

# Nash Equilibria in Discrete Routing Games with Convex Latency Functions <sup>★</sup>

Martin Gairing<sup>1</sup>, Thomas Lücking<sup>1</sup>, Marios Mavronicolas<sup>2</sup>, Burkhard Monien<sup>1</sup>,  
and Manuel Rode<sup>1\*\*</sup>

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

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

<sup>2</sup> Department of Computer Science, University of Cyprus,  
P. O. Box 20537, Nicosia CY-1678, Cyprus. Email: mavronic@ucy.ac.cy

**Abstract.** We study *Nash equilibria* in a discrete routing game that combines features of the two most famous models for non-cooperative routing, the *KP model* [16] and the *Wardrop model* [27]. In our model, users share parallel links. A *user strategy* can be any probability distribution over the set of links. Each user tries to minimize its *expected latency*, where the latency on a link is described by an arbitrary non-decreasing, convex function. The social cost is defined as the sum of the users' expected latencies. To the best of our knowledge, this is the first time that *mixed* Nash equilibria for routing games have been studied in combination with *non-linear* latency functions.

As our main result, we show that for identical users the social cost of any Nash equilibrium is bounded by the social cost of the *fully mixed Nash equilibrium*. A Nash equilibrium is called fully mixed if each user chooses each link with non-zero probability. We present a complete characterization of the instances for which a fully mixed Nash equilibrium exists, and prove that (in case of its existence) it is unique. Moreover, we give bounds on the *coordination ratio* and show that several results for the Wardrop model can be carried over to our discrete model.

## 1 Introduction

**Motivation and Framework.** One of the most important concepts in non-cooperative game theory is the concept of *Nash equilibria* [22]. A Nash equilibrium is a state of the system in which no player can improve its objective by unilaterally changing its *strategy*. A Nash equilibrium is called *pure* if all players choose exactly one strategy, and *mixed* if players choose probability distributions over strategies. The *coordination ratio* is the worst-case ratio of the *social cost*

---

<sup>★</sup> This work has been partially supported by the European Union within the 6th Framework Programme under contract 001907 (DELIS), by the IST Program of the European Union under contract number IST-2001-33116 (FLAGS), by funds from the Joint Program of Scientific and Technological Collaboration between Greece and Cyprus, by research funds at University of Cyprus, and by the VEGA grant No. 2/3164/23.

<sup>\*\*</sup> International Graduate School of Dynamic Intelligent Systems

in a Nash equilibrium state and the minimum social cost. Of special interest to our work is the *fully mixed Nash equilibrium* where each player chooses each strategy with non-zero probability. We consider a hybridization of the two most famous models for non-cooperative routing in literature: the *KP model* [16] and the *Wardrop model* [8, 27].

In the KP model, each of  $n$  users employs a *mixed strategy*, which is a probability distribution over  $m$  parallel *links*, to control the shipping of its *traffic*. Traffic is *unsplittable*. A *capacity* specifies the rate at which each link processes traffic. *Identical users* have the same traffic whereas the traffic of the users may vary arbitrarily in the model of *arbitrary users*. In a Nash equilibrium, each user selfishly routes its traffic on links that minimize its *individual cost*: its *expected latency cost*, given the expected network congestion caused by the other users. The *social cost* of a Nash equilibrium is the expectation, over all random choices of the users, of the maximum *latency* through a link (over all links).

In the Wardrop model, *arbitrary* networks with *latency functions* for edges are considered. Moreover, the traffic is *splittable* into arbitrary pieces. Here, unregulated traffic is modeled as a *network flow*. *Equilibrium flows* are flows with all paths used between a given pair of a *source* and a *destination* having the same latency. The latency functions are convex. Thus, equilibrium flows are optimal solutions to a convex program. An equilibrium in this model can be interpreted as a Nash equilibrium in a game with an infinite number of users, each carrying an infinitesimal amount of traffic from a *source* to a *destination*. The Wardrop model restricts to pure Nash equilibria. The *individual cost* of a user is the sum of the edge *latencies* on a path from the user's source to the its destination. The *social cost* of a Nash equilibrium is the sum of all individual costs.

The routing model considered in this work combines aspects of both the KP model and the Wardrop model. First, we restrict the network structure to that of the KP model (parallel links) and we assume a user's traffic to be unsplittable. On the other hand, we allow arbitrary non-decreasing and convex latency functions, whereas in the KP model latency functions are linear. In our model, the latency function of a link is a function in the total traffic of users assigned to this link. The social cost is defined as the expected sum of all user costs – as opposed to the social cost used in the KP model. Thus, as far as the generality of latency functions and the definition of social cost are concerned, we lean toward the Wardrop model, whereas the network structure and the indivisibility of each user's traffic remain as in the KP model. Restricted to *pure* Nash equilibria, our model has already been studied in [6], and restricted to *linear* latency functions in [18]. It is a particular instance of what is known as *congestion game* [21, 23]. It is known that a pure Nash equilibrium always exists in this setting.

The main results of this work are the identification of the worst-case mixed Nash equilibrium and bounds on the coordination ratio. The convex latency functions define a very general discrete routing game. To the best of our knowledge this is the first time that mixed Nash equilibria are studied in such a game. **Related Work.** The KP model was introduced by Koutsoupias and Papadimitriou [16]. They introduced the notion of coordination ratio and analyzed the

coordination ratio for some special cases. Later, Czumaj and Vöcking [7], and Koutsoupias et al. [15] gave asymptotically tight upper bounds on the coordination ratio for pure and mixed Nash equilibria. Mavronicolas and Spirakis [20] studied further the KP model and introduced the fully mixed Nash equilibrium. They showed that, in case it exists, the fully mixed Nash equilibrium is unique. Gairing et al. [12] conjecture that the fully mixed Nash equilibrium, whenever it exists, has the worst social cost among all Nash equilibria. From here on we will refer to this as the *Fully Mixed Nash Equilibrium Conjecture*. Up to now, the conjecture could be proven only for several particular cases of the KP model [12, 19]. A proof of the conjecture will enable the derivation of upper bounds on the coordination ratio via studying the fully mixed Nash equilibrium.

Lücking et al. [18] considered the KP model with respect to *quadratic social cost*, defined as the sum of weighted individual costs. In this context, they proved the Fully Mixed Nash Equilibrium Conjecture in the case of identical users and identical links. This result is strongly related to results presented in this paper.

A natural problem is the effective computation of a Nash equilibrium. For general strategic games, it is still open as to whether a Nash equilibrium can be computed in polynomial time, even for two player games. Fotakis et al. [11] showed that a pure Nash equilibrium for the KP model can be computed in polynomial time using Graham’s algorithm [13]. Furthermore, they proved that the problem to compute the best or worst pure Nash equilibrium is  $\mathcal{NP}$ -complete. Feldmann et al. [9] showed that any deterministic assignment of users to links can be transformed into a Nash equilibrium in polynomial time without increasing the social cost. In particular, combining this result with known approximation algorithms for the computation of optimal assignments [14] yields a PTAS for the problem to compute a best pure Nash equilibrium.

The Wardrop model was already studied in the 1950’s [2, 27], in the context of road traffic systems. Wardrop [27] introduced the concept of equilibrium to describe user behavior in this kind of traffic networks. For a survey of the early work on this model, see [3]. A lot of subsequent work on this model has been motivated by Braess’s Paradox [5]. Inspired by the new interest in the coordination ratio, Roughgarden and Tardos [24–26] re-investigated the Wardrop model. For a survey of results, we refer to [10] and references therein.

**Results.** With our methods, we can only prove results for identical users. However, for this case we obtain through a very thorough analysis the following

- In the case of its existence, the fully mixed Nash equilibrium is the worst-case Nash equilibrium for any instance with convex latency functions. Therewith, we prove the Fully Mixed Nash Equilibrium Conjecture to hold for the model under consideration, whereas it remains unproven for the KP model in the general case. This broadens some recent results from [18] for a special case of our model, where latency functions are restricted to be linear. We use an appropriate counterexample to show that the convexity assumption we are making for the latency functions cannot be relaxed.
- For arbitrary non-decreasing and non-constant latency functions, the fully mixed Nash equilibrium is unique in the case of its existence.

- We give a complete characterization of instances for which the fully mixed Nash equilibrium exists.
- For pure Nash equilibria we adapt an upper bound on the coordination ratio from Roughgarden and Tardos [26] to our (discrete) model. This bound holds for non-decreasing and non-constant latency functions. Considering polynomial latency functions with non-negative coefficients and of maximum degree  $d$ , this yields an upper bound of  $d + 1$ .
- For identical links with latency function  $f(x) = x^d$ ,  $d \in \mathbb{N}$ , the coordination ratio for mixed Nash equilibria is bounded by the  $(d + 1)$ 'th Bell number. This bound can be approximated arbitrarily but never reached.
- We give a  $\mathcal{O}(m \log n \log m)$  algorithm to compute a pure Nash equilibrium for non-decreasing latency functions.
- For arbitrary users, computing the best-case or worst-case pure Nash equilibrium is  $\mathcal{NP}$ -hard even for identical links with a linear latency function.

**Road Map.** Section 2 introduces notations and terminology. In Section 3, the Fully Mixed Nash Equilibrium Conjecture is proven for the model we consider. The necessity of the convexity assumption is also established there. Furthermore, we determine the conditions under which the fully mixed Nash equilibrium exists. Section 4 presents bounds on coordination ratio and complexity results.

## 2 Discrete Routing Games

**General.** The number of ways a set of  $k$  elements can be partitioned into non-empty subsets is called the  $k$ -th Bell Number [4, 28], denoted by  $B_k$ . It is defined by the recursive formula  $B_0 = 1$  and

$$B_{k+1} = \sum_{0 \leq q \leq k} B_q \cdot \binom{k}{q} \quad \text{for all } k \geq 0. \quad (1)$$

Throughout, denote for any integer  $m \geq 1$ ,  $[m] = \{1, \dots, m\}$ .

We consider a *network* consisting of a set of  $m$  parallel *links*  $1, 2, \dots, m$  from a *source* node to a *destination* node. Each of  $n$  *network users*  $1, 2, \dots, n$ , or *users* for short, wishes to route a particular amount of traffic along a (non-fixed) link from source to destination. Denote as  $w_i$  the *traffic* of user  $i \in [n]$ . Define the  $n \times 1$  *traffic vector*  $\mathbf{w}$  in the natural way. For any subset  $A \subseteq [n]$  of users, denote  $w_A = \sum_{i \in A} w_i$ . If users are *identical*, we assume that  $w_i = 1$  for all  $i \in [n]$ . In this case,  $w_A$  reduces to  $|A|$ . Assume throughout that  $m > 1$  and  $n > 1$ .

A *pure strategy* for user  $i \in [n]$  is some specific link. A *mixed strategy* for user  $i \in [n]$  is a probability distribution over pure strategies; thus, a mixed strategy is a probability distribution over the set of links. The *support* of the mixed strategy for user  $i \in [n]$ , denoted as  $\text{support}(i)$ , is the set of those pure strategies (links) to which  $i$  assigns positive probability. A *pure strategy profile* is represented by an  $n$ -tuple  $\langle \ell_1, \ell_2, \dots, \ell_n \rangle \in [m]^n$ ; a *mixed strategy profile* is represented by an  $n \times m$  *probability matrix*  $\mathbf{P}$  of  $nm$  probabilities  $p(i, j)$ ,  $i \in [n]$  and  $j \in [m]$ , where  $p(i, j)$  is the probability that user  $i$  chooses link  $j$ .

For a probability matrix  $\mathbf{P}$ , define *indicator variables*  $I(i, j) \in \{0, 1\}$ , where  $i \in [n]$  and  $j \in [m]$ , such that  $I(i, j) = 1$  if and only if  $p(i, j) > 0$ . Thus, the

support of the mixed strategy for user  $i \in [n]$  is the set  $\{j \in [m] \mid I(i, j) = 1\}$ . A mixed strategy profile  $\mathbf{P}$  is *fully mixed* [20, Section 2.2] if for all users  $i \in [n]$  and links  $j \in [m]$ ,  $I(i, j) = 1$ . Throughout, we will cast a pure strategy profile as a special case of a mixed strategy profile in which all strategies are pure.

**System, Models and Cost Measures.** Associated with every link  $j \in [m]$ , is a latency function  $f_j : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ ,  $f_j(0) = 0$ , which is non-decreasing and non-constant. Define the  $m \times 1$  *vector of latency functions*  $\Phi$  in the natural way. If  $f_j = f$  for all  $j \in [m]$ , we say that the links are *identical*, otherwise they are *arbitrary*. For a pure strategy profile  $\langle \ell_1, \ell_2, \dots, \ell_n \rangle$ , the *individual latency cost for user  $i \in [n]$* , denoted by  $\lambda_i$ , is defined by  $f_j(\sum_{k \in [n] : \ell_k = j} w_k)$ , with  $j = \ell_i$ . For a mixed strategy profile  $\mathbf{P}$ , denote as  $A_j$  the *expected latency* on link  $j \in [m]$ , i.e.

$$A_j = \sum_{A \subseteq [n]} \prod_{k \in A} p(k, j) \cdot \prod_{k \notin A} (1 - p(k, j)) \cdot f_j(w_A).$$

The *expected latency cost* for user  $i \in [n]$  on link  $j \in [m]$ , denoted by  $\lambda_{ij}$ , is the expectation, over all random choices of the remaining users, of the individual latency cost for user  $i$  had its traffic been assigned to link  $j$ ; thus,

$$\begin{aligned} \lambda_{ij} &= \sum_{\langle \ell_1, \dots, \ell_n \rangle} \prod_{k \in [n] \setminus \{i\}} p(k, \ell_k) \cdot f_j(w_i + \sum_{\substack{k \in [n] \setminus \{i\} \\ \ell_k = j}} w_k) \\ &= \sum_{A \subseteq [n] \setminus \{i\}} \prod_{k \in A} p(k, j) \prod_{k \notin A \cup \{i\}} (1 - p(k, j)) \cdot f_j(w_i + w_A). \end{aligned}$$

For each user  $i \in [n]$ , the *expected individual latency cost*, denoted by  $\lambda_i$ , is the expectation, over all links  $j \in [m]$ , of the expected latency cost for user  $i$  on link  $j$ ; thus,  $\lambda_i = \sum_{j \in [m]} p(i, j) \cdot \lambda_{ij}$ . Associated with a mixed strategy profile  $\mathbf{P}$  and a vector of latency functions  $\Phi$  is the *social cost*, denoted by  $\text{SC}^\Sigma(\Phi, \mathbf{P})$ , which is the sum, over all users, of the expected individual latency costs of the users. Thus,  $\text{SC}^\Sigma(\Phi, \mathbf{P}) = \sum_{i \in [n]} \lambda_i$ . On the other hand, the *social optimum*, denoted by  $\text{OPT}^\Sigma(\Phi)$ , is the least possible value, over all pure strategy profiles  $\mathbf{L}$ , of the social cost. Thus,  $\text{OPT}^\Sigma(\Phi) = \min_{\mathbf{L}} \text{SC}^\Sigma(\Phi, \mathbf{L})$ .

**Nash Equilibria and Coordination Ratio.** We are interested in a special class of mixed strategies called Nash equilibria [22] that we describe below. Say that a user  $i \in [n]$  is *satisfied for the probability matrix  $\mathbf{P}$*  if  $\lambda_{ij} = \lambda_i$  for all links  $j \in \text{support}(i)$ , and  $\lambda_{ij} \geq \lambda_i$  for all  $j \notin \text{support}(i)$ . Otherwise, user  $i$  is *unsatisfied*. Thus, a satisfied user has no incentive to unilaterally deviate from its mixed strategy.  $\mathbf{P}$  is a *Nash equilibrium* [16, Section 2] if and only if all users  $i \in [n]$  are satisfied for  $\mathbf{P}$ . The *coordination ratio* is the maximum value, over all vectors of latency functions  $\Phi$  and Nash equilibria  $\mathbf{P}$ , of the ratio  $\text{SC}^\Sigma(\Phi, \mathbf{P})/\text{OPT}^\Sigma(\Phi)$ .

### 3 Results on Fully Mixed Nash Equilibria

For the model of identical users, we now consider fully mixed Nash Equilibria. We start with a definition and a technical lemma. Both can be proven for the model of arbitrary users, and are useful several times throughout the paper.

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

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

In the same way, we define a function  $\tilde{H}(\tilde{p}, r, \mathbf{w}, g)$  by replacing  $p$  with a vector of  $r$  probabilities all equal to  $\tilde{p}$ . In the case that all users have the same traffic, we omit  $\mathbf{w}$  in the parameter list. Note that  $w_A$  reduces to  $|A|$  in this case.

We prove a natural monotonicity property of the function  $H(p, \mathbf{w}, g)$ .

**Lemma 1.** For every vector of  $r$  probabilities  $p = (p_1, \dots, p_r)$  and every non-decreasing and non-constant function  $g : \mathbb{R} \rightarrow \mathbb{R}$ ,  $H(p, \mathbf{w}, g)$  is strictly increasing in each probability  $p_i, \forall i \in [r]$ .

*Proof.* We prove, that  $H(p, \mathbf{w}, g)$  is strictly increasing in  $p_r$ . The lemma then follows by symmetry of  $H(p, \mathbf{w}, g)$  in all probabilities  $p_i, i \in [r]$ . It is

$$\begin{aligned} H(p, \mathbf{w}, g) &= \sum_{A \subseteq [r]} \prod_{k \in A} p_k \prod_{k \notin A} (1 - p_k) \cdot g(w_A) \\ &= \sum_{A \subseteq [r-1]} \prod_{k \in A} p_k \prod_{k \notin A \cup \{r\}} (1 - p_k) \cdot [g(w_A) + p_r \cdot (g(w_A + w_r) - g(w_A))] \end{aligned}$$

As  $g(w_A + w_r) - g(w_A) \geq 0$  for all  $A \subseteq [r-1]$  ( $g$  is non-decreasing), and  $g(w_A + w_r) - g(w_A) > 0$  for some  $A \subseteq [r-1]$  ( $g$  is non-constant), the claim follows.  $\square$

### 3.1 The Worst-Case Nash Equilibrium

We now focus on the Fully Mixed Nash Equilibrium Conjecture. We first show that for an arbitrary Nash equilibrium  $\mathbf{P}$ , the expected latency of a user  $i$  on a link  $j$  increases if we set all user probabilities on link  $j$  to be the average probability on that link. We then use this result to show that the expected individual latency of user  $i$  in the Nash equilibrium  $\mathbf{P}$  is at most its expected individual latency in the fully mixed Nash equilibrium. By definition, this proves the Fully Mixed Nash Equilibrium Conjecture for our model. We furthermore give an example with a strictly increasing but non-convex latency function for which the Fully Mixed Nash Equilibrium Conjecture does not hold, showing that the assumption of convexity for the latency functions is essential.

**Lemma 2.** Let  $g$  be convex and define  $p = (p_1, \dots, p_n)$  and  $\tilde{p} = \frac{\sum_{i \in [n]} p_i}{n}$ . Then  $H(p, g) \leq \tilde{H}(\tilde{p}, n, g)$ .

*Proof.* Define a set of  $n$  probabilities  $q = (q_1, \dots, q_n)$  by  $q_1 = q_2 = \frac{p_1 + p_2}{2}$  and  $q_i = p_i, \forall i \in [3, n]$ . Then

$$H(p, g) = \sum_{A \subseteq [3, n]} \prod_{k \in A} p_k \prod_{k \notin A \cup \{1, 2\}} (1 - p_k) \cdot F(|A|, p, g),$$

where

$$F(|A|, p, g) = p_1 \cdot p_2 \cdot [g(|A| + 2) - 2g(|A| + 1) + g(|A|)] \\ + (p_1 + p_2) \cdot [g(|A| + 1) - g(|A|)] + g(|A|).$$

Similarly,

$$H(q, g) = \sum_{A \subseteq [3, n]} \prod_{k \in A} q_k \prod_{k \notin A \cup \{1, 2\}} (1 - q_k) \cdot F(|A|, q, g).$$

It suffices to show, that  $F(|A|, q, g) - F(|A|, p, g) \geq 0$ . Indeed,

$$F(|A|, q, g) - F(|A|, p, g) = (q_1 \cdot q_2 - p_1 \cdot p_2) \cdot [g(|A| + 2) - 2g(|A| + 1) + g(|A|)] \\ + (q_1 + q_2 - (p_1 + p_2)) \cdot [g(|A| + 1) - g(|A|)] \\ = \left( \frac{p_1 - p_2}{2} \right)^2 [g(|A| + 2) - 2g(|A| + 1) + g(|A|)] \geq 0,$$

since  $g$  is convex.  $\square$

**Lemma 3.** *Consider the model of identical users and arbitrary links with non-decreasing, non-constant and convex latency functions. If there exists a fully mixed Nash equilibrium  $\mathbf{F}$ , then for every mixed Nash equilibrium  $\mathbf{P}$ ,  $\lambda_i(\mathbf{P}) \leq \lambda_i(\mathbf{F})$  for all  $i \in [n]$ .*

*Proof.* Define  $\Theta_{ij} = \sum_{k \in [n], k \neq i} p(k, j)$  and  $\tilde{p}(j) = \frac{\Theta_{ij}}{n-1}$ . The claim holds if  $\lambda_{ij}(\mathbf{P}) \leq \lambda_i(\mathbf{F})$ , for all  $i \in [n], j \in [m]$ . So assume there exists  $i \in [n]$  and  $j \in [m]$  with  $\lambda_{ij}(\mathbf{P}) > \lambda_i(\mathbf{F})$ . By Lemma 2

$$\lambda_{ij}(\mathbf{P}) \leq \sum_{A \subseteq [n] \setminus \{i\}} \prod_{k \in A} \tilde{p}(j) \prod_{k \notin A \cup \{i\}} (1 - \tilde{p}(j)) \cdot f_j(1 + |A|), \quad \text{and} \\ \lambda_{ij}(\mathbf{F}) = \sum_{A \subseteq [n] \setminus \{i\}} \prod_{k \in A} p_F(j) \prod_{k \notin A \cup \{i\}} (1 - p_F(j)) \cdot f_j(1 + |A|),$$

where  $p_F(j)$  is the probability for any user to choose link  $j$  in the fully mixed Nash equilibrium  $\mathbf{F}$ . Note that the upper bound on  $\lambda_{ij}(\mathbf{P})$  is strictly increasing in  $\tilde{p}(j)$ , since  $f_j$  is non-decreasing and non-constant. Therefore,  $\lambda_{ij}(\mathbf{P}) > \lambda_i(\mathbf{F})$  implies that  $\tilde{p}(j) > p_F(j)$ . Since  $\sum_{j \in [m]} \tilde{p}(j) = \sum_{j \in [m]} p_F(j) = 1$ , there exists a link  $k$  with  $\tilde{p}(k) < p_F(k)$ . However, this implies that  $\lambda_{ik}(\mathbf{P}) < \lambda_i(\mathbf{F})$  and thus  $\lambda_i(\mathbf{P}) < \lambda_i(\mathbf{F})$ .  $\square$

If we look at the different model where latency functions only depend on the user and not on the link, we know that there exists a fully mixed Nash equilibrium with probabilities  $p(i, j) = \frac{1}{m}$  for all  $i \in [n]$  and  $j \in [m]$ . With the same method as in Lemma 3, we can prove that the expected individual latency of a user is bounded by its expected individual latency of this fully mixed Nash equilibrium.

**Theorem 1.** *Consider the model of identical users and arbitrary links with non-decreasing, non-constant and convex latency functions. If the fully mixed Nash equilibrium  $\mathbf{F}$  exists, then for every mixed Nash equilibrium  $\mathbf{P}$ ,  $\text{SC}^\Sigma(\Phi, \mathbf{P}) \leq \text{SC}^\Sigma(\Phi, \mathbf{F})$ .*

*Proof.* Follows from the definition of  $\text{SC}^\Sigma(\Phi, \mathbf{P})$  combined with Lemma 3.  $\square$

The Fully Mixed Nash Equilibrium Conjecture has been proven for the model of identical users, identical links and latency function  $f(x) = x$  by Lücking et al. [18]. Theorem 1 generalizes this result to non-decreasing, non-constant and convex latency functions. We continue to prove that the convexity assumption is essential.

**Proposition 1.** *There exists an instance with identical users, identical links and a non-decreasing, non-convex latency function with a pure Nash equilibrium  $\mathbf{L}$  and fully mixed Nash equilibrium  $\mathbf{F}$  such that  $\lambda_i(\mathbf{L}) > \lambda_i(\mathbf{F})$  for all  $i \in [n]$ .*

*Proof.* Consider an instance with  $m = 2$  links and  $n = 4$  users. Define  $f$  as follows:  $f(1) = 1, f(2) = 2, f(3) = 2 + \epsilon, f(4) = 2 + 2\epsilon$ , where  $\epsilon > 0$ . Then in each pure Nash equilibrium, there are exactly 2 users on each link. Let  $\mathbf{L}$  be such a pure Nash equilibrium. Then  $\lambda_i(\mathbf{L}) = 2$  for all  $i \in [n]$ . Now consider the fully mixed Nash equilibrium  $\mathbf{F}$ . Here  $p(i, j) = \frac{1}{2}$  for all  $i \in [n], j \in [m]$ . Thus,

$$\lambda_i(\mathbf{F}) = \frac{1}{8}(f(1) + 3f(2) + 3f(3) + f(4)) = \frac{15}{8} + \frac{5\epsilon}{8}, \quad \forall i \in [n].$$

For  $\epsilon < \frac{1}{5}$  it follows that  $\lambda_i(\mathbf{L}) > \lambda_i(\mathbf{F})$  for all  $i \in [n]$ . □

### 3.2 Uniqueness of the Fully Mixed Nash Equilibrium

We first show that the probabilities of all users on a certain link are identical in a fully mixed Nash equilibrium. We then use this fact to establish uniqueness of the fully mixed Nash equilibrium.

**Theorem 2 (Uniqueness of the Fully Mixed Nash Equilibrium).** *Consider the model of identical users and arbitrary links with non-decreasing and non-constant latency functions. If a fully mixed Nash equilibrium  $\mathbf{F}$  exists, then it is unique.*

### 3.3 Existence of Fully Mixed Nash Equilibrium

For the special case where all latency functions are equal, i.e.  $f_j = f$  for all  $j \in [m]$ , a fully mixed Nash equilibrium always exists and has probabilities  $p(i, j) = \frac{1}{m}$  for all  $i \in [n], j \in [m]$ . For the general case, the existence of the fully mixed Nash equilibrium is not granted, but depends on the latency functions  $f_j$ . We will now shed light on this dependence. Without loss of generality, assume the links to be ordered non-decreasingly according to  $f_j(1)$ . Let  $g_j : [n-1] \cup \{0\} \rightarrow \mathbb{R}$  be defined by  $g_j(x) = f_j(x+1)$  for all  $j \in [m]$ . For  $k \in [m], j \in [k-1]$ , determine  $p_j(k)$ , such that  $\tilde{H}(p_j(k), n-1, g_j) = f_k(1)$ . Then,  $\tilde{H}(p_j(k), n-1, g_j)$  is the expected individual latency of any user on link  $j$ , if  $p(i, j) = p_j(k)$  for all  $i \in [n]$ . Note, that due to Lemma 1,  $\tilde{H}(p_j(k), n-1, g_j)$  is strictly increasing in  $p_j(k)$ , and hence  $p_j(k)$  is uniquely determined.

**Definition 2.** *Links  $k$  with  $\sum_{j \in [k-1]} p_j(k) > 1$  are called dead links. Links  $k$  with  $\sum_{j \in [k-1]} p_j(k) = 1$  are called special links.*

**Lemma 4.** *Consider the model of identical users and arbitrary links with non-decreasing and non-constant latency functions. If  $j \in [m]$  is a dead link, then in any Nash equilibrium  $\mathbf{P}$ ,  $p(i, j) = 0$  for all  $i \in [n]$ .*

**Lemma 5.** *Consider the model of identical users and arbitrary links with non-decreasing and non-constant latency functions. Let  $S$  be the set of special links. In any Nash equilibrium  $\mathbf{P}$ , there exists at most one user  $i$  with  $p(i, j) > 0$  for some  $j \in S$ .*

**Theorem 3 (Characterization of Fully Mixed Nash Equilibria).** *Consider the model of identical users and arbitrary links with non-decreasing and non-constant latency functions. There exists a fully mixed Nash equilibrium, if and only if there are no special and no dead links.*

Theorem 3 implies that if the fully mixed Nash equilibrium does not exist, then the instance contains dead or special links. But dead links are never used in any Nash equilibrium and could be removed from the instance. We now broaden the result from Theorem 3 by giving an upper bound on the social cost in the case that the fully mixed Nash equilibrium does not exist.

**Theorem 4.** *Consider an instance with special or dead links. Then the social cost of any Nash equilibrium  $\mathbf{P}$  is bounded by the social cost of the fully mixed Nash equilibrium  $\mathbf{F}$  for the instance where the links are restricted to the non-special and non-dead links.*

## 4 Coordination Ratio and Complexity Results

### 4.1 Bounds on Coordination Ratio for Special Latency Functions

We now consider the model of identical users and identical links with latency function  $f(x) = x^d$ ,  $d \in \mathbb{N}$ . In this model, every pure Nash equilibrium has optimal social cost. For mixed Nash equilibria, we now show that the coordination ratio is bounded by the  $(d+1)$ -th Bell Number  $B_{d+1}$  (see Equation (1)). Due to [17],  $B_{d+1} \approx (d+1)^{-\frac{1}{2}} [\gamma(d+1)]^{d+1+\frac{1}{2}} e^{\gamma(d+1)-d-2}$ , where the function  $\gamma(d+1)$  is defined implicitly by  $\gamma(d+1) \cdot \ln(\gamma(d+1)) = d+1$ .

**Theorem 5.** *Consider the model of identical users and identical links with latency function  $f(x) = x^d$ ,  $d \in \mathbb{N}$ . Then,*

$$\sup_{\mathbf{w}, \mathbf{P}} \frac{\text{SC}^\Sigma(\Phi, \mathbf{P})}{\text{OPT}^\Sigma(\Phi)} = B_{d+1}.$$

### 4.2 Bounds on Coordination Ratio for General Latency Functions

In this section, we carry over an upper bound from Roughgarden and Tardos [26, Corollary 2.10] on the coordination ratio for splittable flows and continuous latency functions to our discrete setting. For the proof, which is a straightforward adaptation of the corresponding proof in [26], we make use of the following lemma.

**Lemma 6.** *Let  $g_j : [n] \rightarrow \mathbb{R}$  be a convex function for  $j \in [m]$ . Set  $X = \{x = (x_1, \dots, x_m) \in \mathbb{N}_0^m \mid \sum_{j \in [m]} x_j = n\}$ . Then  $\sum_{j \in [m]} g_j(x_j)$  is minimum among all  $x = (x_1, \dots, x_m) \in X$ , if and only if*

$$g_j(x_j + 1) + g_k(x_k - 1) \geq g_j(x_j) + g_k(x_k) \quad \forall j, k \in [m].$$

Lemma 6 can be shown by the application of convex cost flows [1, Chapter 14].

**Lemma 7.** *Consider the model of identical users and arbitrary links with non-decreasing and non-constant latency functions. If  $xf_j(x) \leq \alpha \sum_{t=1}^x f_j(t)$  for all  $j \in [m]$ , then the social cost of any pure Nash equilibrium is bounded by  $\alpha \text{OPT}^\Sigma(\Phi)$ .*

The following corollary is an example for the application of the upper bound.

**Corollary 1.** *Consider the model of identical users and arbitrary links. If latency functions are polynomials with non-negative coefficients and maximum degree  $d$ , then the coordination ratio for pure Nash equilibria is bounded by  $d + 1$ .*

### 4.3 Computation of Pure Nash Equilibrium and Optimum

In the model of identical users and identical links, the users are evenly distributed to the links in every pure Nash equilibrium, and every pure Nash equilibrium has optimum social cost. In the following, we give an algorithm to compute a pure Nash equilibrium in the model of identical users but arbitrary non-decreasing latency functions. A simple approach is to assign the users one by one to their respective best link. This greedy algorithm, also known as Graham's algorithm, can be implemented with running time  $\mathcal{O}((n + m) \log m)$  if the links are kept in a priority queue according to their latency after the assignment of the next user.

**ALGORITHM 1**

**Input:**  $n$  and any assignment  $x_1, \dots, x_m$

**Output:** Nash Equilibrium  $x_1, \dots, x_m$

```

for  $\delta = n, \lceil \frac{n}{2} \rceil, \lceil \frac{n}{4} \rceil, \dots, 1$  do
  let  $t$  be such that  $f_t(x_t + \delta)$  is minimum;
  while  $\exists s \in [n]$  with  $x_s \geq \delta$  and  $f_s(x_s) > f_t(x_t + \delta)$  do
    let  $s \in [m]$  be such that  $x_s \geq \delta$  and
       $f_s(x_s)$  is maximum w.r.t. this requirement;
     $x_s = x_s - \delta$ ;  $x_t = x_t + \delta$ ;
    let  $t$  be such that  $f_t(x_t + \delta)$  is minimum;

```

Our algorithm takes time  $\mathcal{O}(m \log n \log m)$ , which is better if  $m = o(\frac{n}{\log n})$ . The algorithm takes as input an arbitrary initial assignment of users to links given by  $x_1, \dots, x_m$ , where  $x_j$  is the number of users on link  $j$ . It transforms this assignment into a Nash equilibrium by moving chunks of users at a time. The first chunk contains all users. In each phase the chunk size is cut in half until a chunk consists of one user only. In the sequel we refer to  $x_j$  as the *load* on link  $j \in [m]$ .

**Proposition 2.** *Consider the model of identical users and arbitrary links with non-decreasing latency functions. Then Algorithm 1 computes a pure Nash equilibrium in time  $\mathcal{O}(m \log n \log m)$ .*

The following lemma shows that we can compute an optimal pure assignment in the same way as a Nash equilibrium, but according to other latency functions.

A corresponding result holds for the case of continuous latency functions and splittable flows (see e.g. [26]).

**Lemma 8.** *Consider an instance of the routing model with identical users and  $m$  links with latency function  $f_j(x)$  on link  $j$  for  $j \in [m]$ , such that  $xf_j(x)$  is convex. Set  $h_j(x) = xf_j(x) - (x-1)f_j(x-1)$ . Let  $\mathbf{L}$  be any pure strategy profile.  $\mathbf{L}$  is an optimal assignment with respect to latency functions  $f_j$ , if and only if  $\mathbf{L}$  is a Nash equilibrium with respect to latency functions  $h_j$ .*

Due to Lemma 8, Algorithm 1 can be used to compute an optimal pure assignment by applying it to the instance with latency functions  $h_j$  on link  $j$ .

#### 4.4 Complexity Results

Fotakis et al. [11] proved that computing the best-case or worst-case pure Nash equilibrium in the KP model is  $\mathcal{NP}$ -hard. Keep in mind that in the KP model the social cost of a pure Nash equilibrium is the maximum latency on a link, whereas in our model the social cost is the sum of the individual latency costs. We now show that computing the best-case or the worst-case pure Nash equilibrium in our model is also NP-hard even for identical links with latency function  $f(x) = x$ .

**Proposition 3.** *Consider the model of arbitrary users and identical links with latency function  $f(x) = x$ . Then, computing the best-case or the worst-case pure Nash equilibrium is NP-hard.*

It is easy to see that Graham’s algorithm [13] (known to work for the KP model [11]) still works for the model under consideration to compute a pure Nash equilibrium in polynomial time.

## References

1. R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice Hall, 1993.
2. M. Beckmann, C.B. McGuire, and C.B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
3. M.J. Beckmann. On the theory of traffic flow in networks. *Traffic Quart*, 21:109–116, 1967.
4. E.T. Bell. Exponential numbers. *American Mathematical Monthly*, 41(7):411–419, 1934.
5. D. Braess. Über ein Paradoxon der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.
6. A. Czumaj, P. Krysta, and B. Vöcking. Selfish traffic allocation for server farms. In *Proc. of the 34th Ann. ACM Symp. on Theory of Computing*, pp. 287–296, 2002.
7. A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *Proc. of the 13th Ann. ACM-SIAM Symp. on Discrete Algorithms*, pp. 413–420, 2002.
8. 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*, 73B(2):91–118, 1969.

9. R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Nashification and the coordination ratio for a selfish routing game. In *Proc. of the 30th Int. Colloq. on Automata, Languages, and Programming*, LNCS 2719, pp. 514–526, 2003.
10. R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish routing in non-cooperative networks: A survey. In *Proc. of the 28th Int. Symp. on Mathematical Foundations of Computer Science*, LNCS 2747, pp. 21–45, 2003.
11. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of nash equilibria for a selfish routing game. In *Proc. of the 29th Int. Colloq. on Automata, Languages, and Programming*, LNCS 2380, pp. 123–134, 2002.
12. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Extreme nash equilibria. In *Proc. of the 8th Italian Conference on Theoretical Computer Science*, LNCS 2841, pp. 1–20, 2003. Also accepted to *Theoretical Computer Science*, Special Issue on *Game Theory Meets Theoretical Computer Science*.
13. R.L. Graham. Bounds on multiprocessing timing anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.
14. D.S. Hochbaum and D. Shmoys. A polynomial approximation scheme for scheduling on uniform processors: using the dual approximation approach. *SIAM Journal on Computing*, 17(3):539–551, 1988.
15. E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate equilibria and ball fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.
16. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. of the 16th Int. Symp. on Theoretical Aspects of Computer Science*, LNCS 1563, pp. 404–413, 1999.
17. L. Lovász. *Combinatorial Problems and Exercises*. North-Holland, 1993.
18. T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. In *Proc. of the 21st Int. Symp. on Theoretical Aspects of Computer Science*, LNCS 2996, pp. 547–558, 2004.
19. T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the worst-case nash equilibrium? In *Proc. of the 28th Int. Symp. on Mathematical Foundations of Computer Science*, LNCS 2747, pp. 551–561, 2003.
20. M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proc. of the 33rd Ann. ACM Symp. on Theory of Computing*, pp. 510–519, 2001.
21. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13:111–124, 1996.
22. J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
23. R.W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *Int. Journal of Game Theory*, 2:65–67, 1973.
24. T. Roughgarden. Stackelberg scheduling strategies. In *Proc. of the 33rd Ann. ACM Symp. on Theory of Computing*, pp. 104–113, 2001.
25. T. Roughgarden. The price of anarchy is independent of the network topology. In *Proc. of the 34th Ann. ACM Symp. on the Theory of Computing*, pp. 428–437, 2002.
26. T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
27. J.G. Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II, Vol. 1*, pp. 325–378, 1956.
28. H.S. Wilf. *Generatingfunctionology*. Academic Press, 1994.