Alcom-FT Technical Report Series
ALCOMFT-TR-03-146

Selfish Routing in Non-Cooperative Networks: A Survey ⋆

R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode⋆⋆

Department of Computer Science, Electrical Engineering and Mathematics
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
{obelix,gairing,luck,bm,rode}@uni-paderborn.de

Abstract. We study the problem of n users selfishly routing traffics through a shared network. Users route their traffics by choosing a path from their source to their destination of the traffic with the aim of minimizing their private latency. In such an environment Nash equilibria represent stable states of the system: no user can improve its private latency by unilaterally changing its strategy.
In the first model the network consists only of a single source and a single destination which are connected by m parallel links. Traffics are unsplittable. Users may route their traffics according to a probability distribution over the links. The social optimum minimizes the maximum load of a link. In the second model the network is arbitrary, but traffics are splittable among several paths leading from their source to their destination. The goal is to minimize the sum of the edge latencies.
Many interesting problems arise in such environments: A first one is the problem of analyzing the loss of efficiency due to the lack of central regulation, expressed in terms of the coordination ratio. A second problem is the Nashification problem, i.e. the problem of converting any given non-equilibrium routing into a Nash equilibrium without increasing the social cost. The Fully Mixed Nash Equilibrium Conjecture (FMNE Conjecture) states that a Nash equilibrium, in which every user routes along every possible edge with probability greater than zero, is a worst Nash equilibrium with respect to social cost. A third problem is to exactly specify the sub-models in which the FMNE Conjecture is valid. The well-known Braess's Paradox shows that there exist networks, such that strict sub-networks perform better when users are selfish. A natural question is the following network design problem: Given a network, which edges should be removed to obtain the best possible Nash equilibrium.
We present complexity results for various problems in this setting, upper and lower bounds for the coordination ratio, and algorithms solving the problem of Nashification. We survey results on the validity of the FMNE Conjecture in the model of unsplittable flows, and for the model of splittable flows we survey results for the network design problem.

⋆ Partly supported by the DFG-SFB 376 and by the IST Program of the EU under contract numbers IST-1999-14186 (ALCOM-FT), and IST-2001-33116 (FLAGS).
⋆⋆ International Graduate School of Dynamic Intelligent Systems

1 Introduction

Motivation-Framework. Large-scale traffic and communication networks, like e.g. the internet, telephone networks, or road traffic systems often lack a central regulation for several reasons: The size of the network may be too large, the networks may be dynamically evolving over time, or the users of the network may be free to act according to their private interest, without regard to the overall performance of the system. Besides the lack of central regulation even cooperation of the users among themselves may be impossible due to the fact that the users may not even know each other. Networks with non-cooperative users have already been studied in the early 1950s in the context of road traffic systems [War52,BMW56]. Recently, this kind of networks has become increasingly important in computer science. Modern computational artifacts, like e.g. the internet, are modeled as communication networks with non-cooperative users.

We survey routing problems in communication networks where n sources of traffic, called users, are going to route their traffics through a shared network. Traffics are routed through links of the network at a certain rate depending on the link, and different users may have different objectives, e.g. speed, quality of service, etc. The users choose routing strategies in order to minimize their private costs in terms of their private objectives without cooperating with other users. Such networks are called non-cooperative networks [KLO95]. A famous example of such a network is the internet. Motivated by non-cooperative systems like the internet, combining ideas from game theory and computer science has become increasingly important [FPS00,JV01,Nis99,NR99,Pap01].

Such an environment, which lacks a central control unit due to its size or operational mode, can be modeled as a non-cooperative game [OR94]. Users selfishly choose their private strategies, which in our environment correspond to paths (or probability distributions over the paths) from their sources to their destinations. When routing their traffics according to the strategies chosen, the users will experience an expected latency caused by the traffics of all users sharing edges. Each user tries to minimize its private cost, expressed in terms of its expected individual latency. This often contradicts the goal of optimizing the social cost which measures the global performance of the whole network. The degradation of the global performance due to the selfish behavior of its users is often termed price of anarchy [Pap01,Rou02] and measured in terms of the coordination ratio. The theory of Nash equilibria [Nas51] provides us with an important solution concept for environments of this kind: a Nash equilibrium is a state of the system such that no user can decrease its individual cost by unilaterally changing its strategy. It has been shown by Nash that a Nash equilibrium exists under fairly broad circumstances.

The concept of Nash equilibria has become an important mathematical tool in analyzing the behavior of selfish users in non-cooperative systems [Pap01]. Many algorithms have been developed to compute a Nash equilibrium in a general game (see [MM96] for an overview). Although the theorem of Nash [Nas51] guarantees the existence of a Nash equilibrium the computational complexity of computing a Nash equilibrium in general games is open even if only $n = 2$ users are involved. Papadimitriou [Pap01] states that due to the guaranteed existence of a solution the problem is unlikely to be $\mathcal{NP}$-hard. The problem becomes even more challenging when global objective functions have to be optimized over the set of all Nash equilibria. In this work we survey results for two different models of non-cooperative communication networks.

In the first model, the network consists of a single source and a single destination which are connected by m parallel links of capacities $c_1, \ldots, c_m$. Users $1, \ldots, n$ are going to selfishly route their traffics $w_1, \ldots, w_n$ from the source to the destination. Traffics are unsplittable. Users may choose mixed strategies to route their traffics. This model has been introduced by Koutsoupias and Papadimitriou [KP99]. We denote it as the KP-model. The individual cost of a user is defined as the maximum expected latency over all links it has chosen with positive probability. Depending on how the latency of a link is defined we distinguish between three variations of the model: In the identical link model all links have equal capacity. In the model of related links the latency for a link j is defined to be the quotient of the sum of the traffics through j and the capacity c_j. In the general case of unrelated links traffic i induces load w_{ij} on link j. In all these models the social cost is defined to be the maximum expected latency on a link, where the expectation is taken over all random choices of the users.

The second model, which we denote by Wardrop-model, has already been studied in the 1950's [War52,BMW56] in the context of road traffic systems. It allows to split traffics into arbitrary pieces. Wardrop [War52] 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 [Bec67]. In this environment unregulated traffic is modeled as network flow. Given an arbitrary network with edge latency functions, equilibrium flows have been classified as flows with all flow paths used between a given source-destination pair having equal latency. Equilibrium flows are optimal solutions to a convex program, if the edge latencies are given by convex functions. A lot of subsequent work (see [RT02, Sec. 1.2] for a brief survey) on this model has been motivated by Braess's Paradox [Bra68]. An equilibrium in this model can be interpreted as a Nash equilibrium in a game with infinitely many users, each carrying an infinitesimal amount of traffic from a source to a destination. The individual cost of a user is defined to be the sum of the edge latencies on a path from the user's source to its destination, the social cost is defined to be the sum over all edge latencies in the network. Inspired by the new interest in the coordination ratio, Roughgarden and Tardos [RT02,Rou01,Rou02] re-investigated the Wardrop-model. We survey some of their results in section 6.

The practical relevance of the Wardrop-model is underpinned by its use by traffic engineers, who utilized equilibria in route-guidance systems to prescribe user behavior. Recent analyses of this framework have been done by Schulz and Stier Moses [SS02] for capacitated networks; algorithms and experimental benchmarking on real world problems are given by Jahn et al. [JMSS02].

The two models differ in terms of their definition of social cost. Routing in the KP-model is equivalent to scheduling n jobs on m parallel machines [GLM$^+$02].

The maximum expected latency is then equivalent to the expected makespan of a schedule. The Wardrop-model has its origins in the definition of road traffic systems. Here the total network load expressed in terms of the sum of edge loads is a natural measure for the network performance. In the KP-model there may be Nash equilibria with different social costs, while in the Wardrop-model the costs of Nash equilibria are equal.

It is well known that in non-cooperative networks, due to the lack of coordination, the users may get to a solution, i.e. a Nash equilibrium, that is sub-optimal in terms of the social cost. Koutsoupias and Papadimitriou [KP99] defined the coordination ratio as the ratio of the social cost of a worst Nash equilibrium and the social cost of the global optimal solution. The coordination ratio is a measure for the price of anarchy.

The well known Braess's Paradox [Bra68] shows that there exist networks, such that strict sub-networks perform better when users are selfish. If the goal is to construct a network in such a way that the coordination ratio of the network is small, an interesting network design problem arises: Given a network and the corresponding routing tasks, determine a set of edges which should be removed from the network to obtain a best possible routing at Nash equilibrium. Such network design problems arise e.g. in the routing of road traffic, when traffic engineers want to determine roads that should be closed, or changed to one-way roads, in order to obtain an optimal traffic flow.

In the case that the users are allowed to randomize their strategies, the set of solutions of the routing problem equals the set of all mixed Nash equilibria. A special class of mixed equilibria is the class of Fully Mixed Nash equilibria. A Nash equilibrium is fully mixed, if every user plays each of its pure strategies with positive probability. Gairing et al. [GLM$^+$02] conjecture that, if a fully mixed Nash equilibrium exists, it is the worst equilibrium with respect to the social cost in the KP-model with related links.

In the case that the users are not allowed to randomize their strategies, the set of solutions of the routing problem consists of all pure Nash equilibria. In this environment the problem of Nashification becomes important. The problem of Nashification is to compute an equilibrium routing from a given non-equilibrium one without increasing the social cost. The intention to centrally nashify a non-equilibrium solution is to provide a routing from which no user has an incentive to deviate.

One way to nashify an assignment is to perform a sequence of greedy selfish steps. A greedy selfish step is a user's change of its current pure strategy to its best pure strategy with respect to the current strategies of all other users. Any sequence of greedy selfish steps leads to a pure Nash equilibrium. However, the length of such a sequence may be exponential in n. We present polynomial time algorithms for the problem of Nashification.

Overview. In section 2 we review basic notations. Then, in section 3 we focus on the sub-model of identical links. In section 4 we turn to the more general model of related links. The most general sub-model of the KP-model is considered in section 5. Finally, in section 6 we review results for the Wardrop-model.

2 Basic Notations

Mathematical Preliminaries. For an integer $i \geq 1$, denote $[i] = \{1, \ldots, i\}$. Denote Γ the *Gamma function*; that is, for any natural number i, $\Gamma(i+1) = i!$, while for any arbitrary real number $x > 0$, $\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$. We use the fact that $\Gamma(x+1) = x \cdot \Gamma(x)$. The Gamma function is invertible; both Γ and its inverse Γ^{-1} are increasing.
Game Theoretic Framework. We consider a *non-cooperative network game* in which each *network user* $i \in [n]$, or *user* for short, wishes to route a particular amount of traffic along (non-fixed) paths from *source* s_i to *destination* t_i.

Denote $\mathcal{P}_i$ as the set of all distinct simple paths from s_i to t_i. Then, a *pure strategy* for user $i \in [n]$ is a path in $\mathcal{P}_i$. A *mixed strategy* for user i is a probability distribution over pure strategies. Each user i wants to minimize its *private cost*, which depends on the latency of the shared edges used by user i to route its traffic. A Nash equilibrium is then a setting where no user can decrease its private cost by unilaterally changing its strategy. The performance of the total network is measured in terms of *social cost*. In this paper we consider two models which are different with respect to the following aspects:

1. The structure of the network.
2. Traffics may be splittable or unsplittable.
3. The definition of the private and social cost functions.

2.1 KP-Model

We first consider a model introduced by Koutsoupias and Papadimitriou [KP99], called *KP-model*. Here, the network consists only of a single source node and a single destination node connected by m parallel links. Traffics are *unsplittable*, that is, a pure strategy of a user is some specific link, and a mixed strategy is a probability distribution over the set of links. The private cost of user i is defined to be the maximum expected latency on a link used by user i. The social cost is the expected maximum latency of a link.
General. Due to the simplicity of the network, a *pure strategy profile* $\mathbf{L}$ is represented by an n-tuple $\langle l_1, l_2, \ldots, l_n \rangle \in [m]^n$; a *mixed strategy profile* $\mathbf{P}$ is represented by an $n \times m$ *probability matrix* of nm probabilities p_{ij}, $i \in [n]$ and $j \in [m]$, where p_{ij} is the probability that user i chooses link j. A mixed strategy profile $\mathbf{F}$ is *fully mixed* if for all users $i \in [n]$ and all links $j \in [m]$, $p_{ij} > 0$. The *support* of the mixed strategy for user $i \in [n]$, denoted *support*(i), is the set of those pure strategies (links) to which i assigns positive probability; so, *support*$(i) = \{j \in [m] \mid p_{ij} > 0\}$. For pure strategies we denote *link*$(i) = l_i$.
System, Models and Cost Measures. Denote w_i the *traffic* of user $i \in [n]$. Define the $n \times 1$ *traffic vector* $\mathbf{w}$ in the natural way. Assume, without loss of generality, that $w_1 \geq w_2 \geq \ldots \geq w_n$, and denote $W = \sum_{i=1}^n w_i$ the *total traffic*. Denote $c_j > 0$ the *capacity* of link $j \in [m]$, representing the rate at which the link processes traffic. In the model of *identical links,* all link capacities are equal. Link capacities may vary arbitrarily in the model of *related links*. Without loss of

generality assume $c_1 \geq \ldots \geq c_m$, and denote $C = \sum_{j=1}^{m} c_j$ the *total capacity.* So, the *latency* for traffic w_i through link j equals $\frac{w_i}{c_j}$. In the model of unrelated links, there exists neither an ordering on the traffics nor on the capacities. We denote w_{ij} the *traffic* of user $i \in [n]$ on link $j \in [m]$. Link capacities are not necessary as a problem input when links are unrelated. However, to obtain common expressions in the following definitions we assume that the link capacities are all equal to 1.

Let $\mathbf{P}$ be an arbitrary mixed strategy profile. The *expected latency* of user i on link j is

$$\lambda_{ij} = \frac{w_{ij} + \sum_{k \in [n], k \neq i} p_{kj} w_{kj}}{c_j}.$$

Denote $\mathsf{IC}(\mathbf{w}, \mathbf{P})$ the *maximum expected individual latency* by

$$\mathsf{IC}(\mathbf{w}, \mathbf{P}) = \max_{i \in [n]} \max_{j \in [m] \mid p_{ij} > 0} \lambda_{ij}.$$

The *minimum expected latency* of user i is $\lambda_i = \min_{j \in [m]} \lambda_{ij}$. The *expected load* Λ_j on link j is the ratio between the expected traffic on link j and the capacity of link j. Thus,

$$\Lambda_j = \frac{\sum_{i \in [n]} p_{ij} w_{ij}}{c_j}.$$

The *maximum expected load* $\Lambda = \max_{j \in [m]} \Lambda_j$ is the maximum (over all links) of the expected load Λ_j on a link j. Associated with a traffic vector $\mathbf{w}$ and a mixed strategy profile $\mathbf{P}$ is the *social cost* [KP99, Section 2], denoted $\mathsf{SC}(\mathbf{w}, \mathbf{P})$, which is the expected maximum latency on a link, where the expectation is taken over all random choices of the users. Thus,

$$\mathsf{SC}(\mathbf{w}, \mathbf{P}) = \sum_{\langle l_1, l_2, \ldots, l_n \rangle \in [m]^n} \left(\prod_{k=1}^{n} p_{k l_k} \cdot \max_{j \in [m]} \frac{\sum_{k : l_k = j} w_{kj}}{c_j} \right).$$

Note that $\mathsf{SC}(\mathbf{w}, \mathbf{P})$ reduces to the maximum latency through a link in the case of pure strategies. Moreover, by definition of the social cost, there always exists a pure strategy profile with minimum social cost. So, the *social optimum* [KP99, Section 2] associated with a traffic vector $\mathbf{w}$, denoted $\mathsf{OPT}(\mathbf{w})$, is the *least possible* maximum (over all links) latency through a link, that is,

$$\mathsf{OPT}(\mathbf{w}) = \min_{\langle l_1, l_2, \ldots, l_n \rangle \in [m]^n} \max_{j \in [m]} \frac{\sum_{k : l_k = j} w_{kj}}{c_j}.$$

Nash Equilibria and Coordination Ratio. Say that a user $i \in [n]$ is *satisfied for the probability matrix* $\mathbf{P}$ if $\lambda_{ij} = \lambda_i$ for all links $j \in$ *support*(i), and $\lambda_{ij} \geq \lambda_i$ for all $j \notin$ *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* [Nas51] iff all users $i \in [n]$ are satisfied for $\mathbf{P}$.

Fix any traffic vector $\mathbf{w}$. A *best (worst) Nash equilibrium* is a Nash equilibrium that minimizes (maximizes) $\mathsf{SC}(\mathbf{w}, \mathbf{P})$. The *best social cost* is the social cost of a best Nash equilibrium. The *worst social cost* is the social cost of a worst Nash equilibrium and is denoted by $\mathsf{WC}(\mathbf{w})$. Gairing et.al. [GLM+02] conjecture that in case of its existence the *fully mixed Nash equilibrium*, which is unique, is the worst Nash equilibrium.

Fully Mixed Nash Equilibrium Conjecture [GLM+02]
Consider the model of arbitrary traffics and related links. Then, for any traffic vector $\mathbf{w}$ *such that the fully mixed Nash equilibrium* $\mathbf{F}$ *exists, and for any Nash equilibrium* $\mathbf{P}$*,* $\mathsf{SC}(\mathbf{w}, \mathbf{P}) \leq \mathsf{SC}(\mathbf{w}, \mathbf{F})$.

The *coordination ratio* [KP99] is the *maximum* of $\mathsf{WC}(\mathbf{w})/\mathsf{OPT}(\mathbf{w})$, over all traffic vectors $\mathbf{w}$. Correspondingly, we denote the maximum of $\mathsf{IC}(\mathbf{w}, \mathbf{P})/\mathsf{OPT}(\mathbf{w})$ the *individual coordination ratio.*

Though a mixed Nash equilibrium always exists, as implied by the fundamental theorem of Nash [Nas51], this is not the case for pure Nash equilibria in general settings. However, for the model under consideration, there is always some pure strategy profile, which fulfills the Nash equilibrium condition. A proof can be given with the help of *sequences of greedy selfish steps.* In a *selfish step,* exactly one unsatisfied user is allowed to change its pure strategy. A selfish step is a *greedy selfish step* if the user chooses its best strategy. Selfish steps do not increase the social cost of the initial pure strategy profile. Starting with any pure strategy profile, every sequence of (greedy) selfish steps eventually ends in a pure Nash equilibrium. Moreover, starting with any pure strategy profile with minimum social cost, this also shows that there always exists a pure Nash equilibrium $\mathbf{L}$ with social cost $\mathsf{SC}(\mathbf{w}, \mathbf{L}) = \mathsf{OPT}(\mathbf{w})$.

Theorem 1 ([FKK+02], Theorem 1). *There exists a pure Nash equilibrium.*

Algorithmic Problems. We list a few algorithmic problems related to Nash equilibria that we consider in this work. The definitions are given in the style of Garey and Johnson [GJ79]. A problem instance is a tuple (n, m, w, c), where n is the number of users (traffics), m is the number of links, $w = (w_{i,j})$ is a $n \times m$ matrix of traffics and $c = (c_j)$ is a vector of m link capacities.

$\mathbf{\Pi_1}$: BEST NASH EQUILIBRIUM SUPPORTS
INSTANCE: A problem instance (n, m, w, c).
OUTPUT: A best Nash equilibrium $\mathbf{P}$.

$\mathbf{\Pi_2}$: WORST NASH EQUILIBRIUM SUPPORTS
INSTANCE: A problem instance (n, m, w, c).
OUTPUT: A worst Nash equilibrium $\mathbf{P}$.

The corresponding problem to compute a worst pure Nash equilibrium is denoted as WORST PURE NASH EQUILIBRIUM SUPPORTS. If, additionally, m is constant the problem is denoted as m-WORST PURE NASH EQUILIBRIUM SUPPORTS.

Π_3: NASH EQUILIBRIUM SOCIAL COST
INSTANCE: A problem instance (n, m, w, c); a *Nash equilibrium* $\mathbf{P}$ for (n, m, w, c).
OUTPUT: The social cost of the Nash equilibrium $\mathbf{P}$.

Π_4: NASHIFY
INSTANCE: A problem instance (n, m, w, c); a pure strategy profile $\mathbf{L}$ for the system of the users; an integer $k > 0$.
QUESTION: Is there a sequence of at most k selfish steps that transforms $\mathbf{L}$ to a (pure) Nash equilibrium?

If k is a constant and not part of the input the corresponding decision problem is denoted as k-NASHIFY.

2.2 Wardrop-Model

In the Wardrop-model traffics are splittable into infinitesimally small pieces. Each of these infinitely many pieces may be viewed as the traffic of a single user who chooses a path from its source to its destination. The problem of routing splittable traffic in a congested network has already been studied since the 1950's [War52,BMW56]. Given a network, rates of traffic between pairs of nodes, and a latency function for each edge, the objective is to route traffic such that the sum of all latencies is minimized. More formally, the model is defined as follows:
General. An instance (G, r, l) for the routing problem consists of a network $G = (V, E)$, a set $r = \{(r_i, s_i, t_i) \in \mathbb{R}_{>0} \times V \times V \mid i \in [k]\}$ of routing tasks, and a set $l = \{l_e \mid e \in E\}$ of edge latency functions. A triple $(r_i, s_i, t_i) \in r$ defines the task to route a traffic of rate r_i from s_i to t_i. Traffics are to be routed via paths in $\mathcal{P}_i$. Let $\mathcal{P} = \bigcup_{i=1}^{k} \mathcal{P}_i$. A traffic flow then is a function $f : \mathcal{P} \to \mathbb{R}_{\geq 0}$. For a fixed flow f the flow f_e along edge $e \in E$ is defined as $f_e = \sum_{P \in \mathcal{P}: e \in P} f_P$. f is a feasible solution for the routing problem instance (G, r, l) iff $\sum_{P \in \mathcal{P}_i} f_P = r_i$ for all $i \in [k]$.
System, Models and Cost Measures. The latency of an edge e is given by a non-negative, differentiable and non-decreasing function $l_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. For a fixed flow f the latency of an edge $e \in E$ is defined as $c_e(f_e) = l_e(f_e) f_e$, the product of $e's$ latency when routing traffic f_e times the traffic f_e itself. The latency $c_P(f) = \sum_{e \in P} c_e(f_e)$ of a path P is defined to be the sum of the edge latencies on P.

The social cost $C(f)$ of a flow f is defined as the sum of the edge latencies

$$C(f) = \sum_{e \in E} c_e(f_e) = \sum_{P \in \mathcal{P}} l_P(f) f_P,$$

where $l_P(f) = \sum_{e \in P} l_e(f_e)$.
Nash Equilibria and Coordination Ratio. By definition, a flow f is a Nash equilibrium, if any arbitrarily small amount of flow routed from, say, s_i to t_i via a path $P_1 \in \mathcal{P}_i$ cannot improve the latency experienced on P_1 by switching to a path $P_2 \in \mathcal{P}_i$. More formally, from [DS69] we obtain the following definition:

A flow f in G is a Nash equilibrium, occasionally called a Wardrop equilibrium [DS69], if for all $P_1, P_2 \in \mathcal{P}$ and all $\delta \in [0, f_{P_1}]$, we have $l_{P_1}(f) \leq l_{P_2}(\tilde{f})$, where

$$\tilde{f}_P = \begin{cases} f_P - \delta \text{ if } P = P_1 \\ f_P + \delta \text{ if } P = P_2 \\ f_P \quad\quad \text{if } P \notin \{P_1, P_2\}. \end{cases}$$

In contrast to the notion of mixed Nash equilibria in the KP-model, here flows may be split among different paths. Consequently the social cost $C(f)$ of a flow as defined above is not a measure of expected cost. Nash equilibria in this model are pure, however, there are infinitely many users, each carrying an infinitesimal amount of the overall traffic. As for the KP-model, the *coordination ratio* is defined by the ratio of the social cost of the worst Nash equilibrium and the minimum possible total latency.

3 Routing in the KP-Model with Identical Links

In this section we consider a simple model for our routing game, namely the KP-model where all links have the same capacity. We distinguish between *pure Nash equilibria*, where each user chooses a single link as its strategy, and *mixed Nash equilibria*. Here the strategy of a user is a probability distribution over the links.

3.1 Pure Nash Equilibria

Fotakis *et al.* [FKK+02, Theorem 3] show that computing a pure Nash equilibrium with minimum social cost is $\mathcal{NP}$-hard even in the model of identical links. Since this problem can be formulated as an integer program, it follows that it is $\mathcal{NP}$-complete. However, computing some pure Nash equilibrium $\mathbf{L}$ can be done with help of the LPT-algorithm introduced by Graham [Gra69], using polynomial time and yielding

$$\mathsf{SC}(\mathbf{w}, \mathbf{L}) \leq \left(\frac{4}{3} - \frac{1}{3m}\right) \mathsf{OPT}(\mathbf{w}).$$

Another way to compute a pure Nash equilibrium is, starting from any pure strategy profile, to allow the users to perform greedy selfish steps. Every sequence of greedy selfish steps eventually yields a pure Nash equilibrium. We now give bounds on the maximum length of a sequence of greedy selfish steps, that show that such a sequence may be of exponential length.

Theorem 2 ([EDKM03]). *There exists an instance of n users on m identical links for which the maximum length of a sequence of greedy selfish steps is at least*

$$\frac{\left(\frac{n}{m-1}\right)^{m-1}}{2(m-1)!}.$$

An upper bound on the number of greedy selfish steps is given by

Theorem 3 ([FGL$^+$03]). *For any instance with n users on identical links, the length of any sequence of greedy selfish steps is at most $2^n - 1$.*

Instead of the maximum length one may ask about the minimum length of a sequence of greedy selfish step. In particular one may consider whether a given pure strategy profile can be transformed into a pure Nash equilibrium with at most k selfish steps. This problem is called NASHIFY and was shown to be $\mathcal{NP}$-complete.

Theorem 4 ([GLM$^+$02]). NASHIFY *is $\mathcal{NP}$-complete on identical links even if $m = 2$.*

The proof relies on a reduction from PARTITION. This reduction implies that NASHIFY is $\mathcal{NP}$-complete in the strong sense (cf. [GJ79, Section 4.2]) if m is part of the input. Thus, there is no pseudo-polynomial-time algorithm for NASHIFY (unless $\mathcal{P} = \mathcal{NP}$). In contrast, there is a natural pseudo-polynomial-time algorithm for k-NASHIFY, which exhaustively searches all sequences of k selfish steps; since a selfish step involves a (unsatisfied) user and a link for a total of mn choices, the running time of such an algorithm is $\Theta((mn)^k)$.

In order to compute a Nash equilibrium by using sequences of greedy selfish steps, we can use two types of sequences of polynomial length introduced by Even-Dar *et al.* [EDKM03].

Theorem 5 ([EDKM03]). *For FIFO and Random strategy, the length of a sequence of greedy selfish steps is at most $n(n+1)/2$ before reaching a Nash equilibrium.*

We continue to present algorithm NashifyIdentical that solves NASHIFY when n selfish steps are allowed. A crucial observation is

Lemma 1 ([GLM$^+$02,EDKM03]). *A greedy selfish step of an unsatisfied user i with traffic w_i makes no user k with traffic $w_k \geq w_i$ unsatisfied.*

NashifyIdentical sorts the user traffics in non-increasing order so that $w_1 \geq \ldots \geq w_n$. Then for each user $i := 1$ to n, it removes user i from the link it is currently assigned to, it finds the link j with the minimum latency, and it reassigns user i to the link j. We prove:

Theorem 6 ([GLM$^+$02,EDKM03]). *Let $\mathbf{L} = \langle l_1, \ldots, l_n \rangle$ be a pure strategy profile for n users with traffics $w_1, ..., w_n$ on m identical links with social cost $\mathsf{SC}(\mathbf{w}, \mathbf{L})$. Then algorithm* NashifyIdentical *computes a Nash equilibrium from $\mathbf{L}$ with social cost $\leq \mathsf{SC}(\mathbf{w}, \mathbf{L})$ using $O(n \log n)$ time.*

The proof relies on Lemma 1 and the usage of appropriate data structures. Running the PTAS of Hochbaum and Shmoys [HS87] for scheduling *n jobs* on m identical *machines* yields a pure strategy profile $\mathbf{L}$ such that $\mathsf{SC}(\mathbf{w}, \mathbf{L}) \leq$

$(1+\varepsilon)\mathsf{OPT}(\mathbf{w})$. On the other hand, applying NashifyIdentical to $\mathbf{L}$ yields a Nash equilibrium $\mathbf{L}'$ such that $\mathsf{SC}(\mathbf{w},\mathbf{L}') \leq \mathsf{SC}(\mathbf{w},\mathbf{L})$. Thus, $\mathsf{SC}(\mathbf{w},\mathbf{L}') \leq (1+\varepsilon)\mathsf{OPT}(\mathbf{w})$. Since also $\mathsf{OPT}(\mathbf{w}) \leq \mathsf{SC}(\mathbf{w},\mathbf{L}')$, it follows that:

Theorem 7 ([GLM+02]). *There exists a* PTAS *for* BEST PURE NASH EQUILIBRIUM *for the model of identical links.*

After studying best pure Nash equilibria we now turn our attention to worst pure Nash equilibria. We first show:

Theorem 8 ([SV01,GLM+02]). *Fix any traffic vector* $\mathbf{w}$ *and pure Nash equilibrium* $\mathbf{L}$. *Then,*

$$\frac{\mathsf{SC}(\mathbf{w},\mathbf{L})}{\mathsf{OPT}(\mathbf{w})} \leq 2 - \frac{2}{m+1}.$$

Furthermore, this upper bound is tight.

Theorem 8 shows that the social cost of any Nash equilibrium is at most the factor $2-\frac{2}{m+1}$ away from the social cost of an optimal Nash equilibrium. This also implies, that every Nash equilibrium approximates the social cost of the worst Nash equilibrium within this factor. We now establish, that approximating a worst Nash equilibrium with a better guaranty is $\mathcal{NP}$-hard.

Theorem 9 ([GLM+02]). *It is* $\mathcal{NP}$*-hard to find a Nash equilibrium* $\mathbf{L}$ *with*

$$\frac{\mathsf{WC}(\mathbf{w})}{\mathsf{SC}(\mathbf{w},\mathbf{L}))} < 2 - \frac{2}{m+1} - \varepsilon \quad \text{for any } \varepsilon > 0.$$

It is $\mathcal{NP}$*-hard in the strong sense if the number of links* m *is part of the input.*

The proof of Theorem 9 is based on a reduction from PARTITION. Since WORST CASE PURE NASH EQUILIBRIUM is $\mathcal{NP}$-hard in the strong sense [FKK+02], there exists no pseudo-polynomial algorithm to solve WORST CASE PURE NASH EQUILIBRIUM. However, for a constant number of links m such an algorithm exists [GLM+02].

3.2 Mixed Nash Equilibria

For a mixed Nash equilibrium it is hard even to compute its social cost.

Theorem 10 ([FKK+02]). NASH EQUILIBRIUM SOCIAL COST *is* $\#\mathcal{P}$*-complete even in the model of identical links.*

However, Fotakis *et al.* [FKK+02] show that there exists a fully polynomial, randomized approximation scheme for NASH EQUILIBRIUM SOCIAL COST.

The *coordination ratio* was introduced by Koutsoupias and Papadimitriou [KP99]. They provided a lower bound of $\Omega(\log m/\log\log m)$. This result can be tightened as follows:

Theorem 11 ([KMS02,CV02]). *For m identical links the worst-case coordination ratio is at most*

$$\Gamma^{-1}(m) + \Theta(1) = \frac{\log m}{\log \log m} \cdot (1 + o(1)).$$

Together with the lower bound from [KP99], this bound is tight up to an additive constant.

In the remainder of this section, we consider fully mixed Nash equilibria. Mavronicolas and Spirakis [MS01] showed that in the model of identical links there always exists a unique fully mixed Nash equilibrium.

Lemma 2 ([MS01]). *There is a unique fully mixed Nash equilibrium* $\mathbf{P}$ *with* $p_{ij} = \frac{1}{m}$ *for any user* $i \in [n]$ *and link* $j \in [m]$.

Since we know that in our model the fully mixed Nash equilibrium always exists, we can compare it to a worst Nash equilibrium. The following theorem provides evidence for the *Fully Mixed Nash Equilibrium Conjecture*.

Theorem 12 ([LMM⁺03]). *Consider the model of identical traffics and identical links, and assume that* $m = 2$ *and* n *is even. Then, the* FMNE *Conjecture is valid.*

4 Routing in the KP-Model with Related Links

In this section we are engaged in the two node routing network with related links. Often the term *uniform links* is used to refer to this model in literature. Subsection 4.1 deals with pure Nash equilibria only, whereas the results quoted in subsection 4.2 hold for general (i.e. mixed) Nash equilibria. Subsection 4.3 concentrates on the fully mixed Nash equilibrium.

4.1 Pure Nash Equilibria

In [FKK⁺02] it was shown that the LPT algorithm, which was first explored by Graham [Gra69], can be used to compute some pure Nash equilibrium. Its coordination ratio lies between 1.52 and 1.67 [Fri87]. Keep in mind, that the exact computation of the best Nash equilibrium is $\mathcal{NP}$-complete in the strong sense (see section 3). We now describe another approach to approximate the best Nash equilibrium with arbitrary but constant precision.

A polynomial time algorithm for Nashification. We call the process of converting a given pure strategy profile for related links into a Nash equilibrium without increasing the social cost *Nashification*.

A simple Nashification approach is to perform a greedy selfish step for any user which can improve by this as long as such a user exists. Unfortunately, this can lead to an exponential number of steps, even on identical links (see section 3).

```
NashifyRelated
Input: n users with traffics w_1 ≥ ··· ≥ w_n
       m links with capacities c_1 ≥ ··· ≥ c_m
       Assignment of users to links
Output: Assignment of users to links with less
        or equal maximum latency, which is a NE
{
    // phase 1:
    i := n;  S := {n};
    while i ≥ 1 {
        move user i to link with highest possible index
            without increasing overall maximum latency;
        if i was moved or i ∈ S or link(i) ≤ link(i + 1)
        then S := S ∪ {i};  i := i − 1;
        else {
            move user i to link with smallest possible index
                without increasing overall maximum latency;
            if i was moved then S := S ∪ {i};  i := n;
            else break;
        }
    }
    // phase 2:
    while ∃i ∈ S {
        make greedy selfish step for user i = min(S);
        S := S\{i};
    }
}
```

Fig. 1. NashifyRelated: Converts any assignment into a Nash equilibrium

In Figure 4.1 we present the algorithm NashifyRelated which nashifies any pure routing by a polynomial number of (not necessarily selfish) moves without increasing the maximum latency. A crucial observation for proving the correctness of the algorithm is stated in Lemma 3, which is a generalization of Lemma 1. It shows that greedy-moving a user from its current link to a non-slower link can only make users of smaller size unsatisfied.

Lemma 3 ([FGL$^+$03]). *If user i with traffic w_i performs a greedy selfish step from link j to link k with $c_j \leq c_k$, then no user s with traffic $w_s \geq w_i$ becomes unsatisfied.*

NashifyRelated works in two phases. At every time $link(i)$ denotes the link user i is currently assigned to. The main idea is to fill up slow links with users with small traffics as close to the maximum latency as possible in the first phase (but without increasing the maximum latency) and to perform greedy selfish steps for unsatisfied users in the second phase. During the first phase, set S is used to collect all users that have already been considered by the algorithm. No user is deleted from S in this phase. Throughout the whole algorithm, each user in S is located on a link with non-greater index than any smaller user in S. In other words, the smaller the traffic of a user in S, the slower the link it is assigned to. We may start with $S = \{n\}$, because the above property is trivially fulfilled if S contains only one user. When no further user can be added to S, the first phase terminates. In the second phase we successively perform greedy selfish steps for all unsatisfied users, starting with the largest one. That is, we move each user, who can improve by changing its link, to its best link with respect to the current situation. Because of the special conditions that have been established by phase 1, and by Lemma 3, these greedy selfish steps do not cause other users with larger traffics to become unsatisfied.

Note that during phase 1 the algorithm does not necessarily perform selfish steps. The social cost, however, cannot be increased due to the constraints of the move commands in phase 1.

The following lemma formalizes the above mentioned conditions which are established by phase 1.

Lemma 4 ([FGL+03]). *After phase 1 the following holds:*

(1) All unsatisfied users are in S.
(2) $S = \{n, (n-1), \ldots, (n+1-|S|)\}$, that is, S contains the $|S|$ users with smallest traffics.
(3) $i, i+1 \in S \Rightarrow link(i) \leq link(i+1)$.
(4) Every user $i \in S$ can only improve by moving to a link with smaller index.

Each of the properties in Lemma 4 also holds after each run of the loop in phase 2 ([FGL+03]). In particular, S contains all unsatisfied users (property (1)). But, as S is empty when the algorithm terminates, this implies, that there are no unsatisfied users, i.e., the new assignment, defined by $link(\cdot)$, is a Nash equilibrium.

Implementing the algorithm in a proper way yields a running time as stated in the following theorem.

Theorem 13 ([FGL+03]). *Given any pure strategy profile for the model of related links, algorithm* NashifyRelated *computes a Nash equilibrium with non-increased social cost, performing at most $(m+1)n$ moves in sequential running time $\mathcal{O}(m^2 n)$.*

Combining any approximation algorithm for the computation of good routings with the algorithm NashifyRelated yields a method for approximating the best Nash equilibrium. Particularly, using the PTAS for the Scheduling Problem from Hochbaum and Shmoys [HS88], we get:

Corollary 1. *There is a PTAS for approximating a best pure Nash equilibrium.*

We cannot expect to find an FPTAS for this problem, since the exact computation of the best Nash equilibrium is $\mathcal{NP}$-complete in the strong sense [FKK+02].

Coordination Ratio

Theorem 14 ([CV02]). *The coordination ratio for pure Nash equilibria on m related links with capacities $c_1 \geq \cdots \geq c_m$ is bounded from above by*

$$\Gamma^{-1}(m) + 1 = \frac{\log m}{\log \log m}(1 + o(1)) \quad \textit{as well as} \quad O\left(\log\left(\frac{c_1}{c_m}\right)\right).$$

The upper bound in Theorem 14 can be improved to $\Gamma^{-1}(m)$ [FGL+03]. The following Example shows that this improved bound is asymptotically tight.

Example 1 ([FGL+03]) *Let $k \in \mathbb{N}$, and consider the following instance with k different classes of users:*

- *Class U_1:* $|U_1| = k$ *users with traffics* 2^{k-1}
- *Class U_i:* $|U_i| = 2^{i-1} \cdot (k-1) \prod_{j=1,\ldots,i-1}(k-j)$ *users with traffics* 2^{k-i} *for all* $2 \leq i \leq k$.

In the same way we define $k+1$ different classes of links:

- *Class P_0:* *One link with capacity* 2^{k-1}.
- *Class P_1:* $|P_1| = |U_1| - 1$ *links with capacity* 2^{k-1}.
- *Class P_i:* $|P_i| = |U_i|$ *links with capacity* 2^{k-i} *for all* $2 \leq i \leq k$.

Consider the following assignment:

- *Class P_0:* *All users in U_1 are assigned to this link.*
- *Class P_i:* *On each link in P_i there are $2(k-i)$ users from U_{i+1}, respectively, for all $1 \leq i \leq k-1$.*
- *Class P_k:* *The links from P_k remain empty.*

The above assignment is a pure Nash equilibrium $\mathbf{L}$ with social cost $\mathsf{SC}(\mathbf{w}, \mathbf{L}) = k$ and $\mathsf{OPT}(\mathbf{w}) = 1$.

Lemma 5 ([FGL+03]). *For each $k \in \mathbb{N}$ there exists an instance with a pure Nash equilibrium $\mathbf{L}$ with*

$$k = \frac{\mathsf{SC}(\mathbf{w}, \mathbf{L})}{\mathsf{OPT}(\mathbf{w})} \geq \Gamma^{-1}(m) \cdot (1 + o(1)).$$

4.2 Mixed Nash Equilibria

In this subsection we state upper bounds on the coordination ratio and on the individual coordination ratio for mixed Nash equilibria on related links.

Theorem 15 ([CV02]). *The coordination ratio for m related parallel links is*

$$\Theta\left(\min\left\{\frac{\log m}{\log\log\log m}, \frac{\log m}{\log\left(\frac{\log m}{\log c_1/c_m}\right)}\right\}\right).$$

Whereas the above theorem bounds the coordination ratio depending only on the number of links m and the relation between the fastest and slowest link, we now introduce a structural parameter p, which can yield better bounds on the individual coordination ratio (and the coordination ratio for pure Nash equilibria). We denote $\mathcal{M}_1 = \{j \in [m] \mid w_1 \leq c_j \cdot \mathsf{OPT}(\mathbf{w})\}$ and $p = \sum_{j \in \mathcal{M}_1} c_j / C$. In other words, p is the ratio between the sum of link capacities of links to which the largest traffic can be assigned causing latency at most $\mathsf{OPT}(\mathbf{w})$ and the sum of all link capacities. With the help of p we are able to prove an upper bound on the individual coordination ratio.

Theorem 16 ([FGL⁺03]). *For any mixed Nash equilibrium* $\mathbf{P}$ *the ratio between the maximum expected individual latency* $\mathsf{IC}(\mathbf{w},\mathbf{P}) = \max_{i\in[n]} \lambda_i$ *and* $\mathsf{OPT}(\mathbf{w})$ *for the model of related links is bounded by*

$$\frac{\mathsf{IC}(\mathbf{w},\mathbf{P})}{\mathsf{OPT}(\mathbf{w})} < \begin{cases} \frac{3}{2} + \sqrt{\frac{1}{p} - \frac{3}{4}} & \text{if } \frac{1}{3} \leq p \leq 1, \\ 2 + \sqrt[3]{\frac{1}{p} - 2} & \text{if } \frac{1}{37} \leq p < \frac{1}{3}, \\ \Gamma^{-1}\left(\frac{1}{p}\right) & \text{if } p < \frac{1}{37}. \end{cases}$$

Since $\frac{w_1}{c_1} \leq \mathsf{OPT}(\mathbf{w})$, we have $p \geq \frac{c_1}{C} \geq \frac{1}{m}$. Furthermore, $\mathsf{IC}(\mathbf{w},\mathbf{P}) \geq \Lambda$ holds for every assignment. Thus, from Theorem 16 we can derive an upper bound of $\Gamma^{-1}(m)\mathsf{OPT}(\mathbf{w})$ on both the maximum expected load and the maximum expected individual latency. This leads to an improvement of the upper bound on the coordination ratio for pure Nash equilibria [CV02].

From Example 1 the following lower bound on the coordination ratio subject to p can be obtained.

Lemma 6 ([FGL⁺03]). *For each* $k \in \mathbb{N}$ *there exists an instance with a pure Nash equilibrium* $\mathbf{L}$ *with*

$$k = \frac{\mathsf{SC}(\mathbf{w},\mathbf{L})}{\mathsf{OPT}(\mathbf{w})} \geq \Gamma^{-1}\left(\frac{1}{3p}\right).$$

We can also prove $k \geq \Gamma^{-1}(\frac{1}{p}) - 1$. This shows that the generalized upper bound is tight up to an additive constant for **all** m whereas Lemma 5 shows tightness of $\Gamma^{-1}(m)$ up to an additive constant only for **large** m.

We conclude this section by giving another upper bound on the maximal expected individual latency of a mixed Nash equilibrium, which only depends on the number of links m. The same bound also applies to the social cost of a pure Nash equilibrium.

Theorem 17 ([FGL⁺03]). *For any mixed Nash equilibrium* $\mathbf{P}$ *on* m *related links, the maximum expected individual cost is bounded by*

$$\mathsf{IC}(\mathbf{w},\mathbf{P}) \leq \frac{1+\sqrt{4m-3}}{2}\mathsf{OPT}(\mathbf{w}).$$

This bound is tight if and only if $m \leq 5$. *Only if* $m \leq 3$, *there is a pure Nash equilibrium matching the bound. For* $m \geq 2$, *there is no fully mixed Nash equilibrium matching the bound.*

That $\frac{1+\sqrt{4m-3}}{2}$ is an upper bound on the coordination ratio for pure Nash equilibria, which is an implication of Theorem 17, can also be obtained from Cho and Sahni [CS80] and Schuurman and Vredeveld [SV01], where jump optimal schedules are considered. A schedule is said to be jump optimal if no user on a link with maximum load can improve by moving to another processor. Obviously, the set of pure Nash equilibria is a subset of the set of jump optimal schedules. Thus, the strict upper bound of $\frac{1+\sqrt{4m-3}}{2}$ on the ratio between best and worst makespan of jump optimal schedules [CS80,SV01] also holds for pure Nash equilibria. This bound is not asymptotically tight, but for small numbers of links ($m \leq 19$) better than the asymptotically tight bound $\Gamma^{-1}(m)$.

4.3 Fully Mixed Nash Equilibria

A *fully mixed* Nash equilibrium is a special Nash equilibrium, where all probabilities p_{ij} for a user i to choose link j are strictly positive. Such a Nash equilibrium does not always exist, but if it exists, it is unique and can be efficiently computed.

Theorem 18 ([MS01]). *For the model of related links there is a fully mixed Nash equilibrium* $\mathbf{F}$ *for a traffic vector* $\mathbf{w}$ *if and only if*

$$\left(1-\frac{mc_j}{C}\right)\cdot\left(1-\frac{W}{(n-1)w_i}\right)+\frac{c_j}{C}\in(0,1)\quad \forall i\in[n], j\in[m].$$

If $\mathbf{F}$ *exists,* $\mathbf{F}$ *is unique and has associated Nash probabilities*

$$p_{ij}=\left(1-\frac{mc_j}{C}\right)\cdot\left(1-\frac{W}{(n-1)w_i}\right)+\frac{c_j}{C},$$

for any user $i\in[n]$ *and link* $j\in[m]$.

Gairing *et al.* [GLM+02] conjecture that a fully mixed Nash equilibrium (FMNE) is the worst-case Nash equilibrium (the one with highest social cost) among all Nash equilibria for the same instance. So far, this conjecture is proved to hold in special cases (Theorem 20) and to hold up to a constant multiplicative factor in the general related links case (Theorem 19).

Theorem 19 ([FKK+02]). *Consider an instance with identical traffics for the model of related links for which the fully mixed Nash equilibrium exists. Then the social cost of the worst mixed Nash equilibrium is at most 49.02 times the social cost of any fully mixed Nash equilibrium.*

Theorem 20 ([LMM+03]). *For the model of related links and two identical traffics the* FMNE *Conjecture holds.*

In contrast to the yet unproved claim of the FMNE Conjecture (that the fully mixed Nash equilibrium has the worst social cost), Lemma 7 shows, that each user indeed experiences the worst individual cost in the fully mixed Nash equilibrium. This implies, that the social cost of a pure Nash equilibrium is bounded from above by the social cost of the fully mixed Nash equilibrium.

Lemma 7 **([GLM+02]).** *Fix any traffic vector* $\mathbf{w}$*, mixed Nash equilibrium* $\mathbf{P}$ *and user* i*. Then,* $\lambda_i(\mathbf{w}, \mathbf{P}) \leq \lambda_i(\mathbf{w}, \mathbf{F})$.

Theorem 21 **([GLM+02]).** *Fix any traffic vector* $\mathbf{w}$ *and pure Nash equilibrium* $\mathbf{L}$*. Then,* $\mathsf{SC}(\mathbf{w}, \mathbf{L}) \leq \mathsf{SC}(\mathbf{w}, \mathbf{F})$.

5 Routing in the KP-Model with Unrelated Links

Up to now only little attention has been payed to Nash equilibria in the model of unrelated links. Similar to the model of identical links, a pure Nash equilibrium in the model of unrelated links can be computed by performing sequences of (greedy) selfish steps. However, up to now it is unknown whether, starting with any pure strategy profile, there always exists a sequence of polynomial length ending in a pure Nash equilibrium. A trivial upper bound on the number of selfish steps before reaching a pure Nash equilibrium is m^n.

Starting with a pure strategy profile with minimal social cost, the convergence of any sequence of (greedy) selfish steps implies that there always exists a pure Nash equilibrium with optimal social cost. However, we can not hope to approximate a best pure Nash equilibrium within factor $\frac{3}{2}$ since *Minimum Multiprocessor Scheduling* on unrelated processors is not approximable within a factor $\frac{3}{2} - \varepsilon$ for any $\varepsilon > 0$ [LST87].

Consider the case $n \leq m$, that is, there are at most as many users as links. Then, the minimum expected latency of any user i in a pure Nash equilibrium $\mathbf{L}$ is smaller than the minimum expected latency of user i in the fully mixed Nash equilibrium $\mathbf{F}$.

Proposition 1 **([LMM+03]).** *For the model of unrelated links, let* $\mathbf{w} = (w_{ij})$ *be a traffic matrix such that a fully mixed Nash equilibrium* $\mathbf{F}$ *exists, and let* $\mathbf{L}$ *be a pure Nash equilibrium. Let* $n \leq m$*. Then, for every user* i*,* $\lambda_i(\mathbf{L}) < \lambda_i(\mathbf{F})$.

Clearly, the social cost of any pure Nash equilibrium $\mathbf{L}$ is equal to the maximum of the expected latencies, while the social cost of a fully mixed Nash equilibrium $\mathbf{F}$ is at least the expected latency of any user. Hence, Proposition 1 implies that for $n \leq m$ the social cost of every pure Nash equilibrium $\mathbf{L}$ is at most the social cost of the fully mixed Nash equilibrium $\mathbf{F}$.

Theorem 22 **([LMM+03]).** *For the model of unrelated links, let* $\mathbf{w} = (w_{ij})$ *be a traffic matrix such that a fully mixed Nash equilibrium* $\mathbf{F}$ *exists, and let* $\mathbf{L}$ *be a pure Nash equilibrium. Let* $n \leq m$*. Then,* $\mathsf{SC}(\mathbf{w}, \mathbf{L}) \leq \mathsf{SC}(\mathbf{w}, \mathbf{F})$.

In case of $n = 2$, the minimum expected latency of user $i \in [2]$ in any mixed Nash equilibrium $\mathbf{P}$ is also bounded by its minimum expected latency in the fully mixed Nash equilibrium $\mathbf{F}$.

Proposition 2 **([LMM+03]).** *For the model of unrelated links, let* $\mathbf{w} = (w_{ij})$ *be a traffic matrix such that a fully mixed Nash equilibrium* $\mathbf{F}$ *exists, and let* $\mathbf{P}$ *be a Nash equilibrium. If* $n = 2$ *then* $\lambda_i(\mathbf{P}) \leq \lambda_i(\mathbf{F})$ *for every user* $i \in [2]$.

This implies that the FMNE Conjecture holds for $n = m = 2$. However, for the case $n = 3$ and $m = 2$ there exist instances for which the FMNE Conjecture does not hold.

Theorem 23 ([LMM$^+$03]). *For the model of unrelated links, if $n = 2$ and $m = 2$, then the FMNE Conjecture holds. If $n = 3$ and $m = 2$, then the FMNE Conjecture does not hold.*

6 Routing in the Wardrop-Model

The problem of routing splittable traffics in a congested network has already been studied since the 1950's [War52,BMW56]. In contrast to the notion of mixed Nash equilibria in the previous sections, here flows may be split among different paths. The model can be interpreted as a game with infinitely many users each carrying an infinitesimally small amount of traffic. A Nash equilibrium is defined to be a state in which no such user has an incentive to unilaterally deviate from its chosen path. In the Wardrop-model the following characterization of a Nash equilibrium, occasionally called a Wardrop equilibrium, can be obtained:

Theorem 24 ([War52]). *A flow f is a Nash equilibrium iff for every $i \in [k]$ and $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ we have $l_{P_1}(f) \leq l_{P_2}(f)$.*

Thus, in a Nash equilibrium all paths $P \in \mathcal{P}_i$ which are used by user i have equal latency, say $L_i(f)$. Then $C(f) = \sum_{i=1}^{n} L_i(f) r_i$.

The main differences to the KP-model are the following:

a) The social cost of a flow is defined to be the sum of the edge latencies, as opposed to the maximum of the edge latencies in the KP-model.
b) Traffic rates and flows may be split arbitrarily. Equivalently, the Wardrop-model can be viewed as a system of infinitely many users, each controlling an infinitesimally small amount of traffic. In the Wardrop-model the fact that flows are splittable into arbitrarily small pieces guarantees the existence of a Nash equilibrium.
c) Edge latencies are defined by using quite general, but not arbitrary, edge latency functions, giving the latency on an edge per unit of flow. If all edge latency functions $l_e(f_e)$ are of the form $l_e(f_e) = a_e f_e$, the definition of latency coincides with the definition of latency in the KP-model with related links.
d) Due to the differences in the definition of latencies (a) and due to the fact that all Nash equilibria in the Wardrop-model are pure (b), the semantic of the coordination ratios is different in the two models.

In the Wardrop-model optimal flows can be computed as a solution of the non-linear program.

$$\mathcal{NLP}: \quad \min \ \sum_{e \in E} c_e(f_e) \ \text{ s.t.}$$

$$\begin{aligned} \sum_{P \in \mathcal{P}_i} f_P = r_i & \qquad \forall i \in [k] \\ f_e = \sum_{P \in \mathcal{P}: e \in P} f_P & \qquad \forall e \in E \\ f_P \geq 0 & \qquad \forall P \in \mathcal{P} \end{aligned}$$

From this non-linear program the following necessary and sufficient condition for a flow of being optimal can be derived [BMW56] in the case of convex edge latencies $c_e(f_e)$:

Theorem 25 ([BMW56]). *Let $c_e(f_e)$ be convex for all $e \in E$. A flow f is optimal iff for every $i \in [k]$ and $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ we have $c'_{P_1}(f) \leq c'_{P_2}(f)$, where $c'_P(f) = \frac{d}{dx} c_P(f)$ denotes the derivative of $c_P(f)$.*

Note that from the two characterizations of Nash equilibria and optimal flows an optimal flow f for a problem instance (G, r, l) with convex edge latencies $c_e(f_e) = l_e(f_e) f_e$ can be regarded as a flow at Nash equilibrium with respect to the edge latencies $c'_e(f_e) = l_e(f_e) + l'_e(f_e) f_e$. From this it follows that there always exists a Nash equilibrium in the Wardrop-model, and that, if $f, \tilde{f}$ are Nash equilibria then $C(f) = C(\tilde{f})$.

Roughgarden and Tardos proved the following bicriteria bound for the coordination ratio:

Theorem 26 ([RT02]). *Let the edge latency functions $l_e(f_e)$ be continuous and nondecreasing for all $e \in E$. If f is a flow at Nash equilibrium for (G, r, l) and f^* is a feasible flow for $(G, (1+\delta) r, l)$, then $C(f) \leq \frac{1}{\delta} C(f^*)$.*

In particular ($\delta = 1$), the cost of a flow at Nash equilibrium never increases the cost of a feasible flow for the problem instance obtained from doubling the traffic rates.

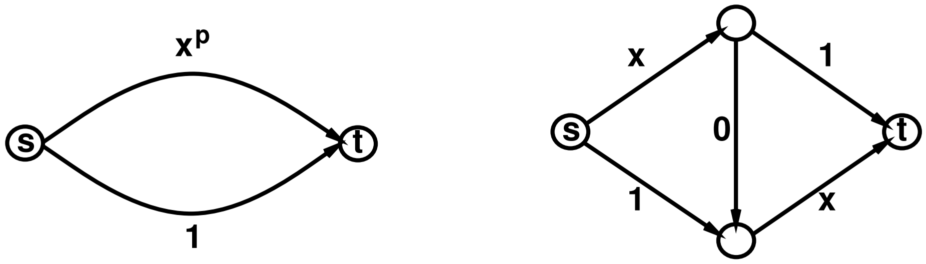

Fig. 2. Pigou's example (left) and Braess's Paradox (right)

Pigou's example [Pig20,BMW56] (see Figure 2) can be used to show that the bound of theorem 26 is tight: In the graph G with two nodes s, t and two edges from s to t with edge latency functions $l_{e_1} = 1$ and $l_{e_2} = x^p$ the optimal solution to route a total of one unit of flow from s to t is to route a fraction of $(p+1)^{-\frac{1}{p}}$ units along e_2 and a flow of $1-(p+1)^{-\frac{1}{p}}$ along edge e_1. The social cost of the optimal solution tends to zero when p tends to infinity. The unique Nash equilibrium in Pigou's example is to route all flow along edge e_2 with social cost 1. For a fixed $\delta > 0$ and any $\varepsilon \in [0, \delta[$, p can be chosen large enough such that $C(f^*) < \delta$ for an optimal flow f^* of $(G, (1 + \delta - \varepsilon), l)$ and such that $C(f^{**})$ is arbitrarily close to δ for an optimal flow f^{**} of $(G, (1+\delta), l)$. Then $\frac{1}{\delta} \cdot C(f^{**})$ is arbitrarily close to $C(f) = 1$.

In the Wardrop-model the coordination ratio cannot be bounded from above by a constant, when arbitrary edge latency functions are allowed. However, when the edge latency functions l_e are restricted to be linear ones, $l_e(f_e) = a_e f_e + b_e$ for some $a_e, b_e \geq 0$, the coordination ratio is bounded from above by $\frac{4}{3}$:

Theorem 27 ([RT02]). *If (G, r, l) has linear edge latency functions $l_e(f_e) = a_e f_e + b_e$ then*

$$\frac{\mathsf{WC}(G, r, l)}{\mathsf{OPT}(G, r, l)} \leq \frac{4}{3}$$

A problem instance based on Braess's Paradox [Bra68] (see Figure 2) is used to show that the bound from the above theorem is tight. Suppose one unit of flow is to be routed from s to t. The globally optimal solution is to route $\frac{1}{2}$ along the upper path and $\frac{1}{2}$ along the lower path, resulting in an overall social cost of $\frac{3}{2}$. The unique Nash equilibrium in this environment is obtained when the total traffic is routed via the edge connecting the two intermediate nodes, resulting in a social cost of 2. The coordination ratio in this example is $\frac{4}{3}$. The interesting fact justifying the title "Paradox" is that after removing the edge with latency 0 the globally optimal solution does not change and the unique Nash equilibrium coincides with the global optimal solution. The network performs better after deletion of an edge.

Motivated by Braess's Paradox Roughgarden [Rou01] studied a natural network design problem:

NETWORK DESIGN
INSTANCE: (G, r, l) with a network $G = (V, E)$, rates r and latencies l.
OUTPUT: A subgraph $H = (V, E_H) \subset G$ such that a flow f at Nash equilibrium for (H, r, l) has minimum social cost over all possible subgraphs H.

In the right graph of figure 2 deleting the edge with latency zero leads to a flow at Nash equilibrium which equals the globally optimal solution.

For linear latency functions it follows from Theorem 27 that the trivial algorithm, which does not remove any edge, is a $\frac{4}{3}$-approximation. Roughgarden proved the following

Theorem 28 ([Rou01]). *If (G, r, l) has linear edge latency functions $l_e(f_e) = a_e f_e + b_e$ then for any $\varepsilon > 0$, there is no $(\frac{4}{3} - \varepsilon)$-approximation algorithm for NETWORK DESIGN unless $\mathcal{P} = \mathcal{NP}$.*

For a more general class of latency functions Roughgarden proved

Theorem 29 ([Rou01]). *If (G, r, l) with $G = (V, E)$ and $|V| = n$ has continuous, nonnegative and nondecreasing edge latency functions then, for any $\varepsilon > 0$, there is no $(\lfloor \frac{n}{2} \rfloor - \varepsilon)$-approximation algorithm for NETWORK DESIGN unless $\mathcal{P} = \mathcal{NP}$.*

Moreover, Roughgarden showed that the approximation quality of the trivial algorithm, which returns the entire graph G, is upper bounded by $\frac{n}{2}$ and that there are instances (generalizations of the Braess graph), for which the trivial algorithm returns an $\lfloor \frac{n}{2} \rfloor$-approximation.

Recently, Roughgarden [Rou02] proved that the worst coordination ratio, that may occur for a network, is equal to the coordination ratio of some simple 2-node 2-edge network. In the example of Pigou the flow at Nash equilibrium incurs 1 unit of social cost, whereas the optimal flow has social cost that tend to zero at a rate of $\Theta(\frac{\log p}{p})$. Thus, as the latency function x^p gets "steeper", the coordination ratio for Pigou's example grows to infinity. The main result of Roughgarden [Rou02] bounds the coordination ratio of problem instance (G, r, l) by a measure of the steepness of the latency functions allowed. This bound is independent of the network topology of G.

A class of nonnegative, differentiable and nondecreasing latency functions $\mathcal{L}$ is called standard, if it contains a non-zero function and if for each $l \in \mathcal{L}$, the function $c_l(x) = x \cdot l(x)$ is convex on $[0, \infty[$. For a non-zero latency function l such that $c_l(x) = x \cdot l(x)$ is convex on $[0, \infty[$ let $c_l'(x) = \frac{d}{dx} c_l(x) = \frac{d}{dx}(x \cdot l(x))$ and define the anarchy value $\alpha(l)$ of l as

$$\alpha(l) = \sup_{r > 0 : l(r) > 0} \frac{1}{\lambda\mu + (1 - \lambda)}$$

where $\lambda \in]0, 1[$ satisfies $c_l'(\lambda r) = l(r)$ and $\mu \in [0, 1]$ is defined as $\mu = l(\lambda r)/l(r)$. Given the anarchy values of single latency functions, the anarchy value $\alpha(\mathcal{L})$ of a standard class of latency functions $\mathcal{L}$ is defined as $\alpha(\mathcal{L}) = \sup_{0 \neq l \in \mathcal{L}} \alpha(l)$.

Let $\rho(G, r, l)$ denote the ratio between the social cost of a Nash equilibrium and the social cost of an optimal flow for (G, r, l) Since in the Wardrop-model all Nash equilibria have equal social cost, $\rho(G, r, l)$ is well defined and expresses the coordination ratio of the problem instance (G, r, l).

That $\alpha(\mathcal{L})$ is always an upper bound on the coordination ratio is stated in the following

Theorem 30 ([Rou02]). *Let $\mathcal{L}$ be a standard class of latency functions with anarchy value $\alpha(\mathcal{L})$. Let (G, r, l) be a problem instance with latency functions drawn from $\mathcal{L}$. Then*

$$\rho(G, r, l) \leq \alpha(\mathcal{L}).$$

Based on Pigou's example Roughgarden constructed instances of the routing problem such that their coordination ratio approaches the anarchy value:

Lemma 8 ([Rou02]). *Given a standard class $\mathcal{L}$ of latency functions containing the constant functions, there are instances (G, r, l), where G is a network of two nodes and two links and latency functions $l \in \mathcal{L}$ such that $\rho(G, r, l)$ is arbitrarily close (from below) to $\alpha(\mathcal{L})$.*

Lemma 8 and Theorem 30 can be combined to achieve

Theorem 31 ([Rou02]). *Let G_2 denote the graph with one source node, one sink node, and two edges directed from source to sink. Let $\mathcal{L}$ be a standard class of latency functions containing the constant functions. Let $\mathcal{I}$ denote the set of all instances (G, r, l) with latency functions in $\mathcal{L}$ and $\mathcal{I}_2 \subseteq \mathcal{I}$ the instances with underlying network G_2, then*

$$\sup_{(G_2,r,l)\in\mathcal{I}_2} \rho(G_2, r, l) = \alpha(\mathcal{L}) = \sup_{(G,r,l)\in\mathcal{I}} \rho(G, r, l)$$

An application of this theory to the set of polynomials taken as a standard class of latency functions results in the following

Theorem 32 ([Rou02]). *Let $\mathcal{L}_p$ be the set of polynomials, let $\mathcal{I}_p$ be the set of instances with latency functions in $\mathcal{L}_p$. Then*

$$\sup_{(G,r,l)\in\mathcal{I}_p} \rho(G, r, l) = \frac{1}{1 - \frac{p}{(p+1)^{\frac{p+1}{p}}}} = \Theta\left(\frac{p}{\log p}\right)$$

Roughgarden generalized the above mentioned results on the coordination ratio for sets of latency functions that do not contain all constant functions and found similar results based on graphs with one source, one sink, and m parallel edges from the source to the sink.

References

[Bec67] M.J. Beckmann. On the theory of traffic flow in networks. *Traffic Quart*, 21:109–116, 1967.

[BHW97] P. Brucker, J. Hurink, and F. Werner. Improving local search heuristics for some scheduling problems. part ii. *Discrete Applied Mathematics*, 72:47–69, 1997.

[BMW56] M. Beckmann, C.B. McGuire, and C.B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[Bra68] D. Braess. Über ein Paradoxon der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.

[CS80] Y. Cho and S. Sahni. Bounds for list schedules on uniform processors. *SIAM Journal on Computing*, 9(1):91–103, 1980.

[CV02] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*, pages 413–420, 2002.

[DS69] S.C. Dafermos and F.T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards, Series B, Vol. 73B, No. 2*, pages 91–118, 1969.

[EDKM03] E. Even-Dar, A. Kesselmann, and Y. Mansour. Convergence time to nash equilibria. In *Proc. of the 30th International Colloquium on Automata, Languages, and Programming (ICALP'03)*, 2003.

[FGL+03] 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 International Colloquium on Automata, Languages, and Programming (ICALP'03)*, 2003.

[FH79] G. Finn and E. Horowitz. A linear time approximation algorithm for multiprocessor scheduling. *BIT*, 19:312–320, 1979.

[FKK+02] 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 International Colloquium on Automata, Languages, and Programming (ICALP'02)*, pages 123–134, 2002.

[FPS00] J. Feigenbaum, C. Papdimitriou, and S. Shenker. Sharing the cost of multicast transmissions. In *Proc. of the 32nd Annual ACM Symposium on the Theory of Computing*, pages 218–227, 2000.

[Fri87] D.K. Friesen. Tighter bounds for lpt scheduling on uniform processors. *SIAM Journal on Computing*, 16(3):554–560, 1987.

[GJ79] M.R. Garey and D.S. Johnson. *Computers and intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman and Company, 1979.

[GLM+02] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Extreme nash equilibria. Technical report, FLAGS-TR-03-10, University of Paderborn, 2002.

[Gra69] R:L. Graham. Bounds on multiprocessing timing anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.

[HM85] A. Haurie and P. Marcotte. On the relatonship between nash-cournot and wardrop equilibria. *Networks*, 15:295–308, 1985.

[HS87] D.S. Hochbaum and D. Shmoys. Using dual approximation algorithms for scheduling problems: Theoretical and practical results. *Journal of the ACM*, 34(1):144–162, 1987.

[HS88] 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.

[JMSS02] O. Jahn, R.H. Möhring, A.S. Schulz, N.E. Stier Moses. System-Optimal Routing of Traffic Flows With User Constraints in Networks With Congestion. MIT Sloan School of Management Working Paper No. 4394-02, 2002.

[JV01] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 364–372, 2001.

[KLO95] Y.A. Korilis, A.A. Lazar, and A. Orda. Architecting noncooperative networks. *IEEE Journal on Selected Areas in Communications*, 13(7):1241–1251, 1995.

[KP99] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, pages 404–413, 1999.

[KMS02] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion In *Proc. of the 9th International Colloquium on Structural Information and Communication Complexity (SIROCCO'02)*, 2002, (accepted for TOCS).

[LMM+03] 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 International Symposium on Mathematical Foundations of Computer Science (MFCS'03)*, 2003.

[LST87] J.K. Lenstra, D.B. Shmoys, and E. Tardos. Approximation algorithms for scheduling unrelated parallel machines. In *Proc. of the 28th Annual Symposium on Foundations of Computer Science (FOCS'87)*, pages 217–224, 1987.

[MM96] R.D. McKelvey and A. McLennan. Computation of equilibria in finite games. In H. Amman, D. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*, 1996.

[MS01] M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 510–519, 2001.

[Nas51] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

[Nis99] N. Nisan. Algorithms for selfish agents. In *Proc. of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, pages 1–15, 1999.

[NR99] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. of the 31st ACM Symposium on Theory of Computing (STOC'99)*, pages 129–140, 1999.

[OR94] M.J. Osborne and A. Rubinstein. *A Course in Game Theory.* The MIT Press, 1994.

[Pap94] C.H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Science*, 48(3):498–532, 1994.

[Pap01] C.H. Papadimitriou. Algorithms, games, and the internet. In *Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 749–753, 2001.

[Pig20] A.C. Pigou. *The economics of welfare.* Macmillan, 1920.

[RT02] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[Rou01] T. Roughgarden. Designing Networks for Selfish Users is Hard. In *Proc. of the 42nd Annual ACM Symposium on Foundations of Computer Science (FOCS'01)*, pages 472–481, 2001.

[Rou02] T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. In *Proc. of the 34th Annual ACM Symposium on Theory of Computing (STOC'02)*, pages 428–437, 2002.

[SS02] A.S. Schulz and N.E. Stier Moses. On The Performance of User Equilibria in Traffic Networks. MIT Sloan School of Management Working Paper No. 4274-02, 2002.

[SV01] P. Schuurman and T. Vredeveld. Performance guarantees of load search for multiprocessor scheduling. In *Proc. of the 8th Conference on Integer Programming and Combinatorial Optimization (IPCO'01)*, pages 370–382, 2001.

[War52] J.G. Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II, Vol. 1*, pages 325–378, 1952.