_j is *non-constant on* $[n]$, or *non-constant* for short, if $\phi_j(1) \neq \phi_j(n)$; since each ϕ_j is non-decreasing, this implies that $\phi_j(1) < \phi_j(n)$.

3.2. *Convexity*

For the case of identical users, we will often assume that the latency functions enjoy some property of discrete convexity on the domain $[n] \cup \{0\}$. Formally, a function $\phi : [n] \cup \{0\} \rightarrow \mathbb{R}_0^+$ is *convex* if for all pairs of integers $x_1, x_2 \in [n-1]$ with $x_1 < x_2$, $\phi(x_1 + 1) - \phi(x_1) \leq \phi(x_2 + 1) - \phi(x_2)$. Clearly, to establish that such a function ϕ is convex, it suffices to prove that for any $x \in [n-1]$, $\phi(x) - \phi(x-1) \leq \phi(x+1) - \phi(x)$. Our definition of a convex function is the particular, one-dimensional case of a corresponding definition of *M-convex* functions due to Murota [38]:

For $B \subseteq \mathbb{N}_0^m$, a function $\phi : B \rightarrow \mathbb{R}$ is called *M-convex* if for all vectors $\mathbf{x}, \mathbf{y} \in B$ and for all dimensions $j \in [m]$ with $x(j) > y(j)$ there exists a dimension $k \in [m]$ with $x(k) < y(k)$ such that $\mathbf{x} - \boldsymbol{\chi}_j + \boldsymbol{\chi}_k \in B$, $\mathbf{y} + \boldsymbol{\chi}_j - \boldsymbol{\chi}_k \in B$, and

$$\phi(\mathbf{x}) + \phi(\mathbf{y}) \geq \phi(\mathbf{x} - \boldsymbol{\chi}_j + \boldsymbol{\chi}_k) + \phi(\mathbf{y} + \boldsymbol{\chi}_j - \boldsymbol{\chi}_k)\,.$$

We will later use a combinatorial property of M-convex functions, which was originally shown by Murota [38, Theorem 2.2]:

Proposition 3.1 (Global Optimality = Local Optimality for M-Convex) *Let ϕ be an M-convex function on $B \subseteq \mathbb{N}_0^m$. Then, the vector $\mathbf{x} \in B$ minimizes ϕ over B if and only if for all pairs of dimensions $j, k \in [m]$, $\phi(\mathbf{x}) \leq \phi(\mathbf{x} - \boldsymbol{\chi}_j + \boldsymbol{\chi}_k)$.*

3.3. *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 a tuple $\mathbf{L} = \langle \ell_1, \ldots, \ell_n \rangle \in [m]^n$; a *mixed assignment* is an $n \times m$ *probability matrix* $\mathbf{P}$ of $n \cdot m$ probabilities $p(i,j)$, for all pairs of a user $i \in [n]$ and a link $j \in [m]$, where $p(i,j)$ is the probability that user i chooses link j; so, for each user $i \in [n]$, $\sum_{j \in [m]} p(i,j) = 1$. We will cast a pure assignment as a special case of a mixed assignment in which all (mixed) strategies are pure. The *support* of the mixed strategy for user $i \in [n]$ in the mixed assignment $\mathbf{P}$, denoted as $\mathsf{Support}_{\mathbf{P}}(i)$, is the set of pure strategies which i chooses with strictly positive probability.

A mixed assignment $\mathbf{F}$ is *fully mixed* [36, Section 2.2] if $f(i,j) > 0$ for all pairs of a user $i \in [n]$ and a link $j \in [m]$. In the *standard* fully mixed assignment $\mathbf{F}$, $f(i,j) = \frac{1}{m}$ for all users $i \in [n]$ and links $j \in [m]$.

Fix now a mixed assignment $\mathbf{P}$. The *load* $\delta_j(\mathbf{P})$ on link $j \in [m]$ is the total traffic of users choosing the link (according to $\mathbf{P}$); so, $\delta_j(\mathbf{P})$ is a random variable. For each link $j \in [m]$, denote as $\theta_j(\mathbf{P})$ the *expected load* on link $j \in [m]$; so, clearly, $\theta_j(\mathbf{P}) = \sum_{k \in [n]} p(k,j)\, w_k$. Moreover, denote as $\theta_{ij}(\mathbf{P}) = \sum_{k \in [n] \setminus \{i\}} p(k,j)\, w_k$, the *expected load* on link $j \in [m]$ excluding user $i \in [n]$.

For a user $i \in [n]$ and a link $j \in [m]$, denote as $\mathbf{p}_{ij}$ the $(n-1)$-dimensional vector $\langle p(1,j), \ldots, p(i-1,j), p(i+1,j), \ldots, p(n,j) \rangle$. The average probability $\widetilde{p}_{ij}$ on link j excluding user i is defined as $\widetilde{p}_{ij} = \frac{\sum_{k \in [n] \setminus \{i\}} p(k,j)}{n-1}$;

clearly, in the case of identical users, $\widetilde{p}_{ij} = \frac{\theta_{ij}(\mathbf{P})}{n-1}$. It is straightforward to verify that $\sum_{j\in[m]} \widetilde{p}_{ij} = 1$. Denote as $\widetilde{\mathbf{p}}_{ij}$ the $(n-1)$-dimensional vector with all entries equal to $\widetilde{p}_{ij}$.

3.4. *Costs*

3.4.1. *Individual Cost and Expected Individual Cost*

For the pure assignment $\mathbf{L}$, the *Individual Cost* for user $i \in [n]$, denoted as $\mathsf{IC}_i(\mathbf{L})$, is

$$\mathsf{IC}_i(\mathbf{L}) = \phi_{\ell_i}\left(\sum_{k\in[n]|\ell_k=\ell_i} w_k\right) ;$$

so, the Individual Cost for user i is the latency on the link it chooses.

Fix now a mixed assignment $\mathbf{P}$. The *Conditional Expected Individual Cost* for user $i \in [n]$ on link $j \in [m]$, denoted as $\mathsf{IC}_{ij}(\mathbf{P})$, is the expectation (according to $\mathbf{P}$) of the Individual Cost for user i had it chosen link j; thus,

$$\mathsf{IC}_{ij}(\mathbf{P}) = \sum_{\mathcal{U}\subseteq[n]\setminus\{i\}} \prod_{k\in\mathcal{U}} p(k,j) \prod_{k\notin\mathcal{U}\cup\{i\}} (1-p(k,j)) \cdot \phi_j(w_i + w_{\mathcal{U}}) .$$

Since the latency function ϕ_j is non-decreasing, it follows that for all pairs of a user $i \in [n]$ and a link $j \in [m]$, $\mathsf{IC}_{ij}(\mathbf{P}) \geq \phi_j(w_i)$.

Using the binomial function, the Conditional Expected Individual Cost is expressed as

$$\mathsf{IC}_{ij}(\mathbf{P}) = \mathsf{BF}(\mathbf{p}_{ij}, \mathbf{w}_{-i}, \phi_{ij}) .$$

For the special case of identical users, this expression reduces to

$$\mathsf{IC}_{ij}(\mathbf{P}) = \mathsf{BF}(\mathbf{p}_{ij}, n-1, , \widehat{\phi}_j) .$$

For each user $i \in [n]$, the *Expected Individual Cost* for user i, denoted as $\mathsf{IC}_i(\mathbf{P})$, is the expectation (according to $\mathbf{P}$) of the Individual Cost of user i. Thus,

$$\mathsf{IC}_i(\mathbf{P}) = \sum_{j\in[m]} p(i,j) \cdot \mathsf{IC}_{ij}(\mathbf{P}) ;$$

so, Expected Individual Cost is a convex combination of Conditional Expected Individual Costs.

3.4.2. *Social Cost*

Associated with an instance $\langle \mathbf{w}, \Phi \rangle$ and a mixed assignment $\mathbf{P}$ is the *Social Cost,* denoted as $\mathsf{SC}^{\Sigma}(\mathbf{w}, \Phi, \mathbf{P})$, which is the sum, over all users, of Expected Individual Costs; so,

$$\begin{aligned}\mathsf{SC}^{\Sigma}(\mathbf{w}, \Phi, \mathbf{P}) &= \sum_{i\in[n]} \mathsf{IC}_i(\mathbf{P}) \\ &= \sum_{i\in[n]} \sum_{j\in[m]} p(i,j) \sum_{\mathcal{U}\subseteq[n]\setminus\{i\}} \prod_{k\in\mathcal{U}} p(k,j) \prod_{k\notin\mathcal{U}\cup\{i\}} (1-p(k,j)) \cdot \phi_j(w_i + w_{\mathcal{U}}) .\end{aligned}$$

3.4.3. *Optimum*

Associated with an instance $\langle \mathbf{w}, \Phi \rangle$ is the *Optimum*, denoted as $\mathsf{OPT}^{\Sigma}(\mathbf{w}, \Phi)$, which is the least possible, over all pure assignments, Social Cost; thus,

$$\mathsf{OPT}^{\Sigma}(\mathbf{w}, \Phi) = \min_{\mathbf{L}\in[m]^n} \mathsf{SC}^{\Sigma}(\mathbf{w}, \Phi, \mathbf{L}) .$$

A pure assignment $\mathbf{L}$ is *optimum* for the instance $\langle \mathbf{w}, \Phi \rangle$ if $\mathsf{SC}^{\Sigma}(\mathbf{w}, \Phi, \mathbf{L}) = \mathsf{OPT}^{\Sigma}(\mathbf{w}, \Phi)$.

We note two obvious lower bounds on Optimum for the case of identical users and identical links with a *monomial* latency function $\phi(\lambda) = \lambda^d$, for some integer $d \geq 1$. Assuming that $n \geq m$, $\mathsf{OPT}^{\Sigma}(n, \langle m, \phi \rangle) \geq n \cdot \left(\frac{n}{m}\right)^d$; assuming that $n \leq m$, $\mathsf{OPT}^{\Sigma}(n, \langle m, \phi \rangle) = n$.

3.5. *Nash Equilibria and Price of Anarchy*

We are interested in a special class of mixed strategies called Nash equilibria [39,40]. Given an instance $\langle \mathbf{w}, \Phi \rangle$ with an associated mixed assignment $\mathbf{P}$, a user $i \in [n]$ is *satisfied* in $\mathbf{P}$ if $\mathsf{IC}_{ij}(\mathbf{P}) = \mathsf{IC}_i(\mathbf{P})$ for all links $j \in \mathsf{Support}_{\mathbf{P}}(i)$, and $\mathsf{IC}_{ij}(\mathbf{P}) \geq \mathsf{IC}_i(\mathbf{P})$ for all links $j \notin \mathsf{Support}_{\mathbf{P}}(i)$. So, a satisfied user has no incentive to unilaterally deviate from its mixed strategy. The mixed assignment $\mathbf{P}$ is a *Nash equilibrium* if all users $i \in [n]$ are satisfied in $\mathbf{P}$.

A *fully mixed Nash equilibrium* [36] is a fully mixed assignment that is a Nash equilibrium. Note that for the case of identical links, the standard fully mixed assignment $\mathbf{F}$ is a (fully mixed) Nash equilibrium since it satisfies that for all users $i \in [n]$ and for all pairs of links $j, l \in [m]$, $\mathsf{IC}_{ij}(\mathbf{P}) = \mathsf{IC}_{il}(\mathbf{P})$; call it the *standard fully mixed Nash equilibrium.*

The *Price of Anarchy,* denoted as PoA^{Σ}, is the *worst-case* ratio $\frac{\mathsf{SC}^{\Sigma}(\mathbf{w},\Phi,\mathbf{P})}{\mathsf{OPT}^{\Sigma}(\mathbf{w},\Phi)}$, over all instances $\langle \mathbf{w}, \Phi \rangle$ and associated Nash equilibria $\mathbf{P}$.

A *worst-case* (or *worst* for short) Nash equilibrium [31] is one which, on any fixed but arbitrary instance, maximizes Social Cost. A *best-case* (or *best* for short) Nash equilibrium [18] is one which, on any fixed but arbitrary instance, minimizes Social Cost. The *Fully Mixed Nash Equilibrium Conjecture,* henceforth abbreviated as the FMNE *Conjecture,* states that a fully mixed Nash equilibrium is a worst-case Nash equilibrium.

4. Fully Mixed Nash Equilibria

We now focus on fully mixed Nash equilibria; we restrict to the special case of identical users. A preliminary property of fully mixed Nash equilibria is shown in Section 4.1. In Section 4.2, we formulate and prove the Fully Mixed Nash Equilibrium Conjecture under the (essential) assumption of convex latency functions. Furthermore, we establish in Section 4.3 that the fully mixed Nash equilibrium may only exist uniquely. In Section 4.4, we prove a combinatorial characterization of instances admitting a fully mixed Nash equilibrium.

4.1. *A Preliminary Property*

We show a preliminary property of fully mixed Nash equilibria, which applies to the case of identical users. Specifically, we prove that all (identical) users choose each fixed link with the same probability.

Lemma 4.1 (Same Probabilities) *Consider the case of identical users and arbitrary links with non-constant latency functions. Fix an instance $\langle n, \Phi \rangle$ with an associated fully mixed Nash equilibrium $\mathbf{F}$. Then, for all pairs of users $i, h \in [n]$ and for all links $j \in [m]$, $f(i,j) = f(h,j)$.*

PROOF. Fix any pair of users $i, h \in [n]$ and a link $j \in [m]$. Since $\mathbf{F}$ is a fully mixed Nash equilibrium,

$$\begin{aligned}\mathsf{IC}_i(\mathbf{F}) &= \mathsf{IC}_{ij}(\mathbf{F}) \\ &= \sum_{\mathcal{U} \subseteq [n] \setminus \{i,h\}} \prod_{k \in \mathcal{U}} f(k,j) \prod_{k \in [n] \setminus (\mathcal{U} \cup \{i,h\})} (1 - f(k,j)) \cdot \left[(1 - f(h,j)) \phi_j(|\mathcal{U}| + 1) + f(h,j) \phi_j(|\mathcal{U}| + 2) \right]\end{aligned}$$

and also

$$\mathsf{IC}_h(\mathbf{F}) = \sum_{\mathcal{U} \subseteq [n] \setminus \{i,h\}} \prod_{k \in \mathcal{U}} f(k,j) \prod_{k \in [n] \setminus (\mathcal{U} \cup \{i,h\})} (1 - f(k,j)) \cdot \left[(1 - f(i,j)) \phi_j(|\mathcal{U}| + 1) + f(i,j) \phi_j(|\mathcal{U}| + 2) \right].$$

Hence,

$$
\begin{aligned}
&\mathsf{IC}_i(\mathbf{F}) - \mathsf{IC}_h(\mathbf{F}) \\
&= \sum_{\mathcal{U}\subseteq[n]\setminus\{i,h\}} \prod_{k\in\mathcal{U}} f(k,j) \prod_{k\in[n]\setminus(\mathcal{U}\cup\{i,h\})} (1-f(k,j))\cdot(f(h,j)-f(i,j))\cdot(\phi_j(|\mathcal{U}|+2)-\phi_j(|\mathcal{U}|+1)) \\
&= (f(h,j)-f(i,j)) \sum_{\mathcal{U}\subseteq[n]\setminus\{i,h\}} \prod_{k\in\mathcal{U}} f(k,j) \prod_{k\in[n]\setminus(\mathcal{U}\cup\{i,h\})} (1-f(k,j))\cdot(\phi_j(|\mathcal{U}|+2)-\phi_j(|\mathcal{U}|+1))\ .
\end{aligned}
$$

Since $\mathbf{F}$ is fully mixed, $\prod_{k\in\mathcal{U}} f(k,j)\prod_{k\in[n]\setminus(\mathcal{U}\cup\{i,h\})}(1-f(k,j)) > 0$ for all (non-empty) subsets $\mathcal{U} \subseteq [n]\setminus\{i,h\}$. Since ϕ_j is non-decreasing and non-constant on $[n]$, there is a (non-empty) subset $\mathcal{U}\subseteq[n]\setminus\{i,h\}$ such that $\phi_j(|\mathcal{U}|+2) > \phi_j(|\mathcal{U}|+1)$. These imply that

$$
\sum_{\mathcal{U}\subseteq[n]\setminus\{i,h\}} \prod_{k\in\mathcal{U}} f(k,j) \prod_{k\in[n]\setminus(\mathcal{U}\cup\{i,h\})} (1-f(k,j))\ (\phi_j(|\mathcal{U}|+2)-\phi_j(|\mathcal{U}|+1)) > 0\,.
$$

It follows that $\mathsf{IC}_i(\mathbf{F}) > \mathsf{IC}_h(\mathbf{F})$ (resp., $\mathsf{IC}_i(\mathbf{F}) < \mathsf{IC}_h(\mathbf{F})$) if and only if $f(h,j) > f(i,j)$ (resp., $f(h,j) < f(i,j)$). Since the link $j\in[m]$ was chosen arbitrarily, this implies that $\mathsf{IC}_i(\mathbf{F}) > \mathsf{IC}_h(\mathbf{F})$ (resp., $\mathsf{IC}_i(\mathbf{F}) < \mathsf{IC}_h(\mathbf{F})$) if and only if for all links $l\in[m]$, $f(h,l) > f(i,l)$ (resp., $f(h,l) < f(i,l)$). Since, however, $\sum_{l\in[m]} f(h,l) = \sum_{l\in[m]} f(i,l) = 1$, the latter is false. It follows that $\mathsf{IC}_i(\mathbf{F}) = \mathsf{IC}_h(\mathbf{F})$. This implies that $f(h,j) = f(i,j)$, as needed. ■

Lemma 4.1 implies that a fully mixed Nash equilibrium $\mathbf{F}$ can be identified with a sequence $f_j \in (0,1)$, $j\in[m]$, such that for all pairs of a user $i\in[n]$ and a link $j\in[m]$, $f(i,j) = f_j$.

4.2. *The* FMNE *Conjecture*

Convex latency functions are considered in Section 4.2.1. Section 4.2.2 considers arbitrary latency functions. Throughout this section, we keep restricting to the case of identical users.

4.2.1. *Convex Latency Functions*

We prove:

Proposition 4.2 *Consider the case of identical users and arbitrary links with non-constant and convex latency functions. Fix an instance $\langle n, \Phi\rangle$ with an associated fully mixed Nash equilibrium $\mathbf{F}$. Then, for each Nash equilibrium $\mathbf{P}$ and for each user $i\in[n]$, $\mathsf{IC}_i(\mathbf{P}) \leq \mathsf{IC}_i(\mathbf{F})$.*

PROOF. By Lemma 4.1, there is, for each link $j\in[m]$, some probability $f_j\in(0,1)$ such that each user chooses link j with probability f_j. Denote as $\mathbf{f}_j$ the $(n-1)$-dimensional vector of probabilities with all entries equal to f_j.

Since Expected Individual Cost is a convex combination of Conditional Expected Individual Costs, it suffices to prove that for all pairs of a user $i\in[n]$ and a link $j\in\mathsf{Support}_{\mathbf{P}}(i)$, $\mathsf{IC}_{ij}(\mathbf{P}) \leq \mathsf{IC}_i(\mathbf{F})$.

Assume, by way of contradiction, that there is a user $i\in[n]$ and a link $j\in\mathsf{Support}_{\mathbf{P}}(i)$ such that $\mathsf{IC}_{ij}(\mathbf{P}) > \mathsf{IC}_i(\mathbf{F})$. Then,

$$
\begin{aligned}
\mathsf{IC}_{ij}(\mathbf{P}) &= \mathsf{BF}(\mathbf{p}_{ij}, n-1, \widehat{\phi}_j) \\
&\leq \mathsf{BF}(\widetilde{\mathbf{p}}_{ij}, n-1, \widehat{\phi}_j) \quad \text{(by Lemma 2.2)}\ .
\end{aligned}
$$

On the other hand,

$$
\begin{aligned}
\mathsf{IC}_i(\mathbf{F}) &= \mathsf{IC}_{ij}(\mathbf{F}) \qquad\qquad \text{(since } \mathbf{F} \text{ is a fully mixed Nash equilibrium)} \\
&= \mathsf{BF}(\mathbf{f}_j, n-1, \widehat{\phi}_j)\,.
\end{aligned}
$$

Since $\mathsf{IC}_{ij}(\mathbf{P}) > \mathsf{IC}_i(\mathbf{F})$, it follows that $\mathsf{BF}(\widetilde{\mathbf{p}}_{ij}, n-1, \widehat{\phi}_j) > \mathsf{BF}(\mathbf{f}_j, n-1, \widehat{\phi}_j)$. Since $\widehat{\phi}_j$ is non-decreasing and non-constant, it follows from Lemma 2.1 that $\widetilde{p}_{ij} > f_j$. Clearly, $\sum_{l\in[m]} f_l = \sum_{l\in[m]} f(i,l) = 1$. Recall that $\sum_{l\in[m]} \widetilde{p}_{il} = 1$. It follows that there is a link $l\in[m]$ such that $\widetilde{p}_{il} < f_l$. Thus,

$$
\begin{aligned}
\mathsf{IC}_{ij}(\mathbf{P}) &\leq \mathsf{IC}_{il}(\mathbf{P}) && \text{(since } j \in \mathsf{Support}_{\mathbf{P}}(i) \text{ and } \mathbf{P} \text{ is a Nash equilibrium)}\\
&= \mathsf{BF}(\mathbf{p}_{il}, n-1, \widehat{\phi}_l) && \\
&\leq \mathsf{BF}(\widetilde{\mathbf{p}}_{il}, n-1, \widehat{\phi}_l) && \text{(by Lemma 2.2)}\\
&< \mathsf{BF}(\mathbf{f}_l, n-1, \widehat{\phi}_l) && \text{(by Lemma 2.1 (since } \widetilde{p}_{il} < f_l\text{))}\\
&= \mathsf{IC}_{il}(\mathbf{F}) && \\
&= \mathsf{IC}_{i}(\mathbf{F}) && \text{(since } \mathbf{F} \text{ is a fully mixed Nash equilibrium)}\\
&< \mathsf{IC}_{ij}(\mathbf{P}) && \text{(by assumption)}\ ,
\end{aligned}
$$

a contradiction. ■

Since Social Cost is the sum of Expected Individual Costs, Proposition 4.2 directly implies:

Theorem 4.3 (Convexity Implies the FMNE Conjecture) *Consider the case of identical users and arbitrary links with non-constant and convex latency functions. Then, the Fully Mixed Nash Equilibrium Conjecture is valid.*

4.2.2. *Arbitrary Latency Functions*

We now provide a counterexample to Theorem 4.3. More specifically, we construct an instance involving identical links with a non-decreasing and non-constant but *not* convex latency function for which Proposition 4.2 does *not* hold. We prove:

Proposition 4.4 *Consider the case of identical users and identical links with an arbitrary latency function. Then, there is an instance $\langle n, \langle m, \phi\rangle\rangle$ with associated pure Nash equilibrium $\mathbf{L}$ and fully mixed Nash equilibrium $\mathbf{F}$ such that for all users $i \in [n]$, $\mathsf{IC}_i(\mathbf{L}) > \mathsf{IC}_i(\mathbf{F})$.*

PROOF. Consider the instance $\langle 4, \langle 2, \phi\rangle\rangle$, where $\phi : [4] \cup \{0\} \rightarrow \mathbb{R}_0^+$ is a (strictly increasing) function with $\phi(1) = 1$, $\phi(2) = 2$, $\phi(3) = \frac{13}{6}$ and $\phi(4) = \frac{7}{3}$. Since $\phi(3) - \phi(2) < \phi(2) - \phi(1)$, ϕ is *not* convex. Consider any arbitrary pure Nash equilibrium $\mathbf{L}$ and the standard fully mixed Nash equilibrium $\mathbf{F}$, where $f(i,j) = \frac{1}{2}$ for each user $i \in [4]$ and link $j \in [2]$.

Fix any arbitrary user $i \in [4]$. We compare the Individual Cost $\mathsf{IC}_i(\mathbf{L})$ and the Expected Individual Cost $\mathsf{IC}_i(\mathbf{F})$ for user i.

On one hand, note that since $\phi(2) < \phi(3)$ and $\phi(1) < \phi(4)$, the definition of Nash equilibrium implies that in $\mathbf{L}$, there is no link chosen by either 3 or 4 users. So, exactly two users choose each link in $\mathbf{L}$; hence, $\mathsf{IC}_i(\mathbf{L}) = \phi(2) = 2$. On the other hand, $\mathsf{IC}_i(\mathbf{F}) = \frac{1}{8}(\phi(1) + 3\phi(2) + 3\phi(3) + \phi(4)) = \frac{95}{48}$. Since $2 > \frac{95}{48}$. it follows that $\mathsf{IC}_i(\mathbf{L}) > \mathsf{IC}_i(\mathbf{F})$, as needed. ■

Since Social Cost is the sum of Expected Individual Costs, Proposition 4.4 directly implies:

Corollary 4.5 (The FMNE Conjecture Needs Convexity) *Consider the case of identical users and identical links with an arbitrary latency function. Then, the FMNE Conjecture is not valid.*

Corollary 4.5 implies that the assumption of convexity made for Theorem 4.3 is essential.

4.3. *Uniqueness*

We show:

Theorem 4.6 (Fully Mixed Nash Equilibrium Uniqueness) *Consider the case of identical users and arbitrary links with non-constant latency functions. Then, a fully mixed Nash equilibrium may exist only uniquely.*

PROOF. Assume, by way of contradiction, that there is an instance $\langle n, \Phi\rangle$ with two distinct associated fully mixed Nash equilibria $\mathbf{F}$ and $\mathbf{G}$. Lemma 4.1 implies that for each link $j \in [m]$, there is a probability $f_j \in (0,1)$ (resp., $g_j \in (0,1)$) such that for all users $i \in [n]$, $f(i,j) = f_j$ (resp., $g(i,j) = g_j$). For each link

$j \in [m]$, denote as $\mathbf{f}_j$ (resp., $\mathbf{g}_j$) the $(n-1)$-dimensional vector of probabilities with all entries equal to f_j (resp., g_j).

Since $\mathbf{F}$ and $\mathbf{G}$ are distinct with $\sum_{l\in[m]} f_l = \sum_{l\in[m]} g_l = 1$, there are two distinct links $j, l \in [m]$ such that $f_j > g_j$ and $f_l < g_l$. Fix any user $i \in [n]$. Clearly, for the fully mixed Nash equilibria $\mathbf{F}$ and $\mathbf{G}$, the Conditional Expected Individual Costs for user i on the link $j \in [m]$ are $\mathsf{IC}_{ij}(\mathbf{F}) = \mathsf{BF}(\mathbf{f}_j, n-1, \widehat{\phi}_j)$ and $\mathsf{IC}_{ij}(\mathbf{G}) = \mathsf{BF}(\mathbf{g}_j, n-1, \widehat{\phi}_j)$. Hence,

$$\begin{aligned}
\mathsf{IC}_{ij}(\mathbf{F}) &= \mathsf{IC}_{il}(\mathbf{F}) && \text{(since } \mathbf{F} \text{ is a fully mixed Nash equilibrium)}\\
&= \mathsf{BF}(\mathbf{f}_l, n-1, \widehat{\phi}_l) \\
&< \mathsf{BF}(\mathbf{g}_l, n-1, \widehat{\phi}_l) && \text{(by Lemma 2.1 (since } f_l < g_l\text{))}\\
&= \mathsf{IC}_{il}(\mathbf{G}) \\
&= \mathsf{IC}_{ij}(\mathbf{G}) && \text{(since } \mathbf{G} \text{ is a fully mixed Nash equilibrium)}\\
&= \mathsf{BF}(\mathbf{g}_j, n-1, \widehat{\phi}_j) \\
&< \mathsf{BF}(\mathbf{f}_j, n-1, \widehat{\phi}_j) && \text{(by Lemma 2.1 (since } g_j < f_j\text{))}\\
&= \mathsf{IC}_{ij}(\mathbf{F})\,,
\end{aligned}$$

a contradiction. ■

Consider now the case of identical users and identical links, and recall the standard fully mixed Nash equilibrium. Then, Theorem 4.6 immediately implies:

Corollary 4.7 *Consider the case of identical users and identical links. Then, the standard fully mixed assignment is the unique fully mixed Nash equilibrium.*

4.4. *Existence*

We present a characterization of instances admitting a fully mixed Nash equilibrium. Recall our earlier assumption that $\phi_1(1) \leq \phi_2(1) \leq \ldots \leq \phi_n(1)$.

Fix an instance $\langle n, \Phi \rangle$. For each link $k \in [m]$ and for each smaller link $j \in [k-1]$, denote as $p_j(k)$ the probability such that

$$\mathsf{BF}(\mathbf{p}_j(k), n-1, \widehat{\phi}_j) = \min\{\phi_j(n), \phi_k(1)\}\ ,$$

where $\mathbf{p}_j(k)$ is the vector with $n-1$ entries equal to $p_j(k)$. We argue that this definition uniquely determines a probability $p_j(k)$. Recall that $\mathsf{BF}(\mathbf{p}_j(k), n-1, \widehat{\phi}_j)$ is the Conditional Expected Individual Cost for a user on link $j \in [m]$ in the case where $p(i,j) = p_j(k)$ for all remaining users $i \in [n-1]$. In particular, $\mathsf{BF}(\mathbf{0}, n-1, \widehat{\phi}_j) = \phi_j(1)$ and $\mathsf{BF}(\mathbf{1}, n-1, \widehat{\phi}_j) = \phi_j(n)$. Note that

$$\begin{aligned}
\phi_j(1) &= \min\{\phi_j(n), \phi_j(1)\} && \text{(since } \phi_j \text{ is non-decreasing)}\\
&\leq \min\{\phi_j(n), \phi_k(1)\} && \text{(since } \phi_j(1) \leq \phi_k(1)\text{)}\ ,
\end{aligned}$$

while $\phi_j(n) \geq \min\{\phi_j(n), \phi_k(1)\}$. So, $\phi_j(1) \leq \min\{\phi_j(n), \phi_k(1)\} \leq \phi_j(n)$. By Lemma 2.1, $\mathsf{BF}(\mathbf{p}_j(k), n-1, \widehat{\phi}_j)$ is strictly increasing in $p_j(k)$. By continuity of the binomial function, this implies that BF attains exactly once the intermediate value $\min\{\phi_j(n), \phi_k(1)\}$. Hence, the definition uniquely determines a probability $p_j(k)$.

We now continue with two important definitions.

Definition 4.1 (Dead and Special Links)

- *A link $k \in [m]$ is* dead *for the instance $\langle n, \Phi \rangle$ if either* (i) $\phi_j(n) < \phi_k(1)$ *for some smaller link $j \in [k-1]$, or* (ii) $\sum_{j\in[k-1]} p_j(k) > 1$.
- *A link $k \in [m]$ is* special *for the instance $\langle n, \Phi \rangle$ if $\sum_{j\in[k-1]} p_j(k) = 1$.*

We continue to prove some properties of dead and special links. The first of them shows that no user chooses a dead link.

Lemma 4.8 (No User Chooses a Dead Link) *Consider the case of identical users and arbitrary links with non-constant and convex latency functions. Fix an instance $\langle n, \Phi \rangle$ with a dead link $k \in [m]$ and an associated Nash equilibrium $\mathbf{P}$. Then, $p(i,k) = 0$ for all users $i \in [n]$.*

PROOF. Assume, by way of contradiction, that there is a user $i \in [n]$ such that $p(i,k) > 0$. Since $\mathbf{P}$ is a Nash equilibrium, this implies that for any link $j \in [k-1]$, $\mathsf{IC}_{ik}(\mathbf{P}) \leq \mathsf{IC}_{ij}(\mathbf{P})$. Since $\phi_k(1) \leq \mathsf{IC}_{ik}(\mathbf{P})$. it follows that $\phi_k(1) \leq \mathsf{IC}_{ij}(\mathbf{P})$.

Since k is a dead link, there are two cases to consider. For each case, we will derive a contradiction.

(i) Assume first that $\phi_l(n) < \phi_k(1)$ for some smaller $l \in [k-1]$. Then, by Lemma 2.1,

$$\begin{aligned}
\mathsf{IC}_{il}(\mathbf{P}) &= \mathsf{BF}(\mathbf{p}_{il}, n-1, \widehat{\phi}_l) \\
&\leq \mathsf{BF}(\widetilde{\mathbf{p}}_{il}, n-1, \widehat{\phi}_l) && \text{(by Lemma 2.2)} \\
&\leq \mathsf{BF}(\mathbf{1}, n-1, \widehat{\phi}_l) && \text{(by Lemma 2.1 (since } \widetilde{p}_{ij} \leq 1)) \\
&= \phi_l(n) \\
&< \phi_k(1) && \text{(by assumption)} \ ,
\end{aligned}$$

a contradiction.

(ii) Assume now that $\sum_{j\in[k-1]} p_j(k) > 1$. Assume that for each smaller link $j \in [k-1]$, $\phi_k(1) \leq \phi_j(n)$ (since otherwise the claim follows from the previous case). Note that $\sum_{j\in[k-1]} \widetilde{p}_j \leq \sum_{j\in[m]} \widetilde{p}_j = 1$. It follows that there is some smaller link $l \in [k-1]$ such that $p_l(k) > \widetilde{p}_{il}$. Hence,

$$\begin{aligned}
\mathsf{IC}_{il}(\mathbf{P}) &= \mathsf{BF}(\mathbf{p}_{il}, n-1, \widehat{\phi}_l) \\
&\leq \mathsf{BF}(\widetilde{\mathbf{p}}_{il}, n-1, \widehat{\phi}_l) && \text{(by Lemma 2.2)} \\
&< \mathsf{BF}(\mathbf{p}_l(k), n-1, \widehat{\phi}_l) && \text{(by Lemma 2.1 (since } \widetilde{p}_{il} < p_l(k))) \\
&= \min\{\phi_l(n), \phi_k(1)\} && \text{(by definition of } p_l(k)) \\
&= \phi_k(1)\,,
\end{aligned}$$

a contradiction.

Since we derived a contradiction in all possible cases, the proof is now complete. ■

We continue to prove:

Lemma 4.9 (At Most One User Chooses a Special Link) *Consider the case of identical users and arbitrary links with non-constant and convex latency functions. Fix an instance $\langle n, \Phi \rangle$ with an associated Nash equilibrium $\mathbf{P}$. Then, there is at most one user $i \in [n]$ with $p(i,k) > 0$ for some special link $k \in [m]$.*

PROOF. Assume, by way of contradiction, that there are two distinct users $i, h \in [n]$ and two (not necessarily distinct) special links $k, r \in [m]$, $k \leq r$, with $p(i,k) > 0$ and $p(h,r) > 0$.

Since $p(h,r) > 0$, it follows that $\widetilde{p}_{ir} \geq \frac{p(h,r)}{n-1} > 0$. Since $k \leq r \leq m$, this implies that $\sum_{j\in[k-1]} \widetilde{p}_{ij} < \sum_{j\in[m]} \widetilde{p}_{ij} = 1$. Since link k is special, $\sum_{j\in[k-1]} p_j(k) = 1$. It follows that there is some link $l \in [k-1]$ such that $p_l(k) > \widetilde{p}_{il}$. So,

$$\begin{aligned}
\phi_k(1) &\leq \mathsf{IC}_{ik}(\mathbf{P}) \\
&\leq \mathsf{IC}_{il}(\mathbf{P}) && \text{(since } p(i,k) > 0 \text{ and } \mathbf{P} \text{ is a Nash equilibrium)} \\
&= \mathsf{BF}(\mathbf{p}_{il}, n-1, \widehat{\phi}_l) \\
&\leq \mathsf{BF}(\widetilde{\mathbf{p}}_{il}, n-1, \widehat{\phi}_l) && \text{(by Lemma 2.2)} \\
&< \mathsf{BF}(\mathbf{p}_l(k), n-1, \widehat{\phi}_l) && \text{(by Lemma 2.1 (since } p_{il} < p_l(k))) \\
&= \min\{\phi_l(n), \phi_k(1)\} && \text{(by definition of } p_l(k)) \\
&\leq \phi_k(1)\,,
\end{aligned}$$

a contradiction. ■

We are now ready to prove:

Theorem 4.10 (Existence of Fully Mixed Nash Equilibria) *Consider the case of identical users and arbitrary links with non-constant and convex latency functions. Then, there is a fully mixed Nash equilibrium if and only if there are neither dead nor special links.*

PROOF. Throughout, fix an instance $\langle n, \Phi \rangle$.

Assume first that $\langle n, \Phi \rangle$ admits a fully mixed Nash equilibrium $\mathbf{F}$. By definition, $f(i,j) > 0$ for all pairs of a user $i \in [n]$ and a link $j \in [m]$. Lemma 4.8 implies that there is no dead link; since $n \geq 2$, Lemma 4.9 implies that there is no special link either, and we are done.

Assume now that there are neither dead nor special links for the instance $\langle n, \Phi \rangle$. We will determine a fully mixed Nash equilibrium $\mathbf{F}$ for $\langle n, \Phi \rangle$ with $f(i,j) = f_j$ for all users $i \in [n]$ and links $j \in [m]$.

For each link $j \in [m]$, define $\Delta\phi_j = \phi_j(n) - \phi_m(1)$. Clearly, $\Delta\phi_j \leq \phi_j(n)$. Since there are no dead links, it follows that $\Delta\phi_j \geq 0$. So, $0 \leq \Delta\phi_j \leq \phi_j(n)$.

Fix now a link $j \in [m-1]$. For any value $x \in [0, \Delta\phi_j]$, denote $\psi_j(x)$ the value such that

$$\mathsf{BF}(\boldsymbol{\psi}_j(x), n-1, \widehat{\phi_j}) = \phi_m(1) + x \,,$$

where $\boldsymbol{\psi}_j(x)$ is the $(n-1)$-dimensional vector with all entries equal to $\psi_j(x)$. We argue that $\psi_j(x)$ is uniquely determined by this definition.

Note that $\mathsf{BF}(\boldsymbol{\psi}_j(x), n-1, \widehat{\phi_j})$ is the Conditional Expected Individual Cost for a user on link $j \in [m-1]$ in the case where all remaining users $i \in [n-1]$ choose link j with probability $\psi_j(x)$. In particular, $\mathsf{BF}(\mathbf{0}, n-1, \widehat{\phi_j}) = \phi_j(1) \leq \phi_m(1)$ and $\mathsf{BF}(\mathbf{1}, n-1, \widehat{\phi_j}) = \phi_j(n)$. Note also that

$$\begin{aligned} \phi_m(1) + x &\leq \phi_m(1) + \Delta\phi_j & \text{(since } x \leq \Delta\phi_j\text{)} \\ &= \phi_m(1) + \phi_j(n) - \phi_m(1) & \\ &= \phi_j(n) \,, & \end{aligned}$$

while

$$\phi_m(1) + x \geq \phi_m(1) \,.$$

By Lemma 2.1, $\mathsf{BF}(\boldsymbol{\psi}_j(x), n-1, \widehat{\phi_j})$ is strictly increasing in $\psi_j(x)$. By continuity of the binomial function, this implies that BF attains exactly once the intermediate value $\phi_m(1) + x$. Hence, the definition uniquely determines $\psi_j(x)$.

Note that the definition of $\psi_j(x)$ and Lemma 2.1 imply together that $\psi_j(x)$ is strictly increasing in $x \in [0, \Delta\phi_j]$, where $j \in [m-1]$; in particular, this implies that for any $x \in (0, \Delta\phi_j)$, $0 < \psi_j(x) < 1$.

Now for the link m, for each $x \in [0, \min_{j \in [m-1]} \Delta\phi_j]$, set

$$\psi_m(x) = 1 - \sum_{j \in [m-1]} \psi_j(x) \,.$$

Clearly, $\psi_m(x)$ is strictly decreasing in x for $x \in [0, \min_{j \in [m-1]} \Delta\phi_j]$. Moreover, $\psi_m(x) < 1$ for all $x \in (0, \min_{j \in [m-1]} \Delta\phi_j]$.

Define

$$\widehat{x} = \max \left\{ x \in [0, \min_{j \in [m-1]} \Delta\phi_j] \mid \psi_m(x) \geq 0 \right\} ;$$

thus, $\psi_m(\widehat{x}) = 0$.

Consider the function $\mathsf{BF}(\boldsymbol{\psi}_m(x), n-1, \widehat{\phi_m})$, where $\boldsymbol{\psi}_m(x)$ is the $(n-1)$-dimensional vector with all entries equal to $\psi_m(x)$, for $x \in [0, \widehat{x}]$. Since $\psi_m(x)$ is strictly decreasing in x for $x \in [0, \widehat{x}]$, Lemma 2.1 implies that $\mathsf{BF}(\boldsymbol{\psi}_m(x), n-1, \widehat{\phi_m})$ is strictly decreasing in x for $x \in [0, \widehat{x}]$.

Note that by definition of $\psi_m(x)$,

$$\psi_m(0) = 1 - \sum_{j \in [m-1]} \psi_j(0) \,.$$

Recall also that for each link $j \in [m-1]$, $\psi_j(0)$ is (uniquely) determined by the equation

$$\mathsf{BF}(\boldsymbol{\psi}_j(0), n-1, \widehat{\phi}_j) = \phi_m(1)\,.$$

Since there are no dead links, this equation is equivalent to

$$\mathsf{BF}(\boldsymbol{\psi}_j(0), n-1, \widehat{\phi}_j) = \min\left\{\phi_j(n), \phi_m(1)\right\}\,.$$

By definition of $p_j(m)$, it follows that $\psi_j(0) = p_j(m)$. Hence,

$$\psi_m(0) = 1 - \sum_{j\in[m-1]} p_j(m)\,.$$

Since there are neither dead nor special links, it follows that $\sum_{j\in[m-1]} p_j(m) < 1$, so that $\psi_m(0) > 0$. Thus,

$$\begin{aligned}\mathsf{BF}(\boldsymbol{\psi}_m(0), n-1, \widehat{\phi}_m) &> \mathsf{BF}(\mathbf{0}, n-1, \widehat{\phi}_m) \quad \text{(by Lemma 2.1 (since } \psi_m(0) > 0))\\ &= \phi_m(1)\,.\end{aligned}$$

On the other hand,

$$\begin{aligned}\mathsf{BF}(\boldsymbol{\psi}_m(\widehat{x}), n-1, \widehat{\phi}_m) &= \mathsf{BF}(\mathbf{0}, n-1, \widehat{\phi}_m) \quad \text{(since } \psi_m(\widehat{x}) = 0)\\ &= \phi_m(1)\\ &< \phi_m(1) + \widehat{x} \qquad \text{(since } \widehat{x} > 0)\,.\end{aligned}$$

Since $\mathsf{BF}(\boldsymbol{\psi}_m(x), n-1, \widehat{\phi}_m)$ is a continuous, strictly decreasing function in x for $x \in [0, \widehat{x}]$, the Mean Value Theorem implies that there is some $x_0 \in (0, \widehat{x})$ such that

$$\mathsf{BF}(\boldsymbol{\psi}_m(x_0), n-1, \widehat{\phi}_m) = \phi_m(1) + x_0\,.$$

We are now ready to determine a fully mixed Nash equilibrium $\mathbf{F}$ for the instance $\langle n, \Phi\rangle$:
– For each user $i \in [n]$ and link $j \in [m]$, set $f(i,j) = \psi_j(x_0)$.
It remains to show that $\mathbf{F}$ is a fully mixed Nash equilibrium:
– We first prove that $\mathbf{F}$ is a fully mixed assignment. We need to prove that for each link $j \in [m]$, $0 < \psi_j(x_0) < 1$.
  · For a link $j \in [m-1]$, note that $0 < x_0 < \widehat{x} \leq \min_{j\in[m-1]} \Delta\phi_j \leq \Delta\phi_j$. Thus, $0 < \psi_j(x_0) < 1$, as needed.
  · For the link m, note that $0 < x_0 < \widehat{x} \leq \min_{j\in[m-1]} \Delta\phi_j$. Thus, $\psi_m(x_0) < 1$. Since $\psi_m(x)$ is strictly decreasing in x with $\psi_m(\widehat{x}) = 0$, while $x_0 < \widehat{x}$, it follows that $\psi_m(x_0) > 0$. So, $0 < \psi_m(x_0) < 1$, as needed.
– We finally prove that $\mathbf{F}$ is a Nash equilibrium. Recall that by construction of x_0, $\mathsf{B}\left(\boldsymbol{\psi}_m(x_0), n-1, \widehat{\phi}_m\right) = \phi_m(1) + x_0$. By definition of the value $\psi_j(x)$ for each link $j \in [m-1]$, $\mathsf{BF}(\boldsymbol{\psi}_j(x_0), n-1, \widehat{\phi}_j) = \phi_m(1) + x_0$. It follows that $\mathsf{BF}(\boldsymbol{\psi}_j(x_0), n-1, \widehat{\phi}_j) = \phi_m(1) + x_0$ for all links $j \in [m]$. Since for each pair of a user $i \in [n]$ and a link $j \in [m]$,

$$\begin{aligned}\mathsf{IC}_{ij}(\mathbf{F}) &= \mathsf{BF}(\mathbf{f}_{ij}, n-1, \widehat{\phi}_j)\\ &= \mathsf{BF}(\boldsymbol{\psi}_j(x_0), n-1, \widehat{\phi}_j) \quad \text{(by construction of } \mathbf{F})\\ &= \phi_m(1) + x_0\,,\end{aligned}$$

  it follows that $\mathsf{IC}_{ij}(\mathbf{F})$ is constant over all links $j \in [m]$, so that $\mathbf{F}$ is a Nash equilibrium.
The proof is now complete. ■

By the definition of dead and special links, it follows that Theorem 4.10 provides an *efficient* characterization of instances admitting a fully mixed Nash equilibrium.

We now broaden Theorem 4.3 by proving an upper bound on the Social Cost for the case where the fully mixed Nash equilibrium does *not* exist. To state and prove this upper bound, we need first to introduce some simple notation. For a given instance $\langle n, \Phi\rangle$, denote as $\mathcal{SD}$ the set of all special and dead links. Moreover, denote as $\langle n, \Phi \setminus \mathcal{SD}\rangle$ the restriction of the instance $\langle n, \Phi\rangle$ to links outside $\mathcal{SD}$. We prove:

Theorem 4.11 *Consider the case of identical users and arbitrary links with non-constant and convex latency functions. Consider an instance $\langle n, \Phi \rangle$ with an associated Nash equilibrium $\mathbf{P}$, and the instance $\langle n, \Phi \setminus \mathcal{SD} \rangle$ with an associated fully mixed Nash equilibrium $\mathbf{F}$. Then,*

$$\mathsf{SC}^{\Sigma}(n, \Phi, \mathbf{P}) \leq \mathsf{SC}^{\Sigma}(n, \Phi \setminus \mathcal{SD}, \mathbf{F}) \; .$$

PROOF. If there are neither dead nor special links, then $\mathsf{SC}^{\Sigma}(n, \Phi \setminus \mathcal{SD}, \mathbf{F}) = \mathsf{SC}^{\Sigma}(n, \Phi, \mathbf{F})$ and the claim follows from Theorem 4.3. So, assume that there are either dead or special links. Lemma 4.8 implies that no user is assigned (with non-zero probability) by $\mathbf{P}$ to a dead link, while Lemma 4.9 implies that at most one user is assigned (with non-zero probability) by $\mathbf{P}$ to a special link. If no user is assigned (with non-zero probability) by $\mathbf{P}$ to a special link, then $\mathsf{SC}^{\Sigma}(n, \Phi \setminus \mathcal{SD}, \mathbf{F}) = \mathsf{SC}^{\Sigma}(n, \Phi, \mathbf{F})$ and the claim follows again from Theorem 4.3. So, assume that there is a single user assigned (with non-zero probability) by $\mathbf{P}$ to a special link.

Consider any user $i \in [n]$. Note that

$$\begin{aligned}\sum_{j \in [m] \setminus \mathcal{SD}} \theta_{ij}(\mathbf{F}) &= \sum_{j \in [m] \setminus \mathcal{SD}} \sum_{k \in [n] \setminus \{i\}} f(k, j) \\ &= \sum_{k \in [n] \setminus \{i\}} \sum_{j \in [m] \setminus \mathcal{SD}} f(k, j) \\ &= \sum_{k \in [n] \setminus \{i\}} 1 \\ &= n - 1 \, ,\end{aligned}$$

while

$$\begin{aligned}\sum_{j \in [m] \setminus \mathcal{SD}} \theta_{ij}(\mathbf{P}) &= \sum_{j \in [m] \setminus \mathcal{SD}} \sum_{k \in [n] \setminus \{i\}} p(k, j) \\ &= \sum_{k \in [n] \setminus \{i\}} \sum_{j \in [m] \setminus \mathcal{SD}} p(k, j) \\ &\leq \sum_{k \in [n] \setminus \{i\}} 1 \\ &= n - 1 \, .\end{aligned}$$

So, $\sum_{j \in [m] \setminus \mathcal{SD}} \theta_{ij}(\mathbf{F}) \geq \sum_{j \in [m] \setminus \mathcal{SD}} \theta_{ij}(\mathbf{P})$. It follows that there is some link $j_0 \in [m] \setminus \mathcal{SD}$ such that $\theta_{ij_0}(\mathbf{F}) \geq \theta_{ij_0}(\mathbf{P})$, or $\widetilde{f}_{ij_0} \geq \widetilde{p}_{ij_0}$ Hence, we obtain that

$$\begin{aligned}\mathsf{IC}_i(\mathbf{P}) &\leq \mathsf{IC}_{ij_0}(\mathbf{P}) && \text{(since } \mathbf{P} \text{ is a Nash equilibrium)} \\ &= \mathsf{BF}(\mathbf{p}_{ij_0}, n-1, \widehat{\phi}_{j_0}) && \\ &\leq \mathsf{BF}(\widetilde{\mathbf{p}}_{ij_0}, n-1, \widehat{\phi}_{j_0}) && \text{(by Lemma 2.2)} \\ &\leq \mathsf{BF}(\widetilde{\mathbf{f}}_{ij_0}, n-1, \widehat{\phi}_{j_0}) && \text{(by Lemma 2.1 (since } \widetilde{p}_{ij_0} \leq \widetilde{f}_{ij}(\mathbf{F}))) \\ &= \mathsf{BF}(\mathbf{f}_{ij_0}, n-1, \widehat{\phi}_{j_0}) && \text{(by Lemma 4.1)} \\ &= \mathsf{IC}_{ij_0}(\mathbf{F}) && \text{(since } j_0 \in [m] \setminus \mathcal{SD}) \\ &= \mathsf{IC}_i(\mathbf{F}) && \text{(since } \mathbf{F} \text{ is a fully mixed Nash equilibrium)} \; .\end{aligned}$$

Since Social Cost is the sum of Expected Individual Costs, the claim now follows. ■

5. Price of Anarchy

We now present our bounds on the Price of Anarchy for the case of identical users. The case of identical links with a monomial latency function is treated in Section 5.1. The more general case of arbitrary links is treated in Section 5.2. We prepare the reader that Section 5.1 deals with mixed Nash equilibria, while

Section 5.2 deals with pure Nash equilibria; the corresponding bound on the Price of Anarchy in Theorem 5.1 applies to a special case of latency functions (namely, *monomial latency functions*), while the bound on the Price of Anarchy in Corollary 5.4 applies to the more general case of *polynomial latency functions*. However, it turns out the bound in Corollary 5.4 is smaller than the bound in Theorem 5.1.

5.1. *Identical Links with a Monomial Latency Function*

In this section, we assume that there is some integer $d \geq 1$ such that the latency function of each link is the monomial function $\phi(x) = x^d$. We prove:

Theorem 5.1 *Consider the case of identical users and identical links with a monomial latency function* $\phi(x) = x^d$*, for any integer* $d \geq 1$*. Then,* $\mathsf{PoA}^{\Sigma} < B_{d+1}$ *and* $\sup_{m=n\rightarrow\infty} \mathsf{PoA}^{\Sigma} = B_{d+1}$.

PROOF. Since the function $\phi(x) = x^d$ is strictly increasing and convex, Theorems 4.3, 4.10 and 4.6 imply together that the worst-case Nash equilibrium is the unique fully mixed Nash equilibrium. Clearly, the standard fully mixed assignment $\mathbf{F}$ is a Nash equilibrium; so, it is the unique fully mixed Nash equilibrium.

We shall proceed as follows. First, we shall derive a formula for the Social Cost of $\mathbf{F}$; then, we shall use this formula to prove the claim. So,

$$
\begin{aligned}
\mathsf{SC}^{\Sigma}(n, \Phi, \mathbf{F}) &= \sum_{i\in[n]} \sum_{j\in[m]} f(i,j) \sum_{\mathcal{U}\subseteq[n]\setminus\{i\}} \prod_{k\in\mathcal{U}} f(k,j) \prod_{k\notin\mathcal{U}\cup\{i\}} (1-f(k,j)) \cdot \widehat{\phi}(|\mathcal{U}|) \\
&= \sum_{i\in[n]} \sum_{j\in[m]} \sum_{B\subseteq[n]|i\in B} \prod_{k\in B} f(k,j) \cdot \prod_{k\notin B} (1-f(k,j)) \cdot \phi(|B|) \\
&= \sum_{j\in[m]} \sum_{B\subseteq[n]} |B| \cdot \prod_{k\in B} f(k,j) \cdot \prod_{k\notin B} (1-f(k,j)) \cdot \phi(|B|) \\
&= \sum_{j\in[m]} \sum_{k\in[n]} \sum_{B\subseteq[n]||B|=k} |B| \cdot \prod_{t\in B} f(t,j) \cdot \prod_{t\notin B} (1-f(t,j)) \cdot \phi(|B|) \\
&= \sum_{j\in[m]} \sum_{k\in[n]} \binom{n}{k} k \left(\frac{1}{m}\right)^k \cdot \left(1-\frac{1}{m}\right)^{n-k} \cdot k^d \\
&= \sum_{j\in[m]} \sum_{k\in[n]} \binom{n}{k} \left(\frac{1}{m}\right)^k \cdot \left(1-\frac{1}{m}\right)^{n-k} \cdot k^{d+1} \\
&= m \cdot \mathsf{BF}(\frac{1}{m}, n, d+1) \\
&= m \cdot \sum_{k\in[d+1]} S(d+1,k) \cdot n^{\underline{k}} \cdot \left(\frac{1}{m}\right)^k \qquad \text{(by Lemma 2.3)} \\
&= \sum_{k\in[d+1]} S(d+1,k) \cdot \frac{n^{\underline{k}}}{m^{k-1}}
\end{aligned}
$$

It follows that

$$
\frac{\mathsf{SC}^{\Sigma}(n, \Phi, \mathbf{F})}{\mathsf{OPT}^{\Sigma}(n, \Phi)} = \sum_{k\in[d+1]} S(d+1,k) \cdot \frac{n^{\underline{k}}}{m^{k-1}\mathsf{OPT}^{\Sigma}(n, \Phi)} .
$$

We proceed by case analysis on the relation between n and m.

(i) Assume first that $n \geq m$. Then, for each index $k \in [d+1]$,

$$
\begin{aligned}
\frac{n^{\underline{k}}}{m^{k-1}\mathsf{OPT}^{\Sigma}(n,\Phi)} &\leq \frac{n^{\underline{k}}}{m^{k-1}\cdot n\left(\frac{n}{m}\right)^{d}} && \left(\text{since } \mathsf{OPT}^{\Sigma}(n,\Phi) \geq n\left(\tfrac{n}{m}\right)^{d}\right) \\
&= \frac{n^{\underline{k}}}{n^{k}}\cdot\left(\frac{m}{n}\right)^{d+1-k} \\
&< 1 && \left(\text{since } n \geq m \text{ and } n^{\underline{k}} < n^{k}\right)\ ,
\end{aligned}
$$

while also

$$
\lim_{m=n\rightarrow\infty} \frac{n^{\underline{k}}}{m^{k-1}\mathsf{OPT}^{\Sigma}(n,\Phi)} = 1\,.
$$

(ii) Assume now that $n < m$. Then, for each index $k \in [d+1]$,

$$
\begin{aligned}
\frac{n^{\underline{k}}}{m^{k-1}\mathsf{OPT}^{\Sigma}(n,\Phi)} &\leq \frac{n^{\underline{k}}}{m^{k-1}n} && \left(\text{since } \mathsf{OPT}^{\Sigma}(n,\Phi) = n\right) \\
&< \frac{n^{\underline{k}}}{n^{k}} && (\text{since } n < m) \\
&< 1 && \left(\text{since } n^{\underline{k}} < n^{k}\right)\ .
\end{aligned}
$$

It follows that $\mathsf{PoA}^{\Sigma} < \sum_{k\in[d+1]} S(d+1,k) = B_{d+1}$, while $\lim_{m=n\rightarrow\infty} \mathsf{PoA}^{\Sigma} = 1$, as needed. ■

5.2. *Arbitrary Links*

In this section, we consider the case of arbitrary links. We start with a preliminary technical claim.

Proposition 5.2 (Global Optimality = Local Optimality) *For any integer $n \geq 2$, consider the set $X_n = \{\langle x_1, \ldots, x_m\rangle \in \mathbb{N}_0^m \mid \sum_{l\in[m]} x_j = n\}$, and a family of convex functions $\phi_l : [n]\cup\{0\} \rightarrow \mathbb{R}_0^+$, where $l \in [m]$. Then, the vector $\langle x_1, \ldots, x_m\rangle \in X_n$ minimizes the function $\sum_{l\in[m]} \phi_l(x_l)$ over X_n if and only if for all pairs of links $j, k \in [m]$,*

$$
\phi_j(x_j) + \phi_k(x_k) \leq \phi_j(x_j+1) + \phi_k(x_k-1)\,.
$$

It is straightforward to verify that Proposition 5.2 is a particular case of Proposition 3.1. We now prove:

Proposition 5.3 *Consider the case of identical users and arbitrary links. Fix an instance $\langle n, \Phi\rangle$ such that for all links $j \in [m]$ and for all integers $k \in [n]$,*

$$
k\phi_j(k) \leq \alpha \sum_{t\in[k]} \phi_j(t)\,.
$$

Then, for any pure Nash equilibrium $\mathbf{L}$,

$$
\mathsf{SC}^{\Sigma}(n,\Phi,\mathbf{L}) \leq \alpha \cdot \mathsf{OPT}^{\Sigma}(n,\Phi)\,.
$$

PROOF. Fix an optimum pure assignment $\mathbf{G}$ for the instance $\langle n, \Phi\rangle$. For all integers $k \in [n]$ and links $j \in [m]$, define the function

$$
\psi_j(k) = \sum_{t\in[k]} \phi_j(t)\,.
$$

We first show that for each link $j \in [m]$, the function ψ_j is convex. Fix a link $j \in [m]$ and an integer $k \in [n-1]$. Clearly, $\psi_j(k+1) - \psi_j(k) = \phi_j(k+1)$ and $\psi_j(k) - \psi_j(k-1) = \phi_j(k)$. Since the function ϕ_j is non-decreasing, it follows that $\psi_j(k+1) - \psi_j(k) \geq \psi_j(k) - \psi_j(k-1)$, so that ψ_j is convex.

We now show that the vector $\langle\delta_1(\mathbf{L}), \ldots, \delta_m(\mathbf{L})\rangle$ minimizes the function $\sum_{j\in[m]} \psi_j(x_j)$ under the restriction that $\sum_{j\in[m]} x_j = n$. By definition of Nash equilibrium, we have that $\phi_j(\delta_j(\mathbf{L})+1) \geq \phi_k(\delta_k(\mathbf{L}))$ for all pairs of links $j, k \in [m]$. Hence,

$$
\begin{aligned}
&\psi_j(\delta_j(\mathbf{L})+1) + \psi_k(\delta_k(\mathbf{L})-1) \\
&= \psi_j(\delta_j(\mathbf{L})) + \phi_j(\delta_j(\mathbf{L})+1) + \psi_k(\delta_k(\mathbf{L})) - \phi_k(\delta_k(\mathbf{L})) && (\text{by definition of } \psi_j(x) \text{ and } \psi_k(x)) \\
&\geq \psi_j(\delta_j(\mathbf{L})) + \psi_k(\delta_k(\mathbf{L})) && (\text{since } \phi_j(\delta_j(\mathbf{L})+1) \geq \phi_k(\delta_k(\mathbf{L})))\ .
\end{aligned}
$$

By Proposition 5.2, this implies that the pure assignment $\mathbf{L}$ induces loads $\delta_1(\mathbf{L}), \ldots, \delta_m(\mathbf{L})$ which minimize the function $\sum_{j\in[m]} \psi_j(\delta_j(\mathbf{L}))$ under the restriction that $\sum_{l\in[m]} x_l = n$; call $\mathbf{L}$ a *minimizing* assignment. Thus,

$$\begin{aligned}
\mathsf{SC}^{\Sigma}(n,\Phi,\mathbf{L}) &= \sum_{j\in[m]} \delta_j(\mathbf{L})\phi_j(\delta_j(\mathbf{L})) \\
&\leq \sum_{j\in[m]} \alpha \cdot \sum_{t\in[\delta_j(\mathbf{L})]} \phi_j(t) && \text{(by assumption on } \phi\text{)} \\
&= \alpha \sum_{j\in[m]} \psi_j(\delta_j(\mathbf{L})) && \text{(by definition of } \psi_j\text{)} \\
&\leq \alpha \sum_{j\in[m]} \psi_j(\delta_j(\mathbf{G})) && \text{(since } \mathbf{L} \text{ is a minimizing assignment)} \\
&= \alpha \sum_{j\in[m]} \sum_{t\in[\delta_j(\mathbf{G})]} \phi_j(t) && \text{(by definition of } \psi_j\text{)} \\
&\leq \alpha \sum_{j\in[m]} \delta_j(\mathbf{G})\phi_j(\delta_j(\mathbf{G})) && \text{(since } \phi_j \text{ is non-decreasing)} \\
&= \alpha \cdot \mathsf{OPT}^{\Sigma}(n,\Phi)\ , && \text{(since } \mathbf{G} \text{ is an optimum assignment)}\ ,
\end{aligned}$$

as needed. ■

We remark that the proof of Proposition 5.3 is a straightforward adaptation of the corresponding proof for [44, Corollary 2.10] from the continuous setting with splittable flows and continuous latency functions to the discrete setting with unsplittable traffics and discrete latency functions. We conclude this section with a simple application of Proposition 5.3.

Corollary 5.4 *Consider the case of identical users. Assume that all latency functions are (non-zero) polynomials with non-negative coefficients and maximum degree d. Then, for pure Nash equilibria,*

$$\mathsf{PoA}^{\Sigma} \leq d+1\,.$$

PROOF. Consider any latency function $\phi(x) = \sum_{k=0}^{d} a_k x^k$ with $a_k \geq 0$ for all indices $k \in [d] \cup \{0\}$. By Proposition 5.3, it suffices to prove that

$$\frac{x\phi(x)}{\sum_{t\in[x]} \phi(t)} \leq d+1$$

for all integers $x \in [n]$. Clearly,

$$\frac{x\phi(x)}{\sum_{t\in[x]} \phi(t)} = \frac{\sum_{k=0}^{d} a_k x^{k+1}}{\sum_{k=0}^{d} a_k \left(\sum_{t\in[x]} t^k\right)}$$

We shall use the following simple inductive claim:

Lemma 5.5 *For all integers $k \geq 0$ and $x \geq 1$, $\sum_{t\in[x]} t^k \geq \frac{x^{k+1}}{k+1}$.*

So, by Lemma 5.5,

$$\begin{aligned}
\frac{x\phi(x)}{\sum_{t\in[x]} \phi(t)} &\leq \frac{\sum_{k=0}^{d} a_k x^{k+1}}{\sum_{k=0}^{d} \frac{a_k x^{k+1}}{k+1}} \\
&\leq d+1\ ,
\end{aligned}$$

as needed. ■

We remark that Corollary 5.4 is a discrete analog of [44, Corollary 2.11], which held for splittable flows and continuous latency functions.

6. Computing Pure Nash Equilibria and Optimal Assignments

In this section, we provide a fast algorithm to compute a pure Nash equilibrium for the case of identical users and arbitrary links. This algorithm is presented and analyzed in Section 6.1. Section 6.2 establishes that this algorithm can also be used to compute an optimum (pure) assignment for the same case.

6.1. *Pure Nash Equilibria*

A simple approach to compute a pure Nash equilibrium is to assign the (identical) users one by one to their respective best links. This approach is motivated by the classical LPT scheduling algorithm due to Graham [27]; the algorithm had been already employed by Fotakis *et al.* [18, Section 3] for the case of the KP model (where latency functions are linear). The resulting greedy algorithm can be implemented to run in time $O((n+m)\log m)$ if the link latencies are maintained in a priority queue, which is updated after the assignment of each user.

We present an algorithm ComputeNash to compute a pure Nash equilibrium under the assumption of arbitrary, non-decreasing latency functions. We shall establish that the running time of ComputeNash is $O(m \log n \log m)$; this improves on the naive approach if $m = o\left(\frac{n}{\log n}\right)$.

We start with an informal description of the algorithm ComputeNash. A pseudocode description of the algorithm ComputeNash appears in Figure 1. The algorithm takes as input an arbitrary initial pure assignment $\mathbf{L}$ and gives as output a pure Nash equilibrium $\mathbf{L}'$. It does so by moving chunks of users at a time. The first chunk contains all users. In each phase, the *chunk size* is halved until a chunk contains one user only, in which case a Nash equilibrium has been reached. All users in a moved chunk improve their Individual Costs.

Algorithm ComputeNash

Input: an instance $\langle \mathbf{w}, \Phi \rangle$, and an (arbitrary) pure assignment $\mathbf{L}$
Output: a Nash equilibrium $\mathbf{L}'$

begin
$\mathbf{L}' \leftarrow \mathbf{L}$:
    for $j = 1, \ldots, m$ **do**
        $\delta_j(\mathbf{L}') \leftarrow \delta_j(\mathbf{L})$;
for $\delta = n, \lceil \frac{n}{2} \rceil, \lceil \frac{n}{4} \rceil, \ldots, 1$ **do**
    while $\exists s, t \in [m]$ with $\delta_s(\mathbf{L}') \geq \delta$ and $\phi_s(\delta_s(\mathbf{L}')) > \phi_t(\delta_t(\mathbf{L}') + \delta)$ **do**
        choose such $t \in [m]$ so that $\phi_t(\delta_t(\mathbf{L}') + \delta)$ is minimum;
        choose such $s \in [m]$ so that $\phi_s(\delta_s(\mathbf{L}'))$ is maximum;
        transfer δ users from s to t:
            $\delta_s(\mathbf{L}') \leftarrow \delta_s(\mathbf{L}') - \delta$;
            $\delta_t(\mathbf{L}') \leftarrow \delta_t(\mathbf{L}') + \delta$;
return $\mathbf{L}'$;
end

Fig. 1. Algorithm ComputeNash

We prove:

Theorem 6.1 *Consider the case of identical users and arbitrary links. Then,* ComputeNash *computes a pure Nash equilibrium in* $O(m \log m \log n)$ *time.*

PROOF. After the last iteration of the algorithm ComputeNash (with $\delta = 1$), it holds that $\phi_s(\delta_s(\mathbf{L}')) \leq \phi_t(\delta_t(\mathbf{L}') + 1)$ for all pairs of links $s, t \in [m]$. This implies that $\mathbf{L}'$ is a (pure) Nash equilibrium. We continue to analyze the running time of the algorithm ComputeNash. To do so, we first prove an invariant of the algorithm.

Lemma 6.2 (Only Increases or Only Decreases) *During each iteration of the* **for** *loop, the load on a link is either increased or decreased but not both.*

PROOF. By way of contradiction, assume otherwise for some iteration of the **for** loop with chunk size δ. Then, two cases are possible.

(i) The load on some link $t \in [m]$ is decreased after it has been increased.
Consider an increase to the load on link t and the earliest decrease to the load on link t following the increase. By the algorithm, there is some link u whose load is increased simultaneously with the decrease to the load on link t. Denote δ_t and δ_u, and $\widehat{\delta}_t$ and $\widehat{\delta}_u$ the loads on links t and u before the increase (to the load on link t) and before the decrease (to the load on link t), respectively. Note that $\widehat{\delta}_t = \delta_t + \delta$.

For the increase to the load on link t, the choice of t by the algorithm implies that $\phi_t(\delta_t + \delta) \leq \phi_u(\delta_u + \delta)$. For the decrease to the load on link t, the algorithm implies that $\phi_t(\widehat{\delta}_t) > \phi_u(\widehat{\delta}_u + \delta)$ or $\phi_t(\delta_t + \delta) > \phi_u(\widehat{\delta}_u + \delta)$. It follows that $\phi_u(\delta_u + \delta) > \phi_u(\widehat{\delta}_u + \delta)$. Since ϕ_u is non-decreasing, it follows that $\widehat{\delta}_u < \delta_u$. Hence, there is a decrease to the load on link u in between the increase to the load on link t and the increase to the load on link u. Take the latest such decrease to the load on link u. By the algorithm, the load on link u before this decrease is $\widehat{\delta}_u + \delta$. Since there is no change to the load on link t since it has been increased, the load on link t before this decrease is still $\delta_t + \delta$. The choice of u by the algorithm for this decrease (to the load on link u) implies that $\phi_u(\widehat{\delta}_u + \delta) \geq \phi_t(\delta_t + \delta)$. A contradiction.

(ii) The load on some link $s \in [m]$ is increased after it has been decreased.
Consider a decrease to the load on link s and the earliest increase to the load on link s following the decrease. By the algorithm, there is some link u whose load is decreased simultaneously with the increase to the load on link s. Denote δ_s and δ_u, and $\widehat{\delta}_s$ and $\widehat{\delta}_u$ the loads on links s and u before the decrease (to the load on link s) and before the increase (to the load on link s), respectively. Note that $\widehat{\delta}_s = \delta_s - \delta$.

For the decrease to the load on link s, the choice of s by the algorithm implies that $\phi_s(\delta_s) \geq \phi_u(\delta_u)$. For the increase to the load on link s, the algorithm implies that $\phi_s(\widehat{\delta}_s + \delta) < \phi_u(\widehat{\delta}_u)$ or $\phi_s(\delta_s) < \phi_u(\widehat{\delta}_u)$. It follows that $\phi_u(\delta_u) < \phi_u(\widehat{\delta}_u)$. Since ϕ_u is non-decreasing, it follows that $\delta_u < \widehat{\delta}_u$. Hence, there is an increase to the load on link u in between the decrease to the load on link s and the decrease to the load on link u. Take the latest such increase to the load on link u. By the algorithm, the load on link u before this increase is $\widehat{\delta}_u - \delta$. Since there is no change to the load on link s since it has been decreased, the load on link s before this increase is still $\delta_s - \delta$. The choice of u by the algorithm for this increase (to the load on link u) implies that $\phi_u((\widehat{\delta}_u - \delta) + \delta) \leq \phi_s((\delta_s - \delta) + \delta)$ or $\phi_u(\widehat{\delta}_u) \leq \phi_s(\delta_s)$. A contradiction.

Since we derived a contradiction in all possible cases, the proof is now complete. ■

We continue to prove:

Lemma 6.3 *In each iteration of the* **for** *loop, there are at most $O(m)$ iterations of the* **while** *loop.*

PROOF. We consider separately the first iteration of the **for** loop, where $\delta = n$. Consider the first (if any) corresponding iteration of the **while** loop, where the load on some link s is decreased and the load on some link t is increased. For any link l, denote δ_l and $\widehat{\delta}_l$ the loads on link l before and after this iteration of the **while** loop, respectively. By the algorithm, $\delta_s \geq n$; it follows that $\delta_l = 0$ for any link $l \neq s$. By the algorithm, $\widehat{\delta}_t = \delta_t + n = n$: it follows that $\widehat{\delta}_l = 0$ for any link $l \neq t$.

By the algorithm, $\phi_s(\delta_s) > \phi_t(\delta_t + n)$, or $\phi_s(n) > \phi_t(n)$. By the choice of link t by the algorithm, it holds that for any link t' such that $\phi_s(\delta_s) > \phi_{t'}(\delta_{t'} + n)$ or $\phi_s(n) > \phi_{t'}(n)$, either $\phi_t(\delta_t + n) \leq \phi_{t'}(\delta_{t'} + n)$ or $\phi_t(n) \leq \phi_{t'}(n)$.

Assume, by way of contradiction, that a second iteration of the **while** loop is now possible. Since $\widehat{\delta}_t = n$ and $\widehat{\delta}_{t'} = 0$ for any link $t' \neq t$, it follows by the algorithm that there is some link $t' \neq t$ such that

$\phi_t(n) > \phi_{t'}(n)$. Recall also that $\phi_s(n) > \phi_t(n)$. It follows that $\phi_{t'}(n) < \phi_s(n)$. Thus, there is a link t' such that $\phi_{t'}(n) < \phi_s(n)$ for which $\phi_t(n) > \phi_{t'}(n)$. A contradiction. It follows that there is at most one iteration of the **while** loop in the first iteration of the **for** loop.

Consider now any subsequent iteration of the **for** loop with chunk size $\delta < n$. The immediately preceding iteration of the **for** loop has parameter δ' such that $\delta = \lceil \frac{\delta'}{2} \rceil$; clearly, $\delta' \leq 2\delta$. Denote as $\widetilde{\delta}_j$ the load on link $j \in [m]$ upon completion of that iteration of the **for** loop (and immediately before the current iteration with chunk size δ). Partition the set of links $[m]$ into the two sets

$$\mathcal{L}_1 = \{j \in [m] \mid \widetilde{\delta}_j < 2\delta\}$$

and

$$\mathcal{L}_2 = \{j \in [m] \mid \widetilde{\delta}_j \geq 2\delta \,.\}$$

Since each iteration of the **while** loop incurs a simultaneous increase and decrease to the loads on two distinct links, the number of iterations of the **while** loop (in the considered iteration of the **for** loop) is equal to the the total number of decreases to link loads (in the considered iteration of the **for** loop). Hence, we proceed to show:

(1) The total number of decreases (in the considered iteration of the **for** loop) to loads of links in the set $\mathcal{L}_1$ is at most m.

(2) The total number of decreases (in the considered iteration of the **for** loop) to loads of links in the set $\mathcal{L}_2$ is at most m.

Proof of (1): Consider a link $j \in \mathcal{L}_1$ whose load is decreased. Lemma 6.2 implies that its load will further not increase. Since the initial load on link j is less than 2δ and each decrease decreases the load by δ, it follows that the load on link j can be decreased at most once. Hence, the total number of decreases to loads of links in $\mathcal{L}_1$ is at most m.

Proof of (2): We will establish that each link can increase simultaneously with a decrease to the load of any link in $\mathcal{L}_2$ at most once. This will imply that the number of decreases to loads of links in the set $\mathcal{L}_2$ is at most m.

Consider any link $t \in [m]$ and any two consecutive increases to its load. (If there is at most one increase to the load on link t, the claim about the increases to the load on link t holds trivially.) By the algorithm, there is some link s whose load is decreased simultaneously with the second increase to the load on link t. Denote by δ_s and δ_t, and $\widehat{\delta}_s$ and $\widehat{\delta}_t$ the loads on links s and t before the first increase (to the load on link t) and before the second increase (to the load on link t), respectively. Note that $\widehat{\delta}_t = \delta_t + \delta$.

– By Lemma 6.2, there can be no increase to the load on link s in the current iteration of the **for** loop. Hence, $\widehat{\delta}_s \leq \delta_s \leq \widetilde{\delta}_s$.

– By Lemma 6.2, there can be no decrease to the load on link t in the current iteration of the **for** loop. Hence, $\widehat{\delta}_t > \delta_t \geq \widetilde{\delta}_t$.

For the second increase to the load on link t, the algorithm implies that $\phi_s(\widehat{\delta}_s) > \phi_t(\widehat{\delta}_t + \delta)$. Since $\widehat{\delta}_t = \delta_t + \delta$, it follows that $\phi_s(\widehat{\delta}_s) > \phi_t(\delta_t + 2\delta)$.

By the post-condition for the previous iteration of the **for** loop, either $\phi_s(\widetilde{\delta}_s) \leq \phi_t(\widetilde{\delta}_t + \delta')$ or $\widetilde{\delta}_s < \delta'$. We proceed to establish the necessity of the second possibility.

We first prove that $\widetilde{\delta}_s < \delta'$. Assume, by way of contradiction, that $\widetilde{\delta}_s \geq \delta'$. This implies that $\phi_s(\widetilde{\delta}_s) \leq \phi_t(\widetilde{\delta}+\delta')$. Since $\delta_s \leq \widetilde{\delta}_s$ and $\widetilde{\delta}_t \leq \delta_t$, and both ϕ_s and ϕ_t are non-decreasing, it follows that $\phi_s(\delta_s) \leq \phi_t(\delta_t+\delta')$. Since $\delta' \leq 2\delta$ and ϕ_t is non-decreasing, $\phi_t(\delta_t + \delta') \leq \phi_t(\delta_t + 2\delta)$. It follows that $\phi_s(\delta_s) \leq \phi_t(\delta_t + 2\delta)$. Since $\phi_s(\widehat{\delta}_s) > \phi_t(\delta_t + 2\delta)$, it follows that $\phi_s(\delta_s) < \phi_s(\widehat{\delta}_s)$. Since ϕ_s is non-decreasing, this implies that $\delta_s < \widehat{\delta}_s$. A contradiction. It follows that $\widetilde{\delta}_s < \delta'$.

Since $\delta' \leq 2\delta$, this implies that $\widetilde{\delta}_s < 2\delta$. By definition of the set $\mathcal{L}_1$, it follows that $s \in \mathcal{L}_1$. This implies that an increase to the load on a link can occur simultaneously with a decrease to the load of any link in the set $\mathcal{L}_2$ only if this is the first such increase. So, the load on a link can increase simultaneously with a decrease to the load of any link in the set $\mathcal{L}_2$ at most once. It follows that the number of decreases to loads of links in the set $\mathcal{L}_2$ is at most m, as needed.

Claims (1) and (2) imply together that the number of iterations of the **while** loop in each iteration of the **for** loop is $O(m)$, as needed. ■

Each iteration of the **for** loop (with chunk size δ) is implemented using two priority queues. In the first queue, the priorities are according to $\phi_j(\delta_j(\mathbf{L}'))$, $j \in [m]$; in the second queue, the priorities are according to $\phi_j(\delta_j(\mathbf{L}') + \delta)$, $j \in [m]$. At the beginning of the iteration, the two priority queues are constructed in time $O(m \log m)$.

In each iteration of the **while** loop, the links s and t (with maximum $\phi_s(\delta_s(\mathbf{L}'))$ and minimum $\phi_t(\delta_t(\mathbf{L}') + \delta)$, respectively) are determined in constant time using the two priority queues. After each iteration of the **while** loop, which updates $\phi_s(\delta_s(\mathbf{L}'))$ and $\phi_t(\delta_t(\mathbf{L}'))$, the two priority queues are updated in time $O(\log m)$ (by two successive deletion and insertion operations).

By Lemma 6.3, the total time for each iteration of the **for** loop is $O(m \log m) + O(m) \cdot O(\log m) = O(m \log m)$. Since there are $\log n$ iterations of the **for** loop, the total running time of ComputeNash is $O(m \log m \log n)$, as needed. ■

6.2. *Optimal Pure Assignments*

We now establish a relation between optimum pure assignments for a given vector of latency functions and pure Nash equilibria for a modified vector of latency functions. More specifically, given a vector Φ of latency functions, construct the vector Ψ of latency functions by defining for each link $l \in [m]$, the latency function $\psi_l : [n] \to \mathbb{R}$ as

$$\psi_l(x) = x\phi_l(x) - (x-1)\phi_l(x-1)$$

for each $x \in [n]$. We prove:

Proposition 6.4 *Consider the case of identical users. Assume that for each link $l \in [m]$, the function $x\phi_l(x)$ is convex. Then, a pure assignment $\mathbf{L}$ is optimum for the instance $\langle n, \Phi \rangle$ if and only if $\mathbf{L}$ is a Nash equilibrium for the instance $\langle n, \Psi \rangle$.*

PROOF. Clearly, $\mathbf{L}$ is a Nash equilibrium for the instance $\langle n, \Psi \rangle$ if and only if for all pairs of links $j, k \in [m]$,

$$\psi_j(\delta_j(\mathbf{L})) \leq \psi_k(\delta_k(\mathbf{L}) + 1)\,;$$

or, by the definition of the latency function vector Ψ,

$$\begin{aligned}&\delta_j(\mathbf{L})\phi_j(\delta_j(\mathbf{L})) - (\delta_j(\mathbf{L}) - 1)\,\phi_j(\delta_j(\mathbf{L}) - 1)\\ &\leq (\delta_k(\mathbf{L}) + 1)\,\phi_k(\delta_k(\mathbf{L}) + 1) - \delta_k(\mathbf{L})\phi_k(\delta_k(\mathbf{L}))\end{aligned}$$

or

$$\begin{aligned}&\delta_j(\mathbf{L})\phi_j(\delta_j(\mathbf{L})) + \delta_k(\mathbf{L})\phi_k(\delta_k(\mathbf{L}))\\ &\leq (\delta_j(\mathbf{L}) - 1)\,\phi_j(\delta_j(\mathbf{L}) - 1) + (\delta_k(\mathbf{L}) + 1)\,\phi_k(\delta_k(\mathbf{L}) + 1)\,.\end{aligned}$$

By Proposition 5.2, these are necessary and sufficient conditions for the minimization of the sum $\sum_{j \in [m]} \delta_j(\mathbf{L})\phi_j(\delta_j(\mathbf{L}))$ over pure assignments $\mathbf{L}$. Since $\mathsf{SC}^{\Sigma}(\mathbf{w}, \Phi, \mathbf{L}) = \sum_{j \in [m]} \delta_j(\mathbf{L})\phi_j(\delta_j(\mathbf{L}))$, these are as well necessary and sufficient conditions for $\mathbf{L}$ to be an optimum pure assignment for the instance $\langle n, \Phi \rangle$. The proof is now complete. ■

We remark that Proposition 6.4 transfers [44, Corollary 2.7] from the continuous setting of the Wardrop model to discrete routing games. Proposition 6.4 immediately implies:

Corollary 6.5 *Consider the case of identical users. Assume that for each link $l \in [m]$, the function $x\phi_l(x)$ is convex. Then, an optimum pure assignment can be computed in time $O(m \log m \log n)$.*

7. Computing Best and Worst Pure Nash Equilibria

In this section, we present some complexity results for the computation of best and worst pure Nash equilibria. More specifically, we will consider the following two decision problems, which are natural decision versions of corresponding optimization problems defined in [18] for the KP model:

BEST PURE NE

INSTANCE: An instance $\langle \mathbf{w}, \Phi \rangle$ and a rational number $B > 0$.

QUESTION: Is there a pure Nash equilibrium $\mathbf{L}$ with $\mathsf{SC}^{\Sigma}(\mathbf{w}, \Phi, \mathbf{L}) \leq B$?

WORST PURE NE

INSTANCE: An instance $\langle \mathbf{w}, \Phi \rangle$ and a rational number $B > 0$.

QUESTION: Is there a pure Nash equilibrium $\mathbf{L}$ with $\mathsf{SC}^{\Sigma}(\mathbf{w}, \Phi, \mathbf{L}) \geq B$?

We will prove that both these problems are $\mathcal{NP}$-complete even for the case of identical links with an identity latency function. The proofs will use polynomial time transformations from the original $\mathcal{NP}$-complete PARTITION problem [29] or its slight variant RESTRICTED PARTITION that we define below.

PARTITION

INSTANCE: A finite set A of items with $|A| \geq 2$, and a size $s(a) \in \mathbb{N}$ for each item $a \in A$.

QUESTION: Is there a subset $A' \subseteq A$ such that $\sum_{a \in A'} s(a) = \sum_{a \in A \setminus A'} s(a)$?

RESTRICTED PARTITION

INSTANCE: A finite set A of items with $|A| \geq 12$, and a size $s(a) \in \mathbb{N}$ for each item $a \in A$ such that $s(a) \leq \frac{1}{8} \sum_{a' \in A} s(a')$.

QUESTION: Is there a subset $A' \subseteq A$ such that $\sum_{a \in A'} s(a) = \sum_{a \in A \setminus A'} s(a)$?

Clearly, PARTITION reduces trivially to RESTRICTED PARTITION by padding the finite set A in the instance of PARTITION with ten new items a' with $s(a') = \sum_{a \in A} s(a)$. It is easy to see that this trivial reduction is parsimonious; since #PARTITION (i.e. the counting version of PARTITION) is $\#\mathcal{P}$-complete [45], it follows that #RESTRICTED PARTITION (i.e. the counting version of RESTRICTED PARTITION) is also $\#\mathcal{P}$-complete.

We shall also consider the counting versions #BEST PURE NE and #WORST PURE NE of BEST PURE NE and WORST PURE NE, respectively.

We start by proving:

Theorem 7.1 *Consider the case of identical links. Then,* BEST PURE NE *is $\mathcal{NP}$-complete.*

PROOF. Clearly, BEST PURE NE $\in \mathcal{NP}$. To prove $\mathcal{NP}$-hardness, we employ a polynomial time transformation from RESTRICTED PARTITION to BEST PURE NE. Given an instance of RESTRICTED PARTITION, we construct an instance $\langle \mathbf{w}, \langle m, \phi \rangle \rangle$ of BEST PURE NE as follows:

– There are $n = |A| + 2$ users with

$$w_i = \begin{cases} s(a_i), & 1 \leq i \leq |A| \\ \dfrac{\sum_{a \in A} s(a)}{2}, & i \in \{|A|+1, |A|+2\} \end{cases}$$

– There are three identical links with identity latency function $\phi(x) = x$.

– $B = \left(\frac{|A|}{2} + 2\right) \sum_{a \in A} s(a)$.

Clearly, this is a polynomial time mapping. We prove that it is a transformation from RESTRICTED PARTITION to BEST PURE NE.

Assume first that the instance of RESTRICTED PARTITION is positive, and consider a subset $A' \subseteq A$ such that $\sum_{a \in A'} s(a) = \sum_{a \in A \setminus A'} s(a)$. Use A' to define a pure assignment $\mathbf{L}$ for the constructed instance of BEST PURE NE as follows:

– For each item $a_i \in A'$, user i is assigned to link 1; for each item $a_i \in A \setminus A'$, user i is assigned to link 2.
– Users $|A| + 1$ and $|A| + 2$ are assigned to link 3.

We now prove that $\mathbf{L}$ is a (pure) Nash equilibrium for the instance $\langle \mathbf{w}, \langle m, \phi \rangle \rangle$ of BEST PURE NE with $\mathsf{SC}(\mathbf{w}, \langle m, \phi \rangle, \mathbf{L}) \leq B$. Clearly,

$$\begin{aligned} \delta_1(\mathbf{L}) &= \sum_{a_i \in A'} w_i && \text{(by definition of } \mathbf{L}) \\ &= \sum_{a_i \in A'} s(a_i) && \text{(by the mapping)} \\ &= \frac{\sum_{a \in A} s(a)}{2}, \end{aligned}$$

and similarly $\delta_2(\mathbf{L}) = \frac{\sum_{a \in A} s(a)}{2}$. On the other hand,

$$\begin{aligned} \delta_3(\mathbf{L}) &= \sum_{i \in \{|A|+1, |A|+2\}} w_i && \text{(by the definition of } \mathbf{L}) \\ &= \sum_{i \in \{|A|+1, |A|+2\}} \frac{\sum_{a \in A} s(a)}{2} && \text{(by the mapping)} \\ &= \sum_{a \in A} s(a). \end{aligned}$$

Note that each user $i \in [|A|]$ is assigned to either link 1 or link 2; thus, $\mathsf{IC}_i(\mathbf{L}) = \delta_1(\mathbf{L}) = \delta_2(\mathbf{L})$. So, user i is satisfied in $\mathbf{L}$. Note also that $\mathsf{IC}_{|A|+1}(\mathbf{L}) = \delta_3(\mathbf{L}) = \delta_1(\mathbf{L}) + w_{|A|+1} = \delta_2(\mathbf{L}) + w_{|A|+1}$ and similarly $\mathsf{IC}_{|A|+2}(\mathbf{L}) = \delta_3(\mathbf{L}) = \delta_1(\mathbf{L}) + w_{|A|+2} = \delta_2(\mathbf{L}) + w_{|A|+2}$; so users $|A| + 1$ and $|A| + 2$ are also satisfied in $\mathbf{L}$, and $\mathbf{L}$ is a Nash equilibrium. Clearly,

$$\begin{aligned} \mathsf{SC}^{\Sigma}(\mathbf{w}, \langle m, \phi \rangle, \mathbf{L}) &= |A| \cdot \delta_1(\mathbf{L}) + 2 \cdot \delta_3(\mathbf{L}) \\ &= |A| \cdot \frac{\sum_{a \in A} s(a)}{2} + 2 \cdot \sum_{a \in A} s(a) \\ &= \left(\frac{|A|}{2} + 2\right) \cdot \sum_{a \in A} s(a) \\ &= B. \end{aligned}$$

So, the instance $\langle \mathbf{w}, \langle m, \phi \rangle \rangle$ of BEST PURE NE is also positive.

Assume now that the instance of RESTRICTED PARTITION is negative. So, for every subset $A' \subseteq A$, $\sum_{a \in A'} s(a) \neq \sum_{a \in A \setminus A'} s(a)$. It follows that for every subset $A' \subseteq A$, either $\sum_{a \in A'} s(a) < \frac{\sum_{a \in A} s(a)}{2}$ or $\sum_{a \in A \setminus A'} s(a) < \frac{\sum_{a \in A} s(a)}{2}$ (but not both).

We now prove that the instance $\langle \mathbf{w}, \langle m, \phi \rangle \rangle$ of BEST PURE NE is also negative. Consider any arbitrary (pure) Nash equilibrium $\mathbf{L}$ for the instance $\langle \mathbf{w}, \langle m, \phi \rangle \rangle$. We will prove that $\mathsf{SC}^{\Sigma}(\mathbf{w}, \langle m, \phi \rangle, \mathbf{L}) > B$. We first show a preliminary property of $\mathbf{L}$.

Assume, by way of contradiction, that users $|A|+1$ and $|A|+2$ were assigned to the same link in $\mathbf{L}$. Say that link were link 3. Then, $\mathsf{IC}_{|A|+1}(\mathbf{L}) \geq w_{|A|+1} + w_{|A|+2} = \sum_{a \in A} s(a)$. Denote as A' the set of users assigned to link 1; then, $\delta_1(\mathbf{L}) = \sum_{a \in A'} s(a)$. Assume, without loss of generality, that $\sum_{a \in A'} s(a) < \frac{\sum_{a \in A} s(a)}{2}$. Then,

$$
\begin{aligned}
\mathsf{IC}_{(|A|+1)1}(\mathbf{L}) &= \sum_{a \in A'} s(a) + w_{|A|+1} \\
&= \sum_{a \in A'} s(a) + \frac{\sum_{a \in A} s(a)}{2} && \text{(by the mapping)} \\
&< \frac{\sum_{a \in A} s(a)}{2} + \frac{\sum_{a \in A} s(a)}{2} && \text{(by assumption)} \\
&= \sum_{a \in A} s(a) \\
&\leq \mathsf{IC}_{|A|+1}(\mathbf{L}),
\end{aligned}
$$

so that user $|A|+1$ is not satisfied in $\mathbf{L}$; this contradicts the fact that $\mathbf{L}$ is a Nash equilibrium. It follows that users $|A|+1$ and $|A|+2$ are assigned to different links in $\mathbf{L}$.

Assume, without loss of generality, that $\delta_1(\mathbf{L}) \geq \delta_2(\mathbf{L}) \geq \delta_3(\mathbf{L})$. Clearly,

$$
\delta_1(\mathbf{L}) + \delta_2(\mathbf{L}) + \delta_3(\mathbf{L}) = \sum_{i \in [n]} w_i = 2 \sum_{a \in A} s(a),
$$

which implies that

$$
\delta_1(\mathbf{L}) \geq \frac{2}{3} \sum_{a \in A} s(a).
$$

Now, since $w_{|A|+1} = w_{|A|+2} = \frac{\sum_{a \in A} s(a)}{2}$ and users $|A|+1$ and $|A|+2$ are assigned to different links, it follows that there is a user $i \in [|A|]$ assigned to link 1; by definition of RESTRICTED PARTITION, it follows that $w_i \leq \frac{1}{8} \sum_{a \in A} s(a)$. Since $\mathbf{L}$ is a Nash equilibrium, user i is satisfied in $\mathbf{L}$, so that

$$
\begin{aligned}
\mathsf{IC}_i(\mathbf{L}) &= \delta_1(\mathbf{L}) \\
&\leq \delta_3(\mathbf{L}) + w_i \\
&\leq \delta_3(\mathbf{L}) + \frac{1}{8} \sum_{a \in A} s(a).
\end{aligned}
$$

Hence,

$$
\begin{aligned}
\delta_3(\mathbf{L}) &\geq \delta_1(\mathbf{L}) - \frac{1}{8} \sum_{a \in A} s(a) \\
&= \left(2 - \frac{1}{8}\right) \sum_{a \in A} s(a) - \delta_2(\mathbf{L}) - \delta_3(\mathbf{L}) && \left(\text{since } \delta_1(\mathbf{L}) + \delta_2(\mathbf{L}) + \delta_3(\mathbf{L}) = 2 \textstyle\sum_{a \in A} s(a)\right) \\
&\geq \frac{15}{8} \sum_{a \in A} s(a) - \delta_1(\mathbf{L}) - \delta_3(\mathbf{L}) && (\text{since } \delta_1(\mathbf{L}) \geq \delta_2(\mathbf{L})) \\
&\geq \frac{15}{8} \sum_{a \in A} s(a) - \delta_3(\mathbf{L}) - \frac{1}{8} \sum_{a \in A} s(a) - \delta_3(\mathbf{L}) && \left(\text{since } \delta_1(\mathbf{L}) \leq \delta_3(\mathbf{L}) + \tfrac{1}{8} \textstyle\sum_{a \in A} s(a)\right),
\end{aligned}
$$

which implies that

$$
\delta_3(\mathbf{L}) \geq \frac{7}{12} \sum_{a \in A} s(a).
$$

Since there are $|A|+2$ users and $\delta_1(\mathbf{L}) \geq \delta_2(\mathbf{L}) \geq \delta_3(\mathbf{L})$, it follows from the definition of Social Cost that

$$
\begin{aligned}
\mathsf{SC}^{\Sigma}(\mathbf{w}, \langle m, \phi\rangle, \mathbf{L}) &\geq (|A|+2)\cdot \delta_3(\mathbf{L}) \\
&\geq (|A|+2)\cdot \frac{7}{12}\sum_{a\in A} s(a) \\
&> \left(\frac{|A|}{2}+2\right)\cdot \sum_{a\in A} s(a) \qquad (\text{since } |A|\geq 12) \\
&= B\,.
\end{aligned}
$$

So, the instance $\langle \mathbf{w}, \langle m, \phi\rangle\rangle$ is also negative. This completes the proof. ■

Note that the reduction from RESTRICTED PARTITION employed in the proof of Theorem 7.1 is parsimonious; since #RESTRICTED PARTITION is $\#\mathcal{P}$-complete, it immediately follows:

Corollary 7.2 *Consider the case of identical links. Then* #BEST PURE NE *is* $\#\mathcal{P}$*-complete.*

We continue to prove:

Theorem 7.3 *Consider the case of identical links. Then,* WORST PURE NE *is* $\mathcal{NP}$*-complete.*

PROOF. Clearly, WORST PURE NE $\in \mathcal{NP}$. To prove $\mathcal{NP}$-hardness, we employ a polynomial time transformation from PARTITION to WORST PURE NE. Given an instance of PARTITION, we construct an instance $\langle \mathbf{w}, \langle m, \phi\rangle\rangle$ of WORST PURE NE as follows:

– There are $3|A|$ users with

$$
w_i = \begin{cases} s(a_i),\ i\in [|A|] \\ \dfrac{1}{4|A|},\ |A|+1\leq i\leq 3|A| \end{cases}
$$

– There are two identical links with identity latency function $\phi(x) = x$.

– $B = 3|A|\left(\frac{\sum_{a\in A} s(a)}{2} + \frac{1}{4}\right)$.

Clearly, this is a polynomial time mapping. We prove that it is a transformation from PARTITION to WORST PURE NE.

Assume first that the instance of PARTITION is positive, and consider a subset $A' \subseteq A$ such that $\sum_{a\in A'} s(a) = \sum_{a\in A\setminus A'} s(a)$. Use A' to define a pure assignment $\mathbf{L}$ for the constructed instance $\langle \mathbf{w}, \langle m, \phi\rangle\rangle$ of WORST PURE NE as follows:

– For each item $a_i \in A'$, user i is assigned to link 1; for each item $a_i \in A\setminus A'$, user i is assigned to link 2.

– Each user i with $|A|+1 \leq i \leq 2|A|$ is assigned to link 1; each user i with $2|A|+1 \leq i \leq 3|A|$ is assigned to link 2.

We now prove that $\mathbf{L}$ is a Nash equilibrium for the instance $\langle \mathbf{w}, \langle m, \phi\rangle\rangle$ of WORST PURE NE with $\mathsf{SC}^{\Sigma}(\mathbf{w}, \langle m, \phi\rangle, \mathbf{L}) \geq B$. Clearly,

$$
\begin{aligned}
\delta_1(\mathbf{L}) &= \sum_{a\in A'} s(a) + \sum_{|A|+1\leq i\leq 2|A|} \frac{1}{4|A|} \quad (\text{by definition of } \mathbf{L} \text{ and the mapping}) \\
&= \sum_{a\in A'} s(a) + \frac{1}{4} \\
&= \frac{\sum_{a\in A} s(a)}{2} + \frac{1}{4} \qquad (\text{by choice of } A')\ ,
\end{aligned}
$$

and similarly

$$
\delta_2(\mathbf{L}) = \frac{\sum_{a\in A} s(a)}{2} + \frac{1}{4}\,.
$$

Since $\delta_1(\mathbf{L}) = \delta_2(\mathbf{L})$, all users are satisfied in $\mathbf{L}$, and $\mathbf{L}$ is a Nash equilibrium. Since $3|A|$ users are assigned to links and $\delta_1(\mathbf{L}) = \delta_2(\mathbf{L})$, it follows that

$$\begin{aligned}\mathsf{SC}^{\Sigma}(\mathbf{w}, \langle m, \phi\rangle, \mathbf{L}) &= 3|A| \cdot \delta_1(\mathbf{L}) \\ &= 3|A| \cdot \left(\frac{\sum_{a\in A} s(a)}{2} + \frac{1}{4} \right) \\ &= B\,.\end{aligned}$$

So, the instance $\langle \mathbf{w}, \langle m, \phi\rangle\rangle$ of WORST PURE NE is also positive.

Assume now that the instance of PARTITION is negative. So, for every subset $A' \subseteq A$, $\sum_{a\in A'} s(a) \neq \sum_{a\in A\setminus A'} s(a)$. It follows that for every subset $A' \subseteq A$, either $\sum_{a\in A'} s(a) < \frac{\sum_{a\in A} s(a)}{2}$ or $\sum_{a\in A\setminus A'} s(a) < \frac{\sum_{a\in A} s(a)}{2}$ (but not both).

We now prove that the instance $\langle \mathbf{w}, \langle m, \phi\rangle\rangle$ of WORST PURE NE is also negative. Consider any arbitrary pure Nash equilibrium $\mathbf{L}$ for the instance $\langle \mathbf{w}, \langle m, \phi\rangle\rangle$. We will prove that $\mathsf{SC}^{\Sigma}(\mathbf{w}, \langle m, \phi\rangle, \mathbf{L}) < B$.

Denote as A' the set of users assigned to link 1. So, $A \setminus A'$ is the set of users assigned to link 2. Assume, without loss of generality, that $\sum_{a\in A'} s(a) < \frac{\sum_{a\in A} s(a)}{2}$. Then, $\sum_{a\in A\setminus A'} s(a) > \frac{\sum_{a\in A} s(a)}{2}$, so that $\sum_{a\in A\setminus A'} s(a) - \sum_{a\in A'} s(a) \geq 1$. We first show a preliminary property of $\mathbf{L}$.

Assume, by way of contradiction, that some user i with $|A|+1 \leq i \leq 3|A|$ is assigned to link 2. Then,

$$\begin{aligned}\mathsf{IC}_i(\mathbf{L}) &= \delta_2(\mathbf{L}) \\ &\geq \sum_{a\in A\setminus A'} s(a) + w_i \\ &\geq \sum_{a\in A'} s(a) + 1 + w_i \\ &> \sum_{a\in A'} s(a) + \sum_{|A|+1\leq k\leq 3|A|} s(a) + w_i \quad \left(\text{since } \textstyle\sum_{|A|+1\leq k\leq 3|A|} s(a) = \frac{1}{2}\right) \\ &\geq \delta_1(\mathbf{L}) + w_i \\ &= \mathsf{IC}_{i1}(\mathbf{L})\,,\end{aligned}$$

which implies that user i is not satisfied in $\mathbf{L}$; this contradicts the fact that $\mathbf{L}$ is a Nash equilibrium. It follows that all users i with $|A|+1 \leq i \leq 3|A|$ are assigned to link 1 in $\mathbf{L}$, so that

$$\begin{aligned}\delta_1(\mathbf{L}) &= \sum_{a\in A'} s(a) + 2|A|\,\frac{1}{4|A|} \quad \text{(by the mapping)} \\ &= \sum_{a\in A'} s(a) + \frac{1}{2}\,.\end{aligned}$$

For each link $j \in [2]$, denote as n_j the number of users i with $a_i \in A$ that are assigned to link j in $\mathbf{L}$. Clearly, $\sum_{j\in[2]} n_j = |A|$. So, it follows from the definition of Social Cost that

$$\begin{aligned}&\mathsf{SC}(\mathbf{w}, \langle m, \phi\rangle, \mathbf{L}) \\ &= (n_1 + 2|A|)\delta_1(\mathbf{L}) + n_2\delta_2(\mathbf{L}) \\ &= (n_1 + 2|A|)\left(\sum_{a\in A'} s(a) + \frac{1}{2}\right) + n_2\left(\sum_{a\in A\setminus A'} s(a)\right) \\ &< 2\,|A|\left(\sum_{a\in A'} s(a) + \frac{1}{2}\right) + (n_1+n_2)\left(\sum_{a\in A\setminus A'} s(a)\right) \quad \left(\text{since } \textstyle\sum_{a\in A'} s(a) + \frac{1}{2} < \sum_{a\in A\setminus A'} s(a)\right) \\ &= 2\,|A|\left(\sum_{a\in A'} s(a) + \frac{1}{2}\right) + |A|\left(\sum_{a\in A\setminus A'} s(a)\right)\end{aligned}$$

$$
\begin{aligned}
&= |A| \left(\sum_{a \in A'} s(a) + \sum_{a \in A \setminus A'} s(a) \right) + |A| \sum_{a \in A'} s(a) + |A| \\
&= |A| \sum_{a \in A} s(a) + |A| \sum_{a \in A'} s(a) + |A| \\
&\leq |A| \sum_{a \in A} s(a) + |A| \left(\frac{\sum_{a \in A} s(a)}{2} - \frac{1}{2} \right) + |A| \qquad \left(\text{since } \textstyle\sum_{a \in A'} s(a) \leq \frac{\sum_{a \in A} s(a)}{2} - \frac{1}{2}\right) \\
&= 3|A| \left(\frac{\sum_{a \in A} s(a)}{2} + \frac{1}{6} \right) \\
&< B\,.
\end{aligned}
$$

So, the instance $\langle \mathbf{w}, \langle m, \phi \rangle \rangle$ of WORST PURE NE is also negative. This completes the proof. ■

Note that the reduction from PARTITION employed in the proof of Theorem 7.3 is parsimonious; since #PARTITION is $\#\mathcal{P}$-complete [45], it immediately follows:

Corollary 7.4 *Consider the case of identical links. Then* #WORST PURE NE *is* $\#\mathcal{P}$*-complete.*

8. Epilogue

We have introduced *discrete routing games* combining features from two of the most prominent models for non-cooperative routing, namely the KP model [31] and the Wardrop model [48].

- We presented a thorough analysis of fully mixed Nash equilibria for discrete routing games. In particular, we proved that, for the case of identical users, the Social Cost of any Nash equilibrium is bounded by the Social Cost of the fully mixed Nash equilibrium. Moreover, we derived a characterization of instances admitting a fully mixed Nash equilibrium, and we proved that a fully mixed Nash equilibrium may exist only uniquely.
- We presented upper bounds on the Price of Anarchy for the case of identical users.
- Still for the case of identical users, we showed that a pure Nash equilibrium can be computed efficiently. For the case of arbitrary users, we proved that computing the best or the worst pure Nash equilibrium is already $\mathcal{NP}$-complete even for identical links with an identity latency function.

We conclude with a collection of interesting open problems about discrete routing games.

(i) Extend the results about uniqueness and existence of the fully mixed Nash equilibrium (namely, Theorems 4.6 and 4.10) from the case of identical users to the case of arbitrary users. In particular, what are the analogs of dead links and special links for such extension?

(ii) Prove or disprove the FMNE Conjecture for discrete routing games.

(iii) Obtain bounds on the Price of Anarchy for the case of arbitrary users and identical links with a monomial latency function. (This will extend Theorem 5.1.)

(iv) Obtain bounds on the Price of Anarchy for the case of arbitrary users and arbitrary links with a polynomial latency function. (This will extend Corollary 5.4.)

(v) Is the fast algorithm we presented in Section 6 to compute a pure Nash equilibrium (for the case case of identical users and arbitrary links) optimal? We note that there are no known lower bounds for this problem.

(vi) Is there a PTAS for either BEST PURE NE or WORST PURE NE? We know that there is a PTAS to compute a best pure Nash equilibrium for the KP model [14]. (This PTAS employs known approximation algorithms to compute an optimum pure assignment [28]; the so called *Nashification* technique is then applied on that to transform it in polynomial time into a pure Nash equilibrium with no increased Social Cost.)

Acknowledgments.

We thank Marina Gelastou and Andreas Pieris for bringing to our attention an inaccuracy in an earlier statement and proof of Lemma 6.2. We also thank Rainer Feldmann for some very helpful discussions on discrete convexity.

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] 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.

[3] B. Awerbuch, Y. Azar, Y. Richter and D. Tsur, "Tradeoffs in Worst-Case Equilibria," *Theoretical Computer Science,* Vol. 361, Nos. 2–3, pp. 200–209, September 2006.

[4] M. Beckmann, C. B. McGuire and C. B. Winsten, *Studies in the Economics of Transportation.* Yale University Press, 1956.

[5] M. J. Beckmann, "On the Theory of Traffic Flow in Networks," *Traffic Quarterly,* Vol. 21, pp. 109–116, 1967.

[6] E. T. Bell, "Exponential Numbers," *American Mathematical Monthly,* Vol. 41, No. 7, pp. 411–419, August/September 1934.

[7] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos and L. Moscardelli, "Tight Bounds for Selfish and Greedy Load Balancing," *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming,* M. Bugliesi, B. Breneel, V. Sassone and I. Wegener eds., pp. 311–322, Vol. 4051, Lecture Notes in Computer Science, Springer-Verlag, July 2006.

[8] 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.

[9] V. Conitzer and T. Sandholm, "New Complexity Results about Nash Equilibria," *Games and Economic Behavior,* Vol. 1, No. 2, pp. 621–641, July 2008.

[10] 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.

[11] A. Czumaj and B. Vöcking, "Tight Bounds for Worst-Case Equilibria," *ACM Transactions on Algorithms,* Vol. 3, No. 1, 2007.

[12] S. C. Dafermos and F. T. Sparrow, "The Traffic Assignment Problem for a General Network," *Journal of Research of the National Bureau of Standards - B. Mathematical Sciences,* Vol. 73B, No. 2, pp. 291–118, April/May/June 1969.

[13] E. Even-Dar, A. Kesselman and Y. Mansour, "Convergence Time to Nash Equilibria," *ACM Transactions on Algorithms,* Vol. 3, No. 3, 2007.

[14] R. Feldmann, M. Gairing, T. Lücking, B. Monien and M. Rode, "Nashification and the Price of Anarchy 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.

[15] R. Feldmann, M. Gairing, T. Lücking, B. Monien and M. Rode, "Selfish Routing in Non-Cooperative Networks: A Survey," *Current Trends in Theoretical Computer Science: The Challenge of the New Century*, Vol. 1, pp. 373–401, World Scientific Publishing, 2004.

[16] R. Feldmann, M. Mavronicolas and A. Pieris, "Facets of the Fully Mixed Nash Equilibrium Conjecture," *Proceedings of the 1st International Symposium on Algorithmic Game Theory*, B. Monien and U. P. Schroeder eds., pp. 145–157, Vol. 4997, Lecture Notes in Computer Science, Springer-Verlag, April/May 2008.

[17] S. Fischer and B. Vöcking, "On the Structure and Complexity of Worst-Case Equilibria," *Theoretical Computer Science,* Vol. 378, No. 2, pp. 165–174, June 2007.

[18] 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. Extended version accepted to *Theoretical Computer Science.*

[19] D. Fotakis, S. Kontogiannis and P. Spirakis, "Symmetry in Network Congestion Games: Pure Equilibria and Anarchy Cost," *Proceedings of the 3rd International Workshop on Approximation and Online Algorithms,* T. Erlebach and G. Persiano eds., pp. 161–175, Vol. 3879, Lecture Notes in Computer Science, Springer-Verlag, October 2005.

[20] D. Fotakis, S. Kontogiannis and P. Spirakis, "Selfish Unsplittable Flow," *Theoretical Computer Science,* Vol. 348, Nos. 2–3, pp. 226–239, December 2005.

[21] M. Gairing, T. Lücking, M. Mavronicolas and B. Monien, "Computing Nash Equilibria for Scheduling on Restricted Parallel Links," *Proceedings of the 36th Annual ACM Symposium on Theory of Computing,* pp. 613–622, June 2004.

[22] M. Gairing, T. Lücking, M. Mavronicolas and B. Monien, "The Price of Anarchy for Polynomial Social Cost," *Theoretical Computer Science,* Vol. 369, No. 1-3, pp. 116–135, December 2006.

[23] M. Gairing, T. Lücking, M. Mavronicolas and B. Monien, "The Price of Anarchy for Restricted Parallel Links," *Parallel Processing Letters,* Vol. 16, No. 1, pp. 117–131, March 2006.

[24] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis, "Structure and Complexity of Extreme Nash Equilibria," *Theoretical Computer Science,* Vol. 343, Nos. 1 & 2, pp. 133–157, October 2005.

[25] M. Gairing, T. Lücking, B. Monien and K. Tiemann, "Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture," *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming,* L. Caires, G. F. Italiano, L. Monteiro, C. Palamidessi and M. Yung eds., pp. 51–65, Vol. 3580, Lecture Notes in Computer Science, Springer-Verlag, July 2005.

[26] M. Gairing, B. Monien and K. Tiemann, "Selfish Routing with Incomplete Information," *Theory of Computing Systems,* Vol. 42, No. 1, pp. 91–130, January 2008.

[27] R. L. Graham, "Bounds on Multiprocessing Timing Anomalies," *SIAM Journal of Applied Mathematics,* Vol. 17, No. 2, pp. 416–429, March 1969.

[28] D. S. Hochbaum and D. Shmoys, "A Polynomial Approximation Scheme for Scheduling on Uniform Processors: Using the Dual Approximation Approach," *SIAM Journal on Computing,* Vol. 17, No. 3, pp. 539–551, June 1988.

[29] R. M. Karp, "Reducibility Among Combinatorial Problems," *Complexity of Computer Computations,* R. E. Miller and J. W. Thatcher eds., Plenum Press, pp. 85–103, 1972.

[30] 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.

[31] E. Koutsoupias and C. 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.

[32] D. C. Kozen, *The Design and Analysis of Algorithms,* Springer-Verlag, 1992.

[33] 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. Also, accepted to *Theoretical Computer Science.*

[34] 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.

[35] M. Mavronicolas, P. Panagopoulou and P. Spirakis, "Cost Sharing Mechanisms for Fair Pricing of Resource Usage," *Proceedings of the 1st International Workshop on Internet and Network Economics,* X. Deng and Y. Ye eds., pp. 210–224, Vol. 3828, Lecture Notes in Computer Science, Springer-Verlag, December 2005. Also, accepted to *Algorithmica*.

[36] M. Mavronicolas and P. Spirakis, "The Price of Selfish Routing," *Algorithmica,* Vol. 48, No. 1, pp. 91–126, 2007.

[37] I. Milchtaich, "Congestion Games with Player-Specific Payoff Functions," *Games and Economic Behavior,* Vol. 13, No. 1, pp. 111–124, March 1996.

[38] K. Murota, "Submodular Flow Problem with Nonseparable Cost Function," *Combinatorica,* Vol. 19, No. 1, pp. 87–109, January 1999.

[39] 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.

[40] J. F. Nash, "Non-Cooperative Games," *Annals of Mathematics,* Vol. 54, No. 2, pp. 286–295, 1951.

[41] C. H. Papadimitriou, "Algorithms, Games and the Internet," *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing,* pp. 749–753, July 2001.

[42] R. W. Rosenthal, "A Class of Games Possessing Pure-Strategy Nash Equilibria," *International Journal of Game Theory,* Vol. 2, No. 1, pp. 65–67, December 1973.

[43] T. Roughgarden, *Selfish Routing and the Price of Anarchy,* The MIT Press, 2005.

[44] T. Roughgarden and É. Tardos, "How Bad is Selfish Routing?," *Journal of the ACM,* Vol. 49, No. 2, pp. 236–259, May 2002.

[45] J. Simon, "On the Difference between the One and the Many (Preliminary Version)", *Proceedings of the 4th International Colloquium on Automata, Languages and Programming,* pp. 480–491, Vol. 52, Lecture Notes in Computer Science, Springer-Verlag, August 1977.

[46] J. Stirling, *Methodus Differentialis,* London, 1730.

[47] S. Suri, C. D. Tóth and Y. Zhou, "Selfish Load Balancing and Atomic Congestion Games," *Algorithmica,* Vol. 47, No. 1, pp. 79–96, 2007.

[48] 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.