The Structure and Complexity of Nash Equilibria for a Selfish Routing Game*

Dimitris Fotakis† *Spyros Kontogiannis*‡ *Elias Koutsoupias*§
Marios Mavronicolas¶ *Paul Spirakis*‖

Abstract

In this work, we study the combinatorial structure and the computational complexity of Nash equilibria for a certain game that models *selfish routing* over a network consisting of m parallel *links*. We assume a collection of n *users,* each employing a *mixed strategy,* which is a probability distribution over links, to control the routing of its own assigned *traffic*. In a *Nash equilibrium,* each user selfishly routes its traffic on those links that minimize its *expected latency cost,* given the 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, over all links, *latency* through a link.

We embark on a systematic study of several algorithmic problems related to the computation of Nash equilibria for the selfish routing game we consider. In a nutshell, these problems relate to deciding the existence of a Nash equilibrium, constructing *a* Nash equilibrium, constructing the *worst* Nash equilibrium (the one with maximum social cost), or computing the social cost of a (given) Nash equilibrium. Our work provides a comprehensive collection of efficient algorithms, hardness results (both as $\mathcal{NP}$-completeness and $\#\mathcal{P}$-completeness results), and structural results for these algorithmic problems. Our results span and contrast a wide range of assumptions on the syntax of the Nash equilibria and on the parameters of the system.

*This work has been partially supported by the IST Program of the European Union under contract numbers IST-1999-14186 (ALCOM-FT) and IST-2001-33116 (FLAGS), and by funds from the Joint Program of Scientific and Technological Collaboration between Greece and Cyprus.

†Department of Information and Communication Systems Engineering, University of the Aegean, 83200 Samos, Greece. Email: `fotakis@aegean.gr`. Most of this work was done while the author was a postdoctoral fellow at the Max Planck Institut für Informatik, 66123 Saarbrücken, Germany.

‡Department of Computer Science, University of Ioannina, 45110 Ioannina, Greece, & Research Academic Computer Technology Institute, N. Kazantzaki Str., University Campus, 26500 Patras, Greece. Email: `kontog@cs.uoi.gr`. Most of this work was done while the author was a postdoctoral fellow at the Max Planck Institut für Informatik, 66123 Saarbrücken, Germany.

§Department of Computer Science, University of California at Los Angeles, & Department of Informatics, University of Athens, Greece. Email: `elias@cs.ucla.edu`

¶Department of Computer Science, University of Cyprus, P. O. Box 20537, Nicosia CY-1678, Cyprus. Email: `mavronic@ucy.ac.cy`

‖Department of Computer Engineering and Informatics, University of Patras, Rion, 265 00 Patras, Greece, & Research Academic Computer Technology Institute, N. Kazantzaki Str., University Campus, 26500 Patras, Greece. Email: `spirakis@cti.gr`

1 Introduction

1.1 Motivation-Framework

Nash equilibrium [25] (see also [26] for a "popular science" style discussion) is arguably the most important solution concept in Game Theory [27]. It may be viewed to represent a steady state of the play of a strategic game in which each player holds an accurate opinion about the (expected) behavior of other players and acts rationally. Despite the apparent simplicity of the concept, computation of Nash equilibria in finite games has been long observed to be difficult (cf. [19, 35]); in fact, it is arguably one of the few, most important algorithmic problems for which no polynomial-time algorithms are known. Indeed, Papadimitriou [30, p. 1] (see also [31]) actively advocates the problem of computing Nash equilibria as one of the most significant open problems in Theoretical Computer Science today:

> "The Nash equilibrium (definition omitted here) is the predominant concept of rationality in Game Theory; it is also a most fundamental computational problem whose complexity is wide open: Is there a polynomial algorithm which, given a two-person game with a finite strategy space, computes a mixed Nash equilibrium? Together with factoring, *the complexity of finding a Nash equilibrium is in my opinion the most important concrete open question on the boundary of $\mathcal{P}$ today.*"

In this work, we embark on a systematic study of the computational complexity of Nash equilibria in the context of a simple *selfish routing* game, originally introduced by Koutsoupias and Papadimitriou [16], that we describe here. We assume a collection of n users, each employing a *mixed strategy,* which is a probability distribution over m parallel *links,* to control the shipping of its own assigned *traffic*. For each link, a *capacity* specifies the rate at which the link processes traffic. In a Nash equilibrium, each user selfishly routes its traffic on those links that minimize its *expected latency cost,* given the network congestion caused by the other users. A user's *support* is the set of those links on which it may ship its traffic with non-zero probability. The *social cost* of a Nash equilibrium is the expectation. over all random choices of the users, of the maximum, over all links, *latency* through a link.

We are interested in algorithmic problems related to the computation of Nash equilibria for the selfish routing game we consider. More specifically, we aim at determining the computational complexity of the following prototype problems, assuming that users' traffics and links' capacities are given:

- Given users' supports, decide whether there exists a Nash equilibrium; if so, determine the corresponding users' (mixed) strategies. (This is an existence and computation problem.)
- Decide whether there exists a Nash equilibrium; if so, determine the corresponding users' supports and (mixed) strategies. (This is an existence and computation problem.)
- Determine the supports of the *worst* (or the *best*) Nash equilibrium. (These are optimization problems.)
- Given a Nash equilibrium, determine its social cost. (This turns out to be a hard counting problem (cf. [36]).)

Our study distinguishes between *pure* Nash equilibria, where each user chooses exactly one link (with probability one), and *mixed* Nash equilibria, where the choices of each user are modeled by a probability distribution over links. We also distinguish in some cases between models of *uniform capacities,* where all link capacities are equal, and of *arbitrary capacities*; also, we do so between models of *identical traffics,* where all user traffics are equal, and of *arbitrary traffics*.

1.2 Contribution

We start with pure Nash equilibria. By the linearity of the expected latency cost functions we consider, the celebrated result of Nash [25] on the existence of Nash equilibria, which follows from Kakutani's Fixed Point Theorem [8] [1], assures that a *mixed,* but not necessarily pure, Nash equilibrium always exists. The first result (Theorem 3.1), remarked by Kurt Mehlhorn, establishes that a pure Nash equilibrium always exists. The proof argues (by contradiction) that the *lexicographically minimum* sorted vector of expected latencies corresponds, indeed, to a Nash equilibrium. The proof itself is *inefficient* (in the sense of Papadimitriou [29]) in that it does not lead to an *efficient* algorithm for constructing a pure Nash equilibrium: one would apparently have to examine all expected latency vectors (and there are exponentially many of them, as many as pure strategies) to choose the lexicographically minimum one. However, it establishes an interesting connection between pure Nash equilibria and the lexicographically minimum expected latency vector that may be further explored.

To this end, we continue to present an efficient, yet simple algorithm (Theorem 3.2) that computes a pure Nash equilibrium. The algorithm proceeds by sorting all user traffics in non-decreasing order and dropping each traffic in order into the link that currently minimizes its expected latency cost. The time complexity of the algorithm is $\Theta(n \max\{\lg n, m\})$.

We proceed to consider the related problems BEST NASH EQUILIBRIUM SUPPORTS and WORST NASH EQUILIBRIUM SUPPORTS of determining either the *best* or the *worst* pure Nash equilibrium (with respect to social cost), respectively. Not surprisingly, we show that the decision versions of both problems are $\mathcal{NP}$-complete in the strong sense (Theorem 3.3 and Theorem 3.4). The $\mathcal{NP}$-completeness proofs employ reductions from 3-PARTITION [4, Problem SP15]; the proofs rely critically on using *non-identical* traffics, while they assume that all link capacities are equal.

We now turn to mixed Nash equilibria. We start with a structural result for the model of uniform capacities. In particular, we show that in a Nash equilibrium of the selfish routing game we consider, there can be no links traversed by a single user which also "crosses" another user over a different link (Proposition 4.1). Using this property, we establish that for the model of uniform capacities, and assuming that there are only *two* users, the worst mixed Nash equilibrium (with respect to social cost) is the fully mixed Nash equilibrium (Theorem 4.2). Although we believe that this result holds for *any* number of users, we have been unable to extend it beyond the case of two (still assuming uniform capacities).

We continue to formulate an efficient and elegant algorithm for computing a mixed Nash equilibrium (Theorem 5.1). More specifically, the algorithm computes a *generalized fully mixed* Nash equilibrium; this is a generalization of *fully mixed* Nash equilibria [18], where there is a set of "empty" links that do not belong to the support of any user, which includes all remaining links. The algorithm incrementally constructs the common support of all users by throwing away "slow" links (ones with

[1] See [37, Notes for Lecture 2] for a modern account of the proof of Nash's Theorem.

small capacity) till it converges to a common support that comprises a fully mixed Nash equilibrium. A crucial requirement for the algorithm to work is that all user traffics be identical.

We have also obtained an analog of Theorem 4.2 for the model of arbitrary capacities. We establish that *any* Nash equilibrium, in particular the worst one, incurs a social cost that does not exceed 33.06 times the social cost of the fully mixed Nash equilibrium (Theorem 6.1). This result is shown by establishing some interesting properties of the support and the expected latency of links in an arbitrary Nash equilibrium and by comparing the tails of the distribution of maximum link latency in the generalized fully mixed Nash equilibrium and in an arbitrary Nash equilibrium. Theorem 4.2 and Theorem 6.1 provide together substantial evidence about the "completeness" of the fully mixed Nash equilibrium: it appears that it suffices, in general, to focus on bounding the social cost of the fully mixed Nash equilibrium and then use reduction results (such as Theorem 4.2 and Theorem 6.1) to obtain bounds for the general case.

We then shift gears to study the computational complexity of NASH EQUILIBRIUM SOCIAL COST[2]. We have obtained both negative and positive results here. First for the bad news. We show that the problem is $\#\mathcal{P}$-complete [36] in general for the case of mixed Nash equilibria (Theorem 7.1). To show that NASH EQUILIBRIUM SOCIAL COST is $\#\mathcal{P}$-complete, we use a reduction from the problem of computing the probability that the sum of n independent random variables does not exceed a given threshold (see e.g. [11, Theorem 2.1] for the $\#\mathcal{P}$-completeness of the latter problem). We prove that this probability can be recovered by two calls to a (hypothetical) oracle returning the social cost of a given mixed Nash equilibrium.

On the positive side, we get around the established hardness of computing *exactly* the social cost of any mixed Nash equilibrium by presenting a fully polynomial-time randomized approximation scheme[3] for computing the social cost of any given mixed Nash equilibrium to any required degree of approximation (Theorem 7.2). The required number of iterations for the Monte Carlo scheme follows appropriately from Chebyshev's inequality and an easy upper and lower bound on the social cost.

We point out that the polynomial algorithms we have presented for the computation of pure and mixed Nash equilibria (Theorem 3.2 and Theorem 5.1, respectively) are the *first* known polynomial algorithms for the problem (for either the general case of a strategic game with a finite number of strategies, or even for a specific game). On the other hand,On the other hand, the hardness results we have obtained (Theorem 3.3, Theorem 3.4, and Theorem 7.1) indicate that optimization and counting problems in Computational Game Theory may be hard even when restricted to *specific,* simple games such as the selfish routing game considered in our work.

We believe that the polynomial algorithms we have derived (Theorem 3.2 and Theorem 5.1) may offer valuable ideas for settling either other tractable instances of the same game or other games of a similar flavor (e.g., selfish routing over a larger network). To this end, we feel that results offering insights into the syntax and structure of Nash equilibria will be handy. In addition, elimination results, such as Proposition 4.1, and their extensions may be the key to reducing the number of candidate Nash equilibria down to polynomial, which will, in turn, imply the tractability of instances of the game that we have not settled (e.g., computing mixed Nash equilibria when both traffics and capacities vary

[2]For the case of pure Nash equilibria, this problem is trivially in $\mathcal{P}$, since it amounts to computing the maximum.

[3]Consider a counting problem Π with solution $\Pi(x)$ on any instance x. An algorithm A is a *fully polynomial-time randomized approximation scheme* [9, 10] for Π, or *FPRAS* for short, if for each instance x, for any error parameter $\varepsilon > 0$, $\mathbb{P}\mathrm{r}\left(|\mathsf{A}(x) - \mathbf{\Pi}(x)| \leq \varepsilon \mathbf{\Pi}(x)\right) \geq \frac{3}{4}$, and the running time of A is polynomial in $|x|$ and $\frac{1}{\varepsilon}$.

arbitrarily). We leave this as a subject for future work.

1.3 Related Work

The selfish routing game considered in this paper was first introduced by Koutsoupias and Papadimitriou [16] as a vehicle for the study of the price of selfishness for routing over non-cooperative networks [14], like the Internet. This game was subsequently studied in the work of Mavronicolas and Spirakis [18], where fully mixed Nash equilibria were introduced and analyzed. In both works, the aim had been to quantify the amount of performance loss in routing due to selfish behavior of the users. (Later studies of the selfish routing game from the same point of view, that of performance, include the works by Koutsoupias *et al.* [15], and by Czumaj and Vöcking [1].) Unlike these previous papers, our work considers the selfish routing game from the point of view of computational complexity and attempts to classify certain algorithmic problems related to the computation of Nash equilibria of the game with respect to their computational complexity.

Extensive surveys of algorithms and techniques from the literature of Game Theory for the computation of Nash equilibria of general bimatrix games in either strategic or extensive form appear in [19, 35]; see also [24, Section 3.1]. All known such algorithms incur exponential running time, with the seminal algorithm of Lemke and Howson [17] being the prime example; see also [32, 34] for still inefficient extensions. Issues of computational complexity for the computation of Nash equilibria in general games have been raised by Megiddo [20], Megiddo and Papadimitriou [21], and Papadimitriou [29]. The $\mathcal{NP}$-hardness of computing a Nash equilibrium of a *general* bimatrix game with maximum payoff has been established by Gilboa and Zemel [5]. For other algorithmic works on the computation of Nash equilibria, see, e.g., [12, 13]. The book by Scarf [33] is devoted to the computation of Nash equilibria over various economic settings. A similar in motivation and spirit to our paper is the very recent paper by Deng *et al.* [2], which proves complexity, approximability and inapproximability results for the problem of computing an exchange equilibrium in markets with indivisible goods. A general account on overlaps between Computer Science and Game Theory appears in [7].

2 Framework

Most of our definitions are patterned after those in [16, Sections 1 & 2] and [18, Section 2].

2.1 Notation

Throughout, denote for any integer $m \geq 2$, $[m] = \{1, \ldots, m\}$. For an event E in a sample space, denote $\mathbb{P}\mathrm{r}(E)$ the probability of event E happening. For a random variable X, denote $\mathbb{E}(X)$ the *expectation* of X and $\mathrm{Var}(X)$ the *variance* of X.

We consider a *network* consisting of a set of m parallel *links* $1, 2, \ldots, m$ from a *source* node to a *destination* node. Each of n *network users* $1, 2, \ldots, n$, or *users* for short, wishes to route a particular amount of traffic along a (non-fixed) link from source to destination. (Throughout, we will be using subscripts for users and superscripts for links.) Denote w_i the *traffic* of user $i \in [n]$. Define the $n \times 1$ *traffic vector* $\mathbf{w}$ in the natural way. Assume throughout that $m > 1$ and $n > 1$.

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

For a probability matrix $\mathbf{P}$, define *indicator variables* $I_\imath^\ell \in \{0, 1\}$, $\imath \in [n]$ and $\ell \in [m]$, such that $I_\imath^\ell = 1$ if and only if $p_i^\ell > 0$. Thus, the support of the mixed strategy for user $\imath \in [n]$ is the set $\{\ell \in [m] \mid I_\imath^\ell = 1\}$. For each link $\ell \in [m]$, define the *view* of link ℓ, denoted $view(\ell)$, as the set of users $i \in [n]$ that potentially assign their traffics to link ℓ; so, $view(\ell) = \{i \in [n] \mid I_i^\ell = 1\}$. A link $\ell \in [m]$ is *solo* [18] if $|view(\ell)| = 1$; thus, there is exactly one user, denoted $s(\ell)$, that considers a solo link ℓ. Denote $\mathcal{S}$ the set of solo links.

2.1.1 Syntactic Classes of Mixed Strategies

A mixed strategies profile $\mathbf{P}$ is *fully mixed* [18] if for all users $\imath \in [n]$ and links $\jmath \in [m]$, $I_\imath^\jmath = 1$. Throughout, we will be considering a pure strategies profile as a special case of a mixed strategies profile. in which all (mixed) strategies are pure. We proceed to define two new variations of fully mixed strategies profiles. A mixed strategies profile $\mathbf{P}$ is *generalized fully mixed* if there exists a subset $\mathsf{Links} \subseteq [m]$ such that for each pair of a user $\imath \in [n]$, and a link $\jmath \in [m]$, $I_\imath^\jmath = 1$ if $\jmath \in \mathsf{Links}$ and 0 if $\jmath \notin \mathsf{Links}$. Thus, the fully mixed strategies profile is the special case of generalized fully mixed strategies profiles where $\mathsf{Links} = [m]$.

2.1.2 Cost Measures

Denote $c^\ell > 0$ the *capacity* of link $\ell \in [m]$, representing the rate at which the link processes traffic. So, the *latency* for traffic w through link ℓ equals w/c^ℓ. In the model of *uniform capacities,* all link capacities are equal to c, for some constant $c > 0$; link capacities may vary arbitrarily in the model of *arbitrary capacities*. For a pure strategies profile $\langle \ell_1, \ell_2, \ldots, \ell_n \rangle$, the *latency cost for user* $\imath \in [n]$, denoted $\lambda_\imath$, is $(\sum_{k:\ell_k=\ell_\imath} w_k)/c^{\ell_i}$; that is, the latency cost for user i is the latency of the link it chooses. For a mixed strategies profile $\mathbf{P}$, denote W^ℓ the *expected traffic* on link $\ell \in [m]$; clearly, $W^\ell = \sum_{i=1}^n p_i^\ell w_i$. Given $\mathbf{P}$, define the $m \times 1$ *expected traffic vector* $\mathbf{W}$ induced by $\mathbf{P}$ in the natural way. Given $\mathbf{P}$, denote Λ^ℓ the *expected latency* on link $\ell \in [m]$; clearly, $\Lambda^\ell = \frac{W^\ell}{c^\ell}$. Define the $m \times 1$ *expected latency vector* $\mathbf{\Lambda}$ in the natural way. For a mixed strategies profile $\mathbf{P}$, the *expected latency cost* for user $\imath \in [n]$ on link $\ell \in [m]$, denoted λ_i^ℓ, is the expectation, over all random choices of the remaining users, of the latency cost for user $\imath$ had its traffic been assigned to link ℓ; thus, $\lambda_\imath^\ell = \frac{w_\imath + \sum_{k=1,k\neq\imath} p_k^\ell w_k}{c^\ell} = \frac{(1-p_\imath^\ell)w_\imath + W^\ell}{c^\ell}$. For each user $\imath \in [n]$, the *minimum expected latency cost,* denoted $\lambda_\imath$, is the minimum, over all links $\ell \in [m]$, of the expected latency cost for user $\imath$ on link ℓ; thus, $\lambda_i = \min_{\ell \in [m]} \lambda_i^\ell$. For a probability matrix $\mathbf{P}$, define the $n \times 1$ *minimum expected latency cost vector* λ induced by $\mathbf{P}$ in the natural way.

Associated with a traffic vector $\mathbf{w}$ and a mixed strategies profile $\mathbf{P}$ is the *social cost* [16, Section 2], denoted $\mathsf{SC}(\mathbf{w}, \mathbf{P})$, which is the expectation, over all random choices of the users, of the maximum (over all links) latency of traffic through a link; thus,

$$\mathsf{SC}(\mathbf{w}, \mathbf{P}) = \mathbb{E}\left(\max_{\ell \in [m]} \frac{\sum_{k:\ell_k=\ell} w_k}{c^\ell}\right) = \sum_{\langle \ell_1, \ell_2, \ldots, \ell_n \rangle \in [m]^n} \left(\prod_{k=1}^{n} p_k^{\ell_k} \cdot \max_{\ell \in [m]} \frac{\sum_{k:\ell_k=\ell} w_k}{c^\ell}\right)$$

Note that $\mathsf{SC}(\mathbf{w}, \mathbf{P})$ reduces to the maximum latency through a link in the case of pure strategies. On the other hand, the *social optimum* [16, Section 2] associated with a traffic vector $\mathbf{w}$, denoted $\mathsf{OPT}(\mathbf{w})$, is the *least possible* maximum (over all links) latency of traffic through a link; thus,

$$\mathsf{OPT}(\mathbf{w}) = \min_{\langle \ell_1, \ell_2, \ldots, \ell_n \rangle \in [m]^n} \max_{\ell \in [m]} \frac{\sum_{k:\ell_k=\ell} w_k}{c^\ell}$$

Note that while $\mathsf{SC}(\mathbf{w}, \mathbf{P})$ is defined in relation to a mixed strategies profile $\mathbf{P}$, $\mathsf{OPT}(\mathbf{w})$ refers to the *optimum* pure strategies profile.

2.1.3 Nash Equilibria

We are interested in a special class of mixed strategies called Nash equilibria [25] that we describe below. Formally, the probability matrix $\mathbf{P}$ is a *Nash equilibrium* [16, Section 2] if for all users $\imath \in [n]$ and links $\ell \in [m]$, $\lambda_\imath^\ell = \lambda_\imath$ if $I_\imath^\ell = 1$, and $\lambda_\imath^\ell > \lambda_\imath$ if $I_\imath^\ell = 0$. Thus, each user assigns its traffic with positive probability only on links (possibly more than one of them) for which its expected latency cost is minimized; this implies that there is no incentive for a user to unilaterally deviate from its mixed strategy in order to avoid links on which its expected latency cost is higher than necessary.

We continue to state some already known properties of Nash equilibria that will be used in our later proofs. Koutsoupias and Papadimitriou [16, Section 2] provide necessary conditions for Nash equilibria.

Proposition 2.1 (Koutsoupias and Papadimitriou [16]) *Fix a Nash equilibrium* $\mathbf{P}$. *Then, for any user* $i \in [n]$ *and link* $\ell \in [m]$, (1) *for all links* $\ell \in [m]$, $W^\ell = \sum_{k=1}^{n} I_k^\ell \left(W^\ell + w_k - c^\ell \lambda_k\right)$, *and* (2) *for all users* $i \in [n]$, $w_i = \sum_{j=1}^{m} I_i^j \left(W^j + w_i - c^j \lambda_i\right)$.

The following result due to Mavronicolas and Spirakis [18, Lemma 6.1] provides a simple characterization of existence and uniqueness of fully mixed Nash equilibria under the model of identical traffics and capacitated parallel links.

Lemma 2.2 (Mavronicolas and Spirakis [18]) *Consider the fully mixed case for* n *users of identical traffic and* m *capacitated parallel links, and let* $C(m) = \sum_{\ell=1}^{m} c^\ell$. *Then, for all links* $\ell \in [m]$,

$$c^\ell \in \left(\frac{C(m)}{m+n-1}, \frac{nC(m)}{m+n-1}\right) ,$$

if and only if there exists a Nash equilibrium, which must be unique and have associated Nash probabilities

$$p_i^\ell = \frac{(m+n-1)c^\ell - C(m)}{(n-1)C(m)}$$

Mavronicolas and Spirakis [18, Lemma 4.2] proved that in a fully mixed Nash equilibrium, the vector of users' minimum expected latency costs is a linear transformation of the vector of users' traffics. Therefore, under the model of identical traffics, all users incur the same minimum expected latency cost in a fully mixed equilibrium. The following lemma is a special case of [18, Lemma 4.2] for the model identical traffics.

Lemma 2.3 (Mavronicolas and Spirakis [18]) *Consider a fully mixed Nash equilibrium for a selfish routing game on n users of identical traffic and m capacitated parallel links. Let λ_i denote the minimum expected latency cost of any user i, and let $C(m) = \sum_{\ell=1}^{m} c^{\ell}$. Then, $\lambda_i = (m + n - 1)/C(m)$.*

Finally, Mavronicolas and Spirakis [18, Lemma 5.1] show that in the model of uniform capacities, all links are equiprobable in a fully mixed Nash equilibrium.

Lemma 2.4 (Mavronicolas and Spirakis [18]) *Consider the fully mixed case under the model of uniform capacities. Then, there exists a unique Nash equilibrium with associated Nash probabilities $p_i^{\ell} = 1/m$ for each user $i \in [n]$ and link $\ell \in [m]$.*

2.2 Algorithmic Problems

In this section, we formally define several algorithmic problems related to Nash equilibria. The definitions are given in the style of Garey and Johnson [4].

Π_1: NASH EQUILIBRIUM SUPPORTS

INSTANCE: A number n of users; a number m of links; for each user i, a rational number $w_i > 0$, called the *traffic* of user i; for each link j, a rational number $c^j > 0$, called the *capacity* of link j.

OUTPUT: Indicator variables $I_i^j \in \{0, 1\}$, where $1 \leq i \leq n$ and $1 \leq j \leq m$, that support a Nash equilibrium for the system of the users and the links.

We continue with two complementary to each other optimization problems (with respect to social cost).

Π_2: BEST NASH EQUILIBRIUM SUPPORTS

INSTANCE: A number n of users; a number m of links; for each user i, a rational number $w_i > 0$, called the *traffic* of user i; for each link j, a rational number $c^j > 0$, called the *capacity* of link j.

OUTPUT: Indicator variables $I_i^j \in \{0, 1\}$, where $1 \leq i \leq n$ and $1 \leq j \leq m$, that support a Nash equilibrium with *minimum* social cost for the system of the users and the links.

Π_3: WORST NASH EQUILIBRIUM SUPPORTS

INSTANCE: A number n of users; a number m of links; for each user i, a rational number $w_i > 0$, called the *traffic* of user i; for each link j, a rational number $c^j > 0$, called the *capacity* of link j.

OUTPUT: Indicator variables $I_i^j \in \{0, 1\}$, where $1 \leq i \leq n$ and $1 \leq j \leq m$, that support a Nash equilibrium with *maximum* cost for the system of the users and the links.

Π_4: NASH EQUILIBRIUM WITH GIVEN SUPPORTS

INSTANCE: A number n of users; a number m of links; for each user i, a rational number $w_i > 0$, called the *traffic* of user i; for each link j, a rational number $c^j > 0$, called the *capacity* of link j; for each pair of user i and link j, an *indicator variable* $I_i^j \in \{0, 1\}$.

QUESTION: Does there exist a Nash equilibrium supported by the indicator variables for the system of the users and the links?

By results of Mavronicolas and Spirakis [18, Theorem 4.7], NASH EQUILIBRIUM WITH GIVEN SUPPORTS can be solved in time $\Theta(mn)$ when restricted to fully mixed strategies. This suggests that for any particular syntactic class of Nash equilibria, the problem NASH EQUILIBRIUM WITH GIVEN SUPPORTS can be solved efficiently by formulating a set of (polynomially computable) necessary and sufficient conditions, and evaluating them on any particular instance of the problem.

$\mathbf{\Pi_5}$: NASH EQUILIBRIUM SOCIAL COST

INSTANCE: A number n of users; a number m of links; for each user i, a rational number $w_i > 0$, called the *traffic* of user i; for each link j, a rational number $c^j > 0$, called the *capacity* of link j; a *Nash equilibrium* $\mathbf{P}$ for the system of the users and the links.

OUTPUT: The social cost of the Nash equilibrium $\mathbf{P}$.

To establish our hardness results, we will use the problem 3-PARTITION, which is $\mathcal{NP}$-complete in the strong sense (see e.g. [4, Problem SP15 and Theorem 4.4]). 3-PARTITION is formally defined as follows:

$\mathbf{\Pi}_0$: 3-PARTITION

INSTANCE: A positive integer $m \geq 2$, a positive integer B, and a set $J = \{w_1, \ldots, w_{3m}\}$ of $3m$ positive integer weights such that $B/4 < w_i < B/2$ for all $i \in [3m]$, and $\sum_{i=1}^{3m} w_i = mB$.

QUESTION: Can $[3m]$ be partitioned into m sets $J_1, \ldots, J_m$ such that $\sum_{i \in J_j} w_i = B$ for all $j \in [m]$?

3 Pure Nash Equilibria

3.1 Existence of a Pure Nash Equilibrium

We start with a preliminary result remarked by Kurt Mehlhorn.

Theorem 3.1 (A Pure Nash Equilibrium Exists) *There exists at least one pure Nash equilibrium.*

Proof: Consider the universe of pure strategies profiles. Each such profile induces a *sorted* expected latency vector $\mathbf{\Lambda} = \langle \Lambda^1, \Lambda^2, \ldots, \Lambda^m \rangle$, such that $\Lambda^1 \geq \Lambda^2 \geq \ldots \geq \Lambda^m$, in the natural way. (Rearrangement of links may be necessary to guarantee that the expected latency vector is sorted.) Consider the lexicographically minimum expected latency vector $\mathbf{\Lambda}_0$ [4] and assume that it corresponds to a pure strategies profile $\mathbf{P}_0$. We will argue that $\mathbf{P}_0$ is a (pure) Nash equilibrium.

Indeed, assume, by way of contradiction, that $\mathbf{P}_0$ is *not* a Nash equilibrium. By definition of Nash equilibrium, there exists a user $\imath \in [n]$ assigned by $\mathbf{P}_0$ to link $\jmath \in [n]$, and a link $\kappa \in [m]$ such that

$$\Lambda^\jmath \quad > \quad \Lambda^\kappa + \frac{w_\imath}{c^\kappa}\,.$$

Construct now from $\mathbf{P}_0$ a pure strategies profile $\widehat{\mathbf{P}_0}$ which is identical to $\mathbf{P}_0$ except that user $\imath$ is now assigned to link κ. Denote $\widehat{\Lambda_0} = \langle \widehat{\Lambda^1}, \widehat{\Lambda^2}, \ldots, \widehat{\Lambda^m} \rangle$ the traffic vector induced by $\widehat{\mathbf{P}_0}$. By construction,

$$\widehat{\Lambda^\jmath} \quad = \quad \Lambda^\jmath - \frac{w_\imath}{c^\jmath} \; < \; \Lambda^\jmath\,,$$

[4] For any two $m \times 1$ vectors $\mathbf{x}$ and $\mathbf{y}$ say that $\mathbf{x}$ *is lexicographically less than* $\mathbf{y}$ if there is an index $k \leq m$ such that for each index $i \leq k$, $x_i = y_i$, while $x_k < y_k$. Clearly, the relation of being lexicographically less induces a total order on a set of vectors. The *lexicographically minimum* vector is the least element of this total order.

while by construction and assumption,

$$\widehat{\Lambda^\kappa} \;=\; \Lambda^\kappa + \frac{w_\imath}{c^\kappa} \;<\; \Lambda^\jmath\,.$$

Since $\mathbf{\Lambda}_0$ is sorted in non-increasing order and $\Lambda^\kappa + \frac{w_\imath}{c^\kappa} < \Lambda^\jmath$, $\Lambda^\jmath$ precedes Λ^κ in $\mathbf{\Lambda}_0$. Clearly, all entries preceding $\Lambda^\jmath$ in $\mathbf{\Lambda}_0$ remain unchanged in $\widehat{\mathbf{\Lambda}_0}$. Consider now the $\jmath$-th entry of $\widehat{\mathbf{\Lambda}_0}$. There are three possibilities.

- The $\jmath$-th entry of $\widehat{\mathbf{\Lambda}_0}$ is $\widehat{\Lambda^\jmath}$. Since $\widehat{\Lambda^\jmath} < \Lambda^\jmath$, it follows that $\widehat{\mathbf{\Lambda}_0}$ is lexicographically less than $\mathbf{\Lambda}_0$. A contradiction.
- The $\jmath$-th entry of $\widehat{\mathbf{\Lambda}_0}$ is $\widehat{\Lambda^\kappa}$. Since $\widehat{\Lambda^\kappa} < \Lambda^\jmath$, it follows that $\widehat{\mathbf{\Lambda}_0}$ is lexicographically less than $\mathbf{\Lambda}_0$. A contradiction.
- The $\jmath$-th entry of $\widehat{\mathbf{\Lambda}_0}$ is some entry of $\mathbf{\Lambda}_0$ that followed $\Lambda^\jmath$ in $\mathbf{\Lambda}_0$ and remained unchanged in $\widehat{\mathbf{\Lambda}_0}$. Since $\mathbf{\Lambda}_0$ is sorted in non-increasing order, any such entry is less than $\Lambda^\jmath$. It follows that $\widehat{\mathbf{\Lambda}_0}$ is lexicographically less than $\mathbf{\Lambda}_0$. A contradiction.

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

We remark that the proof of Theorem 3.1 establishes that the lexicographically minimum expected traffic vector represents a (pure) Nash equilibrium. Since there are exponentially many pure strategies profiles and that many expected traffic vectors, Theorem 3.1 only provides an *inefficient* proof of existence of pure Nash equilibria (cf. Papadimitriou [29]). An *efficient* algorithmic proof of existence of pure Nash equilibria is the subject of the following section.

3.2 Computing a Pure Nash Equilibrium

We show:

Theorem 3.2 NASH EQUILIBRIUM SUPPORTS *is in $\mathcal{P}$ when restricted to pure equilibria.*

Proof: We present a polynomial-time algorithm A_{pure} that computes the supports of a pure Nash equilibrium. Roughly speaking, the algorithm A_{pure} works in a greedy fashion; it considers each of the user traffics in non-increasing order and assigns it to the link that minimizes (among all links) the latency cost of the user had its traffic been assigned to that link. (Figure 3.2 presents pseudocode for the algorithm A_{pure}.)

Clearly, the supports computed by A_{pure} represent a pure strategies profile. We will show that this profile is a Nash equilibrium. We argue inductively on the number $\imath$ of iterations, $1 \leq \imath \leq n$, of the main loop of A_{pure}. We prove that the system of users and links is is Nash equilibrium after each such iteration.

For the basis case, where $\imath = 1$, consider the user 1 assigned to link ℓ. After the first iteration, $W^\ell = \frac{w_1}{c^\ell}$, while $W^\jmath = 0$ for $\jmath \neq \ell$. By the way link ℓ was chosen, for each link $\jmath \in [m]$, with $\jmath \neq \ell$, $\frac{w_1}{c^\ell} \leq \frac{w_1}{c^\jmath}$, or $W^\ell < W^\jmath + \frac{w_1}{c^\jmath}$. This establishes that the system of users and links is in Nash equilibrium after iteration 1.

The algorithm A_{pure}**:**

- For each pair of a user $\imath \in [n]$ and a link $\jmath \in [m]$, $I_\imath^\jmath := 0$.
- For each link $\jmath \in [m]$, $W^\jmath := 0$.
- Sort the traffics in non-increasing order so that $w_1 \geq w_2 \geq \ldots \geq w_n$.
- For each user $\imath := 1$ to n, **do**
  - $\ell := \arg\min_{1 \leq \jmath \leq m} \left(W^\jmath + \frac{w_\imath}{c^\jmath} \right)$;
  - $I_\imath^\ell := 1$;
  - $W^\ell := W^\ell + \frac{w_i}{c^\ell}$

Figure 1: The algorithm A_{pure}.

Assume inductively that the system of users and links is in Nash equilibrium after iteration $\imath - 1$. Consider what happens after iteration $\imath$, at which user $\imath$ was assigned to link ℓ. Note that iteration $\imath$ only increased the expected traffic on link ℓ and did not increase the expected traffic on any other link. Since, by induction hypothesis, the system was in Nash equilibrium right before before iteration $\imath$, no user other than $\imath$ wanted to switch its link right before iteration $\imath$. It follows that no user other than $\imath$ and those already assigned to link ℓ wants to switch its link right after iteration $\imath$. (Users already assigned to link ℓ must also be exempted since the expected traffic on the link they were assigned increases, and this may make it possible for them to want to switch links.) Note that among all users assigned to link ℓ, including user $\imath$, the traffic of user $\imath$ is the smallest; thus, it suffices to show that user $\imath$ does not want to switch links after iteration $\imath$. By the way link ℓ was chosen for user $\imath$, it must be that for any other link $\jmath \neq \ell$, $W^\ell + \frac{w_\imath}{c^\ell} < W^\jmath + \frac{w_\imath}{c^\jmath}$, where $W^\jmath$ is the expected traffic on link $\jmath$ right before iteration $\imath$. Since, by the algorithm, $W^\jmath + \frac{w_\imath}{c^\jmath}$ is the expected traffic on link $\jmath$ right after iteration $\imath$, it follows that user $\imath$ does not want to switch links after iteration $\imath$, as needed. ■

(This nice observation is due to B. Monien [22].) We remark that A_{pure} can be viewed as a variant of Graham's Longest Processing Time (LPT [6]) algorithm for assigning tasks to identical machines. Nevertheless, since in our case the links may have different capacities, our algorithm instead of choosing the link that will first become idle, it actually chooses the link that minimizes the completion time of the specific task (i.e., the load of a machine prior to the assignment of the task under consideration, plus the overhead of this task). Clearly, this greedy algorithm leads to an assignment which is, as we establish, a Nash equilibrium.

3.3 The Complexity of Computing the Best and the Worst Pure Nash Equilibrium

We prove that it is $\mathcal{NP}$-hard to compute the best Nash equilibrium and the worst pure Nash equilibrium.

Theorem 3.3 *The decision version of* BEST NASH EQUILIBRIUM SUPPORTS *is $\mathcal{NP}$-complete in the strong sense even for identical links.*

Proof: By the definition of the social cost, a selfish routing game on identical links admits a Nash equilibrium of social cost at most B iff it admits a *pure* Nash equilibrium of social cost at most B (see also the proof of Theorem 4.2). Hence we can restrict our attention to pure Nash equilibria.

It is straightforward to decide in $\mathcal{NP}$ whether a selfish routing game admits a pure Nash equilibrium of social cost at most B. To establish $\mathcal{NP}$-completeness in the strong sense, we use a reduction from 3-PARTITION, which is $\mathcal{NP}$-complete in the strong sense (see e.g. [4, Problem SP15 and Theorem 4.4]).

Given an instance (m, B, J) of 3-PARTITION, we construct a selfish routing game $(\mathbf{w}, m)$ with $3m$ users on m identical links. For every $i \in [3m]$, the traffic of user i is w_i. By construction, $(\mathbf{w}, m)$ admits a pure Nash equilibrium of social cost B iff (m, B, J) is a YES-instance of 3-PARTITION.

More precisely, if (m, B, J) is a YES-instance of 3-PARTITION, let $J_1, \ldots, J_m$ be a partition of $[3m]$ into m sets such that $\sum_{i \in J_j} w_i = B$ for all $j \in [m]$, and let $\mathbf{P}$ be the pure strategies profile where for every $i \in [3m]$, $p_i^j = 1$ if $i \in J_j$, and $p_i^j = 0$ otherwise. Since all links have latency B in $\mathbf{P}$, $\mathbf{P}$ is a pure Nash equilibrium with social cost B. For the converse, let us assume that $(\mathbf{w}, m)$ admits a pure Nash equilibrium $\mathbf{P}$ of social cost at most B. Since $\sum_{i=1}^{3m} w_i = mB$, all links have latency (and traffic) B in $\mathbf{P}$. Therefore, setting $J_j = \{i \in [3m] : p_i^j = 1\}$, $j \in [m]$, yields a YES-certificate for the corresponding instance of 3-PARTITION. ■

Theorem 3.4 *When restricted to pure equilibria, the decision version of* WORST NASH EQUILIBRIUM SUPPORTS *is $\mathcal{NP}$-complete in the strong sense even for identical links.*

Proof: Membership in $\mathcal{NP}$ is straightforward. To establish $\mathcal{NP}$-completeness in the strong sense, we use a reduction from 3-PARTITION.

Given an instance (m, B, J) of 3-PARTITION, we construct a selfish routing game $(\mathbf{w}, m+1)$ with $3m+2$ users on $m+1$ identical links. For every $i \in [3m]$, the traffic of user i is w_i. The traffic of users $3m+1$ and $3m+2$ is $w_{3m+1} = w_{3m+2} = B$. By the definition of 3-PARTITION, $B/4 < w_i < B/2$ for all $i \in [3m]$, and $\sum_{i=1}^{3m+2} w_i = (m+2)B$. We show that $(\mathbf{w}, m+1)$ admits a pure Nash equilibrium of social cost at least $2B$ iff (m, B, J) is a YES-instance of 3-PARTITION.

If (m, B, J) is a YES-instance of 3-PARTITION, let $J_1, \ldots, J_m$ be a partition of $[3m]$ into m sets with $\sum_{i \in J_j} w_i = B$ for all $j \in [m]$, and let $\mathbf{P}$ be the pure strategies profile assigning users $3m+1$ and $3m+2$ to link $m+1$, and the remaining users according to $J_1, \ldots, J_m$. Formally, $p_{3m+1}^{m+1} = p_{3m+2}^{m+1} = 1$ and $p_{3m+1}^j = p_{3m+2}^j = 0$ for all $j \in [m]$, and for all $i \in [3m]$ and $j \in [m]$, $p_i^j = 1$ if $i \in J_j$, and $p_i^j = 0$ otherwise. Every link $j \in [m]$ has latency B and link $m+1$ has latency $2B$. Since no user has an incentive to deviate from her strategy, $\mathbf{P}$ is a pure Nash equilibrium of social cost $2B$.

For the converse, let us assume that $(\mathbf{w}, m+1)$ admits a pure Nash equilibrium $\mathbf{P}$ of social cost at least $2B$. Without loss of generality, let $m+1$ be a link with latency at least $2B$ in $\mathbf{P}$. Since all users have traffic at most B and no user assigned to $m+1$ can decrease her latency cost by switching to a different link, all links have latency at least B in $\mathbf{P}$. Since the total traffic is equal to $(m+2)B$, the latency (and traffic) of the first m links is precisely B, and the latency of link $m+1$ is precisely

$2B$. Furthermore, none of the first $3m$ users is assigned to link $m+1$ because $B/4 < w_i < B/2$ for all $i \in [3m]$. Otherwise, a user $i \in [3m]$ assigned to link $m+1$ could decrease her latency cost by switching to a link $j \in [m]$. Therefore, $\mathbf{P}$ assigns users $3m+1$ and $3m+2$ to link $m+1$ and the first $3m$ users to the first m links. For every $j \in [m]$, let $J_j = \{i \in [3m] : p_i^j = 1\}$. Then $J_1, \ldots J_m$ comprises a YES-certificate for the corresponding instance of 3-PARTITION. ■

4 A Characterization of the Worst Mixed Nash Equilibrium

We start with a structural property of mixed Nash equilibria. In the following proposition, we say that a user *crosses* another user if their supports cross each other, i.e. their supports are neither disjoint nor the same.

Proposition 4.1 *Let* $\mathbf{P}$ *be any Nash equilibrium under the model of uniform capacities. Then* $\mathbf{P}$ *induces no solo link considered by a user that crosses another user.*

Proof: Assume that $\mathbf{P}$ induces a solo link ℓ considered by a user $s(\ell)$ that crosses another user; thus, there exists another link $\ell_0 \in support(s(\ell))$ and a user $i_0 \in view(\ell_0)$, so that $p_{i_0}^{\ell_0} > 0$. By the definition of users' expected latency cost,

$$\lambda_{s(\ell)}^{\ell_0} \geq w_{s(\ell)} + p_{i_0}^{\ell_0} w_{i_0} > w_{s(\ell)} = \lambda_{s(\ell)}^{\ell} ,$$

which contradicts the hypothesis that $\mathbf{P}$ is a Nash equilibrium. ■

We then use Proposition 4.1 to provide a syntactic characterization of the worst mixed Nash equilibrium under the model of uniform capacities.

Theorem 4.2 *Consider the model of uniform capacities and assume that* $n = 2$*. Then, the worst Nash equilibrium is the fully mixed Nash equilibrium.*

Proof: Assume, without loss of generality, that $w_1 \geq w_2$. Thus, for any assignment of the two traffics to the m links, the only possible maxima are w_1 (occurring if users 1 and 2 choose different links) and $w_1 + w_2$ (occurring if users 1 and 2 choose the same link). Consider any mixed Nash equilibrium $\mathbf{P}$, which is a $2 \times m$ matrix. Thus,

$$
\begin{aligned}
\mathsf{SC}\left(\mathbf{w},\mathbf{P}\right) &= w_1 \sum_{\ell_1,\ell_2\in[m],\ell_1\neq\ell_2} p_1^{\ell_1}p_2^{\ell_2} + (w_1+w_2)\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell} \\
&= w_1 \sum_{\ell_1,\ell_2\in[m],\ell_1\neq\ell_2} p_1^{\ell_1}p_2^{\ell_2} + w_1\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell} + w_2\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell} \\
&= w_1 \left(\sum_{\ell_1,\ell_2\in[m],\ell_1\neq\ell_2} p_1^{\ell_1}p_2^{\ell_2} + \sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell}\right) + w_2\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell} \\
&= w_1 \sum_{\ell_1,\ell_2\in[m]} p_1^{\ell_1}p_2^{\ell_2} + w_2\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell} \\
&= w_1 \sum_{\ell_1\in[m]} p_1^{\ell_1}\left(\sum_{\ell_2\in[m]} p_2^{\ell_2}\right) + w_2\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell} \\
&= w_1 \sum_{\ell_1\in[m]} p_1^{\ell_1}\cdot 1 + w_2\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell} \\
&= w_1\cdot 1 + w_2\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell} \\
&= w_1 + w_2\sum_{\ell\in[m]} p_1^{\ell}p_2^{\ell}\,.
\end{aligned}
$$

We will show that $\mathsf{SC}(\mathbf{w},\mathbf{P})$ is maximized when $\mathbf{P}$ is the fully mixed equilibrium. We proceed by case analysis.

1. Assume first that $\mathbf{P}$ is pure. We observe that it is not possible for both users to have the same pure strategy (since then the latency cost of a user on any other strategy would be smaller than its current latency cost, contradicting the equilibrium). This implies that the social cost of any pure Nash equilibrium is $\max\{w_1,w_2\}=w_1$. Hence, the social cost of any mixed Nash equilibrium is no less than the cost of any pure Nash equilibrium, which implies that the *worst* mixed Nash equilibrium is no better than any pure Nash equilibrium. We continue to analyze the cost of a mixed Nash equilibrium.

2. Assume now that $\mathbf{P}$ is not pure. There are two cases to consider.

   (a) Assume first that $\mathit{support}(1)\cap\mathit{support}(2)=\emptyset$. Then, clearly, the minimum expected latency cost of user 1 is w_1 (since user 2 does not consider any of the links in its support) and the minimum expected latency cost of user 2 is w_2 (since user 1 does not consider any of the links in its support). Since $\mathsf{SC}(\mathbf{w},\mathbf{P})=w_1+w_2\sum_{\ell\in[m]}p_1^{\ell}p_2^{\ell}$ and there is no link $\ell\in[m]$ such that both $p_1^{\ell}\neq 0$ and $p_2^{\ell}\neq 0$, it follows that $\mathsf{SC}(\mathbf{w},\mathbf{P})=w_1$, which is no worse than the social cost of any pure Nash equilibrium.

   (b) Assume now that $\mathit{support}(1)\cap\mathit{support}(2)\neq\emptyset$. We will show that in this case $\mathbf{P}$ is the fully mixed Nash equilibrium.

   We observe that Proposition 4.1 implies that $\mathit{support}(1)=\mathit{support}(2)$. Otherwise, there would a solo link and the two supports would cross. We will show that, in fact, $\mathit{support}(1)=\mathit{support}(2)=[m]$.

Assume, by way of contradiction, that there exists some link $\ell'' \in [m] \setminus support(1)$ $(= [m] \setminus support(2))$. Then, clearly, the expected latency cost of user 1 on link ℓ'' is equal to w_1, which is *less* than $w_1 + w_2 p_2^\ell$, its expected latency cost on link ℓ (since $p_2^\ell > 0$), and this contradicts equilibrium. It follows that in this case $support(1) = support(2) = [m]$, so that $\mathbf{P}$ is a fully mixed Nash equilibrium. By Lemma 2.4, $p_1^\ell = p_2^\ell = \frac{1}{m}$, so that

$$\begin{aligned} \mathsf{SC}(\mathbf{w}, \mathbf{P}) &= w_1 + w_2 \sum_{\ell \in [m]} p_1^\ell p_2^\ell \\ &= w_1 + w_2 \cdot m \cdot \frac{1}{m^2} \\ &= w_1 + w_2 \cdot \frac{1}{m} \end{aligned}$$

The previous analysis establishes that the worst mixed Nash equilibrium is the fully mixed Nash equilibrium, with corresponding social cost $w_1 + w_2 \cdot \frac{1}{m}$. ■

5 The Generalized Fully Mixed Nash Equilibrium

In this section, we prove that every selfish routing game on users of identical traffic and capacitated links admits a unique generalized fully mixed Nash equilibrium computable in $O(m \log m)$ time. Throughout this section, we assume, without loss of generality, that the links are indexed in non-increasing capacity order, i.e. $c^1 \geq c^2 \geq \ldots \geq c^m$, and that all traffics are equal to 1. For any integer $k \in [m]$, we let $C(k)$ denote the total capacity of the fastest k links, i.e. $C(k) = \sum_{j=1}^{k} c^j$. For any $k \in [m]$, we say that the set $[k]$ consisting of the fastest k links is a *fast link set*.

We start with a polynomial upper bound on the complexity of computing a generalized fully mixed Nash equilibrium.

Theorem 5.1 *Assume that all traffics are identical. Then,* NASH EQUILIBRIUM SUPPORTS *is in* $\mathcal{P}$ *when restricted to generalized fully mixed equilibria.*

Proof: We present a polynomial-time algorithm A_{gfm} that computes the support of a generalized fully mixed equilibrium. A generalized fully mixed Nash equilibrium corresponds to a fully mixed equilibrium when the game is restricted to the links in its support. The following modification of Lemma 2.2 gives a simple characterization of the games that admit a fully mixed Nash equilibrium.

Proposition 5.2 (B. Monien [22]) *Consider a selfish routing game* Γ *on* n *users of identical traffic and* m *capacitated parallel links. Then* Γ *admits a fully mixed Nash equilibrium, which must be unique, if and only if* $c^m > C(m)/(m + n - 1)$.

Proof: Since the links are indexed in non-decreasing capacity order, $c^m > C(m)/(m+n-1)$ implies that $c^\ell > C(m)/(m + n - 1)$ for all $\ell \in [m]$. This in turn implies that $c^\ell < nC(m)/(m + n - 1)$ for all $\ell \in [m]$, since otherwise the total link capacity would be greater than $C(m)$. Then the proposition follows from Lemma 2.2. ■

The algorithm A_{gfm}**:**

- Let $c^1 \geq c^2 \geq \cdots \geq c^m$;
- $\mathsf{Users} := [n]$; $m' := m$; $C(m') = \sum_{\ell=1}^{m'} c^\ell$;
- **while** $m' \geq 0$ **do**
  - **if** $c^{m'} > C(m')/(m' + n - 1)$ **then**
    - **output** the fully mixed strategies profile on Users and $[m']$ and **exit**;
  - **else**
    - $C(m' - 1) = C(m') - c^{m'}$; $m' := m' - 1$;
    - **continue**;

Figure 2: The algorithm A_{gfm} in pseudocode.

The algorithm A_{gfm} (Figure 2) finds the largest fast link set $[m']$ for which the capacity of the slowest link is greater than $C(m')/(m' + n - 1)$. By Proposition 5.2, the restriction of the routing game to $[m']$ admits a fully mixed Nash equilibrium. A_{gfm} outputs the fully mixed Nash equilibrium for the restriction of the routing game to $[m']$ and terminates. The fast link set $[m']$ output by A_{gfm} is never empty because for $m' = 1$, $c^1 > c^1/n$, for all $n \geq 2$. To establish correctness, it suffices to show that the fully mixed Nash equilibrium for the fast link set $[m']$ output by A_{gfm} remains a (generalized fully mixed) Nash equilibrium when the game is extended to the entire set of links $[m]$.

Lemma 5.3 *The fully mixed Nash equilibrium for the system of all users and the fast link set* $[m']$ *output by* A_{gfm} *remains a Nash equilibrium when the game is extended to the entire set of links* $[m]$.

Proof: Let $\mathbf{P}$ be the fully mixed Nash equilibrium for the selfish routing game on Users and $[m']$. We prove that $\mathbf{P}$ (completed with $p_i^j = 0$ for all $i \in [n]$ and $j \in [m] \setminus [m']$) is a (generalized fully mixed) Nash equilibrium with support $[m']$ for the selfish routing game on Users and $[m]$.

The claim is trivial if $m' = m$. Otherwise, Lemma 2.3 implies that the minimum expected latency cost of any user i is $\lambda_i = (m' + n - 1)/C(m')$. Since A_{gfm} does not include link $m' + 1$ in the support of $\mathbf{P}$, $c^{m'+1} \leq C(m' + 1)/(m' + n)$. Therefore,

$$c^{m'+1}(m' + n - 1) \leq C(m') \Rightarrow \frac{1}{c^{m'+1}} \geq \frac{m'+n-1}{C(m')} = \lambda_i$$

Furthermore, for all $j > m' + 1$, $1/c^j \geq \lambda_i$ because $c^{m'+1} \geq c^j$. Therefore, no user has an incentive to deviate to some slower link in $[m] \setminus [m']$ and $\mathbf{P}$ is a generalized fully mixed Nash equilibrium for the routing game on Users and $[m]$. ■

The time complexity of A_{gfm} is dominated by the time to sort the links in non-increasing capacity order. The **while** loop is executed at most $m - 1$ times and each iteration takes constant time. When the support is found, the fully mixed strategies profile $\mathbf{P}$ is computed in $O(m)$ time using Lemma 2.2, since all users have the same strategy. Therefore, the time complexity of A_{gfm} is $O(m \log m)$. ■

Next we establish the uniqueness of the generalized fully mixed equilibrium.

Theorem 5.4 *Consider a selfish routing game Γ on n users of identical traffic and m capacitated parallel links. The equilibrium output by A_{gfm} is the unique generalized fully mixed equilibrium of Γ.*

Proof: Let $\mathbf{P}$ be the generalized fully mixed Nash equilibrium output by A_{gfm}, let $[m']$ be its support, and let $\mathbf{P}'$ be any generalized fully mixed Nash equilibrium of Γ with support $S \subseteq [m]$.

We first show that S is a fast link set. To reach a contradiction, let us assume that S is not a fast link set, i.e. there is some $k \in [m-1]$ such that $k \notin S$ and $k+1 \in S$. By hypothesis, $\mathbf{P}'$ is a fully mixed Nash equilibrium of the restriction of Γ to S. Thus by Lemma 2.3, the minimum expected latency cost λ_i of any user i is equal to $(|S|+n-1)/\left(\sum_{j\in S} c^j\right)$, and by Proposition 5.2, $c^{k+1} > \sum_{j\in S} c^j/(|S|+n-1)$. Using $c^k \geq c^{k+1}$, we obtain that

$$\frac{1}{c^k} < \frac{|S|+n-1}{\sum_{j\in S} c^j} = \lambda_i \tag{1}$$

Since the left-hand side of (1) is equal to the latency cost of any user deviating to link k, (1) contradicts the hypothesis that $\mathbf{P}'$ is a Nash equilibrium of Γ.

¿From now on, we assume that S is a fast link set. By the analysis of A_{gfm}, no fast link set $[q]$, where $q > m'$, is the support of a generalized fully mixed Nash equilibrium of Γ. Hence $S \subseteq [m']$. We prove that $S = [m']$ by contradiction. Let us assume that $S \subset [m']$, and let $q < m'$ be the last link in S. Since $\mathbf{P}'$ is a fully mixed Nash equilibrium of the restriction of Γ to $[q]$, by Lemma 2.3, the minimum expected latency cost of any user i in $\mathbf{P}'$ is $\lambda'_i = (q+n-1)/C(q)$. On the other hand, by the description of A_{gfm},

$$c^{m'}(m'+n-1) > C(m') \tag{2}$$

Adding $\sum_{j=q+1}^{m'} c^j \geq (m'-q)c^{m'}$ to (2), we obtain that

$$c^{m'}(q+n-1) > C(q) \Rightarrow \tfrac{1}{c^{m'}} < \tfrac{q+n-1}{C(q)} = \lambda'_i\,,$$

which contradicts the hypothesis that $\mathbf{P}'$ is a Nash equilibrium of Γ. ■

We conclude this section with an alternative characterization of the support of the generalized fully mixed Nash equilibrium.

Proposition 5.5 *Consider a selfish routing game on n users of identical traffic and m capacitated parallel links. The support of the generalized fully mixed Nash equilibrium coincides with the set $S = \{\ell \in [m] : c^\ell > C(\ell)/(\ell+n-1)\}$.*

Proof: Let $m' \in [m]$ be the largest index such that $c^{m'} > C(m')/(m'+n-1)$. By the analysis of A_{gfm}, $[m']$ is the support of the generalized fully mixed Nash equilibrium. To establish that $S \subseteq [m']$, we observe that the links excluded from $[m']$ by A_{gfm} are also excluded from S. To show that $[m'] \subseteq S$, we observe that if some link $\ell \geq 2$ belongs to S, then link $\ell - 1$ also belongs to S. In particular,

$$C(\ell) < (\ell+n-1)c^\ell \Rightarrow C(\ell-1) < (\ell+n-2)c^\ell \leq (\ell+n-2)c^{\ell-1}$$

Since $m' \in S$ by definition, every link in $[m']$ belongs to S. ■

6 Approximating the Social Cost of the Worst Nash Equilibrium

In this section, we consider selfish routing games on users of identical traffic and capacitated links and prove that the social cost of the generalized fully mixed Nash equilibrium is within a constant factor of the social cost of the worst Nash equilibrium. The remainder of this section is devoted to the proof of the following:

Theorem 6.1 *For a selfish routing game on users of identical traffic and capacitated parallel links, the social cost of the worst Nash equilibrium is at most* 33.06 *times the social cost of the generalized fully mixed Nash equilibrium.*

6.1 Outline of the Proof

We start with some basic properties of Nash equilibria allowing the comparison of the social cost of the worst Nash equilibrium to the social cost of the generalized fully mixed equilibrium (cf. Section 6.2). A few important properties are that the set of non-solo links in the support of any Nash equilibrium is a subset of the support of the generalized fully mixed Nash equilibrium (cf. Proposition 6.7), and that the expected latency of any non-solo link in any Nash equilibrium is at most twice the expected latency of the same link in the generalized fully mixed Nash equilibrium (cf. Proposition 6.8).

After justifying some simplifying assumptions and introducing some notation (cf. Section 6.3), we proceed to analyze the tails of the distribution of maximum link latency in the generalized fully mixed Nash equilibrium (cf. Section 6.4) and in an arbitrary Nash equilibrium (cf. Section 6.5). In Section 6.4, we consider the generalized fully mixed Nash equilibrium and establish a lower bound on the probability that the maximum link latency is no less than a given value (cf. Lemma 6.11). Thus we obtain a strong lower bound on the social cost of the generalized fully mixed Nash equilibrium (cf. Lemma 6.12). In Section 6.5, we consider an arbitrary Nash equilibrium and establish an upper bound on the probability that the maximum link latency is no less than a given value (cf. Lemma 6.13). Combining Lemma 6.11 with Lemma 6.13, we derive an upper bound on the social cost of any Nash equilibrium in terms of our lower bound on the social cost of the generalized fully mixed Nash equilibrium (cf. Lemma 6.14).

6.1.1 Notation

Throughout this section, we consider a selfish routing game on $n \geq 2$ users of identical traffic and $m \geq 2$ capacitated parallel links. We assume, without loss of generality, that the links are indexed in non-increasing capacity order, i.e. $c^1 \geq c^2 \geq \ldots \geq c^m$, and that all traffics are equal to 1. For any $k \in [m]$, we let $C(k)$ denote the total capacity of the fastest k links, i.e. $C(k) = \sum_{j=1}^{k} c^j$. We let $e = 2.718\ldots$ denote the basis of natural logarithms and $\exp(x) = e^x$.

We let $\overline{\mathbf{P}}$ denote the generalized fully mixed Nash equilibrium, and let $\mathbf{P}$ denote an arbitrary Nash equilibrium. In general, we use overlined symbols to refer to the quantities related to $\overline{\mathbf{P}}$ and plain symbols to refer to the quantities related to $\mathbf{P}$. For example, $\overline{\Lambda}^\ell$ denotes the expected latency of link ℓ in the generalized fully mixed Nash equilibrium, and Λ^ℓ denotes the expected latency of ℓ in $\mathbf{P}$.

6.2 Basic Properties of Nash Equilibria

In this section, we prove some basic properties of the generalized fully mixed Nash equilibrium and of Nash equilibria in general. We start with a lower bound on the expected link latencies and a preliminary lower bound on the social cost of the generalized fully mixed Nash equilibrium.

Proposition 6.2 *Let $[m']$ denote the support of the generalized fully mixed Nash equilibrium $\overline{\mathbf{P}}$, and let $\overline{\Lambda}^\ell$ denote the expected latency of link ℓ in $\overline{\mathbf{P}}$. Then, $\overline{\Lambda}^\ell > \frac{m'+n-1}{C(m')} - \frac{1}{c^\ell}$ for all $\ell \in [m']$, and $\overline{\Lambda}^\ell = 0$ for all $\ell \in [m] \setminus [m']$. Moreover, $\overline{\Lambda}^1 \geq \cdots \geq \overline{\Lambda}^m$.*

Proof: Lemma 2.2 implies that for any user $i \in [n]$, $\overline{p}_i^\ell = \frac{c^\ell}{C(m')} + \frac{m'c^\ell - C(m')}{(n-1)C(m')}$ for all $\ell \in [m']$, while $\overline{p}_i^\ell = 0$ for all $\ell \in [m] \setminus [m']$. Since $\overline{p}_i^\ell$'s are the same for all users, for any link $\ell \in [m']$,

$$\overline{\Lambda}^\ell = \frac{n}{c^\ell}\left(\frac{c^\ell}{C(m')} + \frac{m'c^\ell - C(m')}{(n-1)C(m')}\right) = \frac{n}{n-1}\left(\frac{m'+n-1}{C(m')} - \frac{1}{c^\ell}\right) > \frac{m'+n-1}{C(m')} - \frac{1}{c^\ell} \ ,$$

while $\overline{\Lambda}^\ell = 0$, for all $\ell \in [m] \setminus [m']$. In addition, since $\Lambda^\ell = \frac{n}{n-1}(\frac{m'+n-1}{C(m')} - \frac{1}{c^\ell}) > 0$ for all $\ell \in [m']$, and $c^1 \geq \cdots \geq c^{m'}$, we conclude that $\overline{\Lambda}^1 \geq \cdots \geq \overline{\Lambda}^m$. ■

Proposition 6.3 *The social cost of the generalized fully mixed Nash equilibrium with support $[m']$ is at least $\max\{\overline{\Lambda}^1, \frac{m'+n-1}{2C(m')}\}$.*

Proof: Since the social cost cannot be less than the expected latency of any link, we obtain a lower bound of $\overline{\Lambda}^1$. If we assume that $\frac{m'+n-1}{2C(m')} > \overline{\Lambda}^1$, then using that $\overline{\Lambda}^1 > \frac{m'+n-1}{C(m')} - \frac{1}{c^1}$, we obtain that $\frac{1}{c^1} > \frac{m'+m-1}{2C(m')}$. Since the social cost is at least $1/c^1$, we obtain a lower bound of $\frac{m'+m-1}{2C(m')}$. ■

We continue with some basic properties that hold for all Nash equilibria. First we show that the minimum expected latency cost of any user in any Nash equilibrium does not exceed the ratio of $k+n-1$ to the total capacity of the fastest k links, for any $k \in [m]$.

Proposition 6.4 *Let λ_i be the minimum expected latency cost of any user i in any Nash equilibrium $\mathbf{P}$. Then for any $k \in [m]$, $\lambda_i \leq (k+n-1)/C(k)$.*

Proof: To reach a contradiction, let us assume that there is some $k \in [m]$ and some user i such that $\lambda_i > (k+n-1)/C(k)$. Since $\mathbf{P}$ is a Nash equilibrium, for every link $j \in [m]$, $\lambda_i \leq \lambda_i^j = \Lambda^j + (1-p_i^j)/c^j$. Therefore, for every link $j \in [m]$,

$$c^j \frac{k+n-1}{C(k)} < c^j \Lambda^j + 1 - p_i^j \tag{3}$$

Summing up (3) over the fastest k links and using the definitions of $C(k)$ and Λ^j, we obtain that

$$\begin{aligned} k+n-1 &< \sum_{q\in[n]} \sum_{j=1}^{k} p_q^j + k - \sum_{j=1}^{k} p_i^j \\ &= \sum_{q\in[n]\setminus\{i\}} \sum_{j=1}^{k} p_q^j + k \quad \leq \quad n-1+k \ , \end{aligned}$$

a contradiction. ■

We proceed to state two useful corollaries of Proposition 6.4. The first corollary shows that in any Nash equilibrium, the expected latency of any link is at most $(m' + n - 1)/C(m')$. The second corollary establishes the same upper bound on the (observed) latency of any solo link.

Corollary 6.5 *Let Λ^ℓ be the expected latency of any link $\ell \in [m]$ in any Nash equilibrium $\mathbf{P}$. Then, $\Lambda^\ell \leq (m'+n-1)/C(m')$, where $[m']$ is the support of the generalized fully mixed Nash equilibrium.*

Proof: If $\Lambda^\ell > 0$, there is a user i with $p_i^\ell > 0$. Since $\mathbf{P}$ is a Nash equilibrium, $\lambda_i = \lambda_i^\ell \geq \Lambda^\ell$. Applying Proposition 6.4 with $k = m'$, we obtain that $\Lambda^\ell \leq \lambda_i \leq (m' + n - 1)/C(m')$. ■

Corollary 6.6 *For every solo link ℓ in a Nash equilibrium $\mathbf{P}$, the (observed) latency of ℓ is at most $(m' + n - 1)/C(m')$, where $[m']$ is the support of the generalized fully mixed Nash equilibrium.*

Proof: Since ℓ is solo, its (observed) latency is at most $1/c^\ell$. Let i be the only user in $view(\ell)$. Since $\mathbf{P}$ is a Nash equilibrium, $\lambda_i = \lambda_i^\ell = 1/c^\ell$. Applying Proposition 6.4 with $k = m'$, we obtain that $1/c^\ell = \lambda_i \leq (m' + n - 1)/C(m')$. ■

In combination with Proposition 5.5, the following proposition shows that if a Nash equilibrium contains some link $\ell \notin [m']$ in its support, then ℓ is solo. Since the (observed) latency of a solo link is at most $2\,\mathsf{SC}(\mathbf{1}, \overline{\mathbf{P}})$ (see Proposition 6.3 and Corollary 6.6), we can ignore all links excluded from the support of the generalized fully mixed Nash equilibrium.

Proposition 6.7 *The support of any Nash equilibrium $\mathbf{P}$ does not include any link ℓ with $c^\ell < C(\ell)/(\ell + n - 1)$. If the support of $\mathbf{P}$ includes a link ℓ with $c^\ell = C(\ell)/(\ell + n - 1)$, then ℓ is solo.*

Proof: Let ℓ be any link in the support of $\mathbf{P}$. Then there is a user i with $p_i^\ell > 0$ and $1/c^\ell \leq \lambda_i$. Applying Proposition 6.4 with $k = \ell$, we obtain that $\lambda_i \leq (\ell + n - 1)/C(\ell)$. Therefore, for any link ℓ in the support of $\mathbf{P}$, $c^\ell \geq C(\ell)/(\ell + n - 1)$. Moreover, if $c^\ell = C(\ell)/(\ell + n - 1)$, then $1/c^\ell = \lambda^i$, which can happen only if ℓ is solo. ■

Next we show that the expected latency of any non-solo link in any Nash equilibrium is at most twice the expected latency of the same link in the generalized fully mixed Nash equilibrium.

Proposition 6.8 *Let $\mathbf{P}$ be any Nash equilibrium and $\overline{\mathbf{P}}$ be the generalized fully mixed Nash equilibrium. For every link ℓ non-solo in $\mathbf{P}$,*

$$\Lambda^\ell < \left(1 + \frac{1}{|view(\ell)| - 1}\right) \overline{\Lambda}^\ell \ ,$$

where Λ^ℓ (resp. $\overline{\Lambda}^\ell$) denotes the expected latency of ℓ in $\mathbf{P}$ (resp. $\overline{\mathbf{P}}$), and $view(\ell)$ is the set of users whose support in $\mathbf{P}$ includes ℓ.

Proof: Let $[m']$ be the support of $\overline{\mathbf{P}}$. For simplicity of notation, let $k^\ell = |view(\ell)|$. Since $k^\ell > 1$, there is a user $i \in view(\ell)$ with $p_i^\ell \in (0, \frac{\Lambda^\ell c^\ell}{k^\ell}]$. Therefore, the minimum expected latency cost of i is

$$\lambda_i = \lambda_i^\ell = \Lambda^\ell + \tfrac{1-p_i^\ell}{c^\ell} \geq \tfrac{k^\ell - 1}{k^\ell} \Lambda^\ell + \tfrac{1}{c^\ell} \tag{4}$$

Applying Proposition 6.4 with $k = m'$, we obtain that $\lambda_i \leq (m' + n - 1)/C(m')$. Combining this inequality with (4), we obtain that

$$\tfrac{k^\ell - 1}{k^\ell} \Lambda^\ell \leq \tfrac{m'+n-1}{C(m')} - \tfrac{1}{c^\ell} < \overline{\Lambda}^\ell ,$$

where the last inequality follows from Proposition 6.2. Therefore, $\Lambda^\ell < (1 + \frac{1}{k^\ell - 1})\overline{\Lambda}^\ell$. ■

6.3 Preliminaries

In the following, we assume, without loss of generality, that the support of the generalized fully mixed Nash equilibrium coincides with the entire set of links, i.e. $m' = m$, and that $n \geq 3 \ln m$.

For the former assumption, Proposition 5.5 and Proposition 6.7 imply that we can ignore every link not in the support of the generalized fully mixed Nash equilibrium. More precisely, we can ignore every link ℓ with $c^\ell < C(\ell)/(\ell + n - 1)$, because it is not included in the support of any Nash equilibrium (see Proposition 6.7). A link ℓ with $c^\ell = C(\ell)/(\ell + n - 1)$ is not included in the support of the generalized fully mixed Nash equilibrium (see Proposition 5.5), but it may be included in the support of some other Nash equilibrium as a solo link (see Proposition 6.7). However, the (observed) latency of any solo link is at most $\frac{m'+n-1}{C(m')}$ (see Corollary 6.6). Since the social cost of the generalized fully mixed Nash equilibrium is at least $\frac{m'+n-1}{2C(m')}$ (see Proposition 6.3), solo links cannot increase the maximum link latency above $2\,\mathsf{SC}(\mathbf{1}, \overline{\mathbf{P}})$. Therefore we can ignore, without loss of generality, all links excluded from the support of the generalized fully mixed Nash equilibrium, and assume that $m' = m$. For simplicity of notation, we use m instead of m' in what follows.

The latter assumption excludes some trivial cases only and can be made without loss of generality. In particular, if $n < 3 \ln m$, the social cost of any Nash equilibrium is at most $16\,\frac{m+n-1}{C(m)}$, while the social cost of the generalized fully mixed Nash equilibrium is at least $\frac{m+n-1}{2C(m)}$ (see Proposition 6.3)

For each link j and any $x \geq \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}\}$, the function $f_j(x)$ gives a lower bound on the probability that j's latency in the generalized fully mixed equilibrium is at least x (cf. Lemma 6.11):

$$f_j(x) = \begin{cases} \left(\dfrac{\overline{\Lambda}^j}{2ex} \right)^{2xc^j} & \text{if } x \leq n/c^j \\ 0 & \text{if } x > n/c^j \end{cases} \tag{5}$$

The following proposition gives a useful property of $f_j(x)$.

Proposition 6.9 *For any $j \in [m]$, and for all $z \geq 0$ and $y \geq 1$, $f_j(zy) \leq [f_j(z)]^y$.*

Proof: The claim is trivial if $f_j(zy) = 0$. Otherwise, $z \leq zy \leq n/c^j$. Using $y \geq 1$, we conclude that

$$f_j(zy) = \left(\frac{\overline{\Lambda}^j}{2ezy} \right)^{2zyc^j} \leq \left[\left(\frac{\overline{\Lambda}^j}{2ez} \right)^{2zc^j} \right]^y = [f(z)]^y$$

■

For each link j and any $x > (m+n-1)/C(m)$, the function $h_j(x)$ gives an upper bound on the probability that j's latency in Nash equilibrium $\mathbf{P}$ is at least x (cf. Lemma 6.13):

$$h_j(x) = \begin{cases} \left(\dfrac{e\Lambda^j}{x}\right)^{xc^j} & \text{if } x \leq |view(j)|/c^j \\ 0 & \text{if } x > |view(j)|/c^j \end{cases} \tag{6}$$

In the proof of Lemma 6.14, we use the following proposition and compare the social cost of Nash equilibrium $\mathbf{P}$ to the social cost of the generalized fully mixed Nash equilibrium $\overline{\mathbf{P}}$.

Proposition 6.10 *For any $j \in [m]$ and all $x \geq 8/c^m$, $h_j(x) \leq \left[f_j(\frac{7x}{16e^2})\right]^{8e^2/7}$.*

Proof: The claim is trivial if $h_j(x) = 0$. Otherwise, $x \leq |view(j)|/c^j \leq n/c^j$. Thus $|view(j)| \geq 8$, because $x \geq 8/c^m$ and $c^j \geq c^m$. Using Proposition 6.8, we obtain that

$$h_j(x) = \left(\frac{e\Lambda^j}{x}\right)^{xc^j} \leq \left(\frac{8e\overline{\Lambda}^j}{7x}\right)^{xc^j} = \left[\left(\frac{\overline{\Lambda}^j}{2e\frac{7x}{16e^2}}\right)^{2c^j\frac{7x}{16e^2}}\right]^{8e^2/7} = \left[f_j(\tfrac{7x}{16e^2})\right]^{8e^2/7}$$

For the last equality, we observe that $\frac{7x}{16e^2} \leq x \leq n/c^j$. ■

6.4 A Lower Bound on the Social Cost of the Generalized Fully Mixed Equilibrium

We proceed to establish a lower bound on the probability that the maximum latency in the generalized fully mixed Nash equilibrium is no less than given value.

Lemma 6.11 *Let $\overline{L}_{\max}$ denote the maximum link latency in the generalized fully mixed equilibrium. Then, for all $x \geq \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}\}$,*

$$\mathbb{P}\mathrm{r}[\overline{L}_{\max} \geq x] \geq 1 - \exp\left(-\sum_{j=1}^{m} f_j(x)\right)$$

Proof: For each $i \in [n]$ and $j \in [m]$, let $\overline{X}_i^j$ be the random variable indicating whether user i routes her traffic on link j in the generalized fully mixed Nash equilibrium. By Lemma 2.2, all users have the same probability of routing their traffic on every link. Therefore, $\overline{X}_i^j$ takes the value 1 with probability $\frac{\overline{\Lambda}^j c^j}{n}$ and the value 0 otherwise. For each $j \in [m]$, let $\overline{X}^j = \sum_{i=1}^n \overline{X}_i^j$ be the random variable denoting the number of users routing their traffic on link j. Then $\overline{L}_{\max} = \max_{j\in[m]}\{\overline{X}^j/c^j\}$.

For any link j and any x, where $\max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}\} \leq x \leq n/c^j$, we let $k^j = \lceil xc^j \rceil$. Then,

$$\begin{aligned} \mathbb{P}\mathrm{r}[\overline{X}^j \geq k^j] &\geq \binom{n}{k^j}\left(\frac{\overline{\Lambda}^j c^j}{n}\right)^{k^j}\left(1-\frac{\overline{\Lambda}^j c^j}{n}\right)^{n-k^j} \\ &\geq \frac{n^{k^j}}{(k^j)^{k^j}}\left(\frac{\overline{\Lambda}^j c^j}{n}\right)^{k^j}\left(1-\frac{k^j}{n}\right)^{n-k^j} \\ &\geq \left(\frac{\overline{\Lambda}^j c^j}{ek^j}\right)^{k^j} \\ &\geq \left(\frac{\overline{\Lambda}^j}{2ex}\right)^{2xc^j} \end{aligned}$$

For the second inequality, we use that $k^j \geq \overline{\Lambda}^j c^j$, because $k^j \geq xc^j$ and $x \geq \overline{\Lambda}^1 \geq \overline{\Lambda}^j$ (see Proposition 6.2). For the third inequality, we use the fact that for all $k \in [n]$, $(1-\frac{k}{n})^{n-k} \geq e^{-k}$. For the last inequality, we use that $2xc^j \geq \lceil xc^j \rceil$, since $xc^j \geq 1/2$. In particular, $xc^j \geq 1/2$ follows from the facts that: (i) $c^j > C(m)/(m+n-1)$ because link j is in the support of the generalized fully mixed Nash equilibrium, and (ii) $x \geq \frac{m+n-1}{2C(m)}$.

On the other hand, $\mathbb{P}\mathrm{r}[\overline{X}^j \geq \lceil xc^j \rceil] = 0$, for all $x > n/c^j$. Therefore, for any link j and any $x \geq \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}\}$,

$$\mathbb{P}\mathrm{r}[\overline{X}^j \geq \lceil xc^j \rceil] \geq f_j(x) \tag{7}$$

where $f_j(x)$ is defined in (5). Using the fact that in "balls and bins" experiments, the occupancy numbers are negatively associated (see e.g. [3]), we obtain that

$$\begin{aligned}
\mathbb{P}\mathrm{r}[\overline{L}_{\max} < x] = \mathbb{P}\mathrm{r}\left[\bigwedge_{j=1}^{m} (\overline{X}^j < \lceil xc^j \rceil)\right] &\leq \prod_{j=1}^{m} \mathbb{P}\mathrm{r}\left[\overline{X}^k < \lceil xc^j \rceil\right] \\
&\leq \prod_{j=1}^{m} (1 - f_j(x)) \\
&\leq \exp\left(-\sum_{j=1}^{m} f_j(x)\right)
\end{aligned}$$

For the first inequality, we use that the random variables $\overline{X}^1, \ldots, \overline{X}^m$ are negatively associated (see e.g. [3, Proposition 29 and Theorem 33]). The second inequality follows from (7). For the third inequality, we use that for all $x \geq 0$, $1 - x \leq e^{-x}$. ∎

For simplicity of notation, we introduce the function $g(x) = \sum_{j=1}^{m} f_j(x)$. The function $g(x)$ is non-negative in $[0, \infty)$, and has $g(0) = m$ and $g(x) = 0$ for all $x > n$. There is a point x^*, where $\overline{\Lambda}^m/(2e^2) \leq x^* \leq \overline{\Lambda}^1/(2e^2)$, such that $g(x)$ is non-decreasing in $[0, x^*)$ and non-increasing in $(x^*, \infty]$. The function $g(x)$ is not continuous due to the jump discontinuity in the definition of $f_j(x)$'s. However, these jumps are negligible provided that n is not very small. More precisely, for every $j \in [m]$, the j-th term of $g(x)$ jumps from $f_j(n/c^j) < (\frac{1}{2e})^{2n}$ to 0 at n/c^j. Thus each jump of $g(x)$ is less than $m(\frac{1}{2e})^{2n}$. Assuming that $n \geq 3 \ln m$ and that $m \geq 2$, we obtain that each jump of $g(x)$ is less than 0.0018. Therefore, for any $\alpha \in (1, e)$, there is at least one point $x \in (x^*, n)$ such that $g(x) \in (\ln(\alpha) - 0.0018, \ln(\alpha)]$. In the following, we let μ_α denote the smallest such value:

$$\mu_\alpha \equiv \arg\min\{x \in (x^*, n) : g(x) \in (\ln(\alpha) - 0.0018, \ln(\alpha)]\} \tag{8}$$

By the definition of μ_α, $g(\mu_\alpha) \geq \ln(\alpha) - 0.0018$. Moreover, since $g(x)$ is non-increasing in (x^*, ∞), for all $x \geq \mu_\alpha$, $g(x) \leq \ln(\alpha)$.

For simplicity of notation, we let $\mu^*_\alpha = \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}, \mu_\alpha\}$. The following lemma establishes a lower bound of $(1 - \frac{e^{0.0018}}{\alpha})\mu^*_\alpha$ on the social cost of the generalized fully mixed Nash equilibrium.

Lemma 6.12 *For any $\alpha \in (1, e)$, $\mathsf{SC}(\mathbf{1}, \overline{\mathbf{P}}) \geq (1 - \frac{e^{0.0018}}{\alpha})\mu^*_\alpha$, where $\mu^*_\alpha = \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}, \mu_\alpha\}$.*

Proof: In Proposition 6.3, we show that $\mathsf{SC}(\mathbf{1}, \overline{\mathbf{P}}) \geq \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}\}$. If $\mu_\alpha > \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}\}$, we apply Lemma 6.11 with $x = \mu_\alpha$ and obtain that

$$\mathbb{P}\mathrm{r}[\overline{L}_{\max} \geq \mu_\alpha] \geq 1 - \exp(-g(\mu_\alpha)) \geq 1 - \frac{e^{0.0018}}{\alpha}$$

because $g(\mu_\alpha) \geq \ln(\alpha) - 0.0018$. Therefore, the social cost of the generalized fully mixed Nash equilibrium is at least $(1 - \frac{e^{0.0018}}{\alpha}) \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}, \mu_\alpha\}$. ■

6.5 An Upper Bound on the Social Cost of any Nash Equilibrium

Next we consider an arbitrary Nash equilibrium $\mathbf{P}$ and obtain an upper bound on the probability that the maximum link latency is no less than a given value.

Lemma 6.13 *Let $\mathbf{P}$ be any Nash equilibrium, and let $L_{\max}$ denote the maximum link latency in $\mathbf{P}$. Then, for all $x > (m+n-1)/C(m)$,*

$$\mathbb{P}\mathrm{r}[L_{\max} \geq x] \leq \sum_{j=1}^{m} h_j(x)$$

Proof: For each $i \in [n]$ and $j \in [m]$, let X_i^j be the random variable indicating whether user i routes her traffic on link j in $\mathbf{P}$. X_i^j takes the value 1 with probability p_i^j and the value 0 otherwise. For each $j \in [m]$, let $X^j = \sum_{i \in view(j)} X_i^j$ be the random variable denoting the number of users routing their traffic on link j. Then $L_{\max} = \max_{j \in [m]}\{X^j/c^j\}$. By linearity of expectation, $\mathbb{E}[X^j] = \Lambda^j c^j$.

For any link j and any $x \in (\Lambda^j, |view(j)|/c^j]$, we apply the Chernoff bound[5] and obtain that

$$\mathbb{P}\mathrm{r}[X^j \geq xc^j] \leq \left(\frac{e\Lambda^j}{x}\right)^{xc^j}$$

On the other hand, $\mathbb{P}\mathrm{r}[X^j \geq xc^j] = 0$ for all $x > |view(j)|/c^j$. Thus for any link j and any $x > \Lambda^j$, $\mathbb{P}\mathrm{r}[X^j \geq xc^j] \leq h_j(x)$, where $h_j(x)$ is defined in (6).

Observing that $x > (m+n-1)/C(m)$ implies that $x > \Lambda^j$ for all $j \in [m]$ (see Corollary 6.5), and applying the union bound, we conclude that for any $x > (m+n-1)/C(m)$,

$$\mathbb{P}\mathrm{r}[L_{\max} \geq x] = \mathbb{P}\mathrm{r}\left[\bigvee_{j=1}^{m} (X^j \geq xc^j)\right] \leq \sum_{j=1}^{m} h_j(x)$$

■

In the following lemma, we obtain an upper bound on the social cost of $\mathbf{P}$ in terms of μ_α^*.

[5] We use the following standard form of the Chernoff bound (see e.g. [23]) with $1+\delta = x/\Lambda^j$: Let $X_1, X_2, \ldots, X_n$ be independent $0/1$ random variables, let $X = \sum_{i=1}^n X_i$, and let $\mathbb{E}[X]$ denote the expectation of X. Then for any $\delta > 0$,

$$\mathbb{P}\mathrm{r}[X \geq (1+\delta)\mathbb{E}[X]] < \left(\frac{e^\delta}{(1+\delta)^{(1+\delta)}}\right)^{\mathbb{E}[X]}$$

Lemma 6.14 *Let* $\mathbf{P}$ *be any Nash equilibrium. For any* $\alpha \in (1, e)$*,*

$$\mathsf{SC}(\mathbf{1}, \mathbf{P}) \leq \left(\frac{16e^2}{7} - \frac{2\ln(\alpha)^{8e^2/7}}{\ln\ln(\alpha)} \right) \mu_\alpha^*$$

Proof: Using Lemma 6.13, we bound the social cost of $\mathbf{P}$ as follows:

$$\begin{aligned} \mathsf{SC}(\mathbf{1}, \mathbf{P}) = \mathbb{E}[L_{\max}] &= \int_0^\infty \mathbb{P}\mathrm{r}[L_{\max} \geq x]dx \\ &\leq \tfrac{16}{7}\, e^2 \mu_\alpha^* + \int_{\frac{16}{7} e^2 \mu_\alpha^*}^\infty \mathbb{P}\mathrm{r}[L_{\max} \geq x]dx \\ &\leq \tfrac{16}{7}\, e^2 \mu_\alpha^* + \int_{\frac{16}{7} e^2 \mu_\alpha^*}^\infty \sum_{j=1}^m h_j(x)dx \end{aligned} \tag{9}$$

For the second equality, we use that the expectation of a non-negative random variable X is given by $\mathbb{E}[X] = \int_0^\infty \mathbb{P}\mathrm{r}[X \geq x]dx$. The first inequality holds because $\mathbb{P}\mathrm{r}[L_{\max} \geq x] \leq 1$ for all $x \geq 0$. For the second inequality, we apply Lemma 6.13. Using $\mu_\alpha^* \geq \frac{m+n-1}{2C(m)} > 0$ we obtain that $\frac{16}{7} e^2 \mu_\alpha^* > (m+n-1)/C(m)$ as required by the hypothesis of Lemma 6.13.

To conclude the proof, we establish an upper bound on the last term of (9):

$$\begin{aligned} \int_{\frac{16}{7} e^2 \mu_\alpha^*}^\infty \sum_{j=1}^m h_j(x)dx &\leq \int_{\frac{16}{7} e^2 \mu_\alpha^*}^\infty \sum_{j=1}^m \left[f_j(\tfrac{7x}{16e^2})\right]^{8e^2/7} dx \\ &= \tfrac{16}{7}\, e^2 \mu_\alpha^* \int_1^\infty \textstyle\sum_{j=1}^m \left[f_j(\mu_\alpha^* y)\right]^{8e^2/7} dy \\ &\leq \tfrac{16}{7}\, e^2 \mu_\alpha^* \int_1^\infty \sum_{j=1}^m [f_j(\mu_\alpha^*)]^{8e^2 y/7} dy \\ &\leq \tfrac{16}{7}\, e^2 \mu_\alpha^* \int_1^\infty [g(\mu_\alpha^*)]^{8e^2 y/7}\, dy \\ &\leq \tfrac{16}{7}\, e^2 \mu_\alpha^* \int_1^\infty \ln(\alpha)^{8e^2 y/7} dy \\ &= -\frac{2\ln(\alpha)^{8e^2/7}}{\ln\ln(\alpha)} \mu_\alpha^* \end{aligned}$$

For the first inequality, we apply Proposition 6.10 for all $j \in [m]$. Since $\mu_\alpha^* \geq \frac{m+n-1}{2C(m)} > \frac{1}{2c^m}$, because m is in the support of the generalized fully mixed Nash equilibrium, $x \geq \frac{16}{7} e^2 \mu_\alpha^* \geq 8/c^m$ as required by the hypothesis of Proposition 6.10. The first equality follows by changing the variable of integration to $y = \frac{7x}{16e^2 \mu_\alpha^*}$. The second inequality follows from Proposition 6.9, since $y \geq 1$. For the third inequality, we use the fact that for all $x_1, \ldots, x_m \geq 0$ and all $z \geq 1$, $x_1^z + \cdots + x_m^z \leq (x_1 + \cdots + x_m)^z$. Therefore, for all $y \geq 1$,

$$\sum_{j=1}^m [f_j(\mu_\alpha^*)]^{8e^2 y/7} \leq \left[\sum_{j=1}^m f_j(\mu_\alpha^*) \right]^{8e^2 y/7} = [g(\mu_\alpha^*)]^{8e^2 y/7}$$

where $g(x) = \sum_{j=1}^m f_j(x)$. For the fourth inequality, we use that $g(\mu_\alpha^*) \leq \ln(\alpha)$, since $\mu_\alpha^* \geq \mu_\alpha$. For the last equality, we calculate the integral and use that $\alpha \in (1, e)$. ■

Combining Lemma 6.12 with Lemma 6.14, we obtain that the social cost of any Nash equilibrium $\mathbf{P}$ is within a constant factor of the social cost of the generalized fully mixed Nash equilibrium $\overline{\mathbf{P}}$. More precisely, for any $\alpha \in (1, e)$,

$$\mathsf{SC}(\mathbf{1}, \mathbf{P}) \leq \left(\frac{16e^2}{7} - \frac{2\ln(\alpha)^{8e^2/7}}{\ln\ln(\alpha)} \right) \frac{\alpha}{\alpha - e^{0.0018}} \,\mathsf{SC}(\mathbf{1}, \overline{\mathbf{P}})$$

Using $\alpha = 2.17$, we conclude that $\mathsf{SC}(\mathbf{1}, \mathbf{P}) \leq 33.06\,\mathsf{SC}(\mathbf{1}, \overline{\mathbf{P}})$. ■

7 Computing the Social Cost of a Mixed Nash Equilibrium

7.1 The Complexity of Computing the Social Cost of a Mixed Nash Equilibrium

On the negative side, we prove that it is $\#\mathcal{P}$-complete to compute the social cost of a given mixed Nash equilibrium.

Theorem 7.1 NASH EQUILIBRIUM SOCIAL COST *is $\#\mathcal{P}$-complete even for identical links.*

Proof: Membership in $\#\mathcal{P}$ follows from the definition of social cost and the fact that the probabilities in a mixed Nash equilibrium are rational (see e.g. Proposition 2.1 and [16, Section 2]). To show that NASH EQUILIBRIUM SOCIAL COST is $\#\mathcal{P}$-complete, we use a reduction from the problem of computing the probability that the sum of n independent random variables does not exceed a given threshold.

More precisely, let $J = \{w_1, \ldots, w_n\}$ be a set of n integer weights and let $C \geq \sum_{i=1}^n w_i/2$ be an integer. Counting the number of J's subsets with total weight at most C is $\#\mathcal{P}$-complete because it is equivalent to counting the number of feasible solutions of the corresponding KNAPSACK instance (see e.g. [28]). Therefore, given n independent random variables $Y_1(w_1, 1/2), \ldots, Y_n(w_n, 1/2)$, where $Y_i(w_i, 1/2)$ takes the value w_i with probability $1/2$ and the value 0 otherwise, it is $\#\mathcal{P}$-complete to compute the probability that $Y = \sum_{i=1}^n Y_i$ is at most C (see also [11, Theorem 2.1]). Next we show that $\mathbb{P}\mathrm{r}[Y \leq C]$ can be recovered by two calls to an oracle returning the social cost of a given mixed Nash equilibrium.

Given n random variables $Y_1(w_1, 1/2), \ldots, Y_n(w_n, 1/2)$ and an integer $C \geq \sum_{i=1}^n w_i/2$, we construct a selfish routing game on $n+1$ users and 3 identical links. For every $i \in [n]$, the traffic of user i is w_i and the traffic of user $n+1$ is C. We consider a mixed strategies profile $\mathbf{P}$ where user $n+1$ selects link 3 with certainty (i.e. $p_{n+1}^3 = 1$ and $p_{n+1}^1 = p_{n+1}^2 = 0$), and the remaining users select one of the first two links equiprobably (i.e. $p_i^1 = p_i^2 = 1/2$ and $p_i^3 = 0$ for all $i \in [n]$). Since $C \geq \sum_{i=1}^n w_i/2$, $\mathbf{P}$ is a Nash equilibrium. Since w_i's are integral, the social cost of $\mathbf{P}$ is

$$\mathsf{SC}_1 = C + 2 \sum_{B=C+1}^{\infty} \mathbb{P}\mathrm{r}[Y \geq B]\,, \tag{10}$$

where we use that the expectation of a non-negative integral random variable X is given by $\mathbb{E}[X] = \sum_{i=1}^{\infty} \mathbb{P}\mathrm{r}[X \geq i]$ (see e.g [23, Proposition C.7]) and that for all $B > \sum_{i=1}^n w_i/2$, the events that link j, $j \in \{1, 2\}$, has latency at least B are mutually exclusive.

Increasing the traffic of user $n+1$ to $C+1$, we obtain a slightly different game for which $\mathbf{P}$ remains a Nash equilibrium. As before, the social cost of $\mathbf{P}$ for the new game is

$$\mathsf{SC}_2 = C + 1 + 2 \sum_{B=C+2}^{\infty} \mathrm{I\!Pr}[Y \geq B] \tag{11}$$

Combining (10) and (11), we obtain that

$$\mathsf{SC}_2 - \mathsf{SC}_1 = 1 - 2\,\mathrm{I\!Pr}[Y \geq C+1] \tag{12}$$

Since w_i's and C are integers, $\mathrm{I\!Pr}[Y \leq C] = 1 - \mathrm{I\!Pr}[Y \geq C+1]$. Therefore, (12) implies that $\mathrm{I\!Pr}[Y \leq C] = (\mathsf{SC}_2 - \mathsf{SC}_1 + 1)/2$. ■

7.2 Approximating the Social Cost of a Mixed Nash Equilibrium

On the positive side, we get around the $\#\mathcal{P}$-completeness result by formulating a FPRAS that approximates the social cost of any given mixed Nash equilibrium.

Theorem 7.2 *Consider the model of uniform capacities. Then there exists a fully polynomial-time randomized approximation scheme for* NASH EQUILIBRIUM SOCIAL COST.

Proof: We define an efficiently samplable random variable which accurately estimates the social cost of the given Nash equilibrium $\mathbf{P}$ on the given traffic vector $\mathbf{w}$. More precisely, we perform the following experiment, where N is a fixed integer that will be specified later:

> "Repeat N times the random experiment of assigning each user to a link in its support according to the given Nash probabilities. For each experiment $i \in [N]$, let $L^i_{\max}$ be the (measured) maximum link latency. Output the mean $\sum_{i=1}^{N} L^i_{\max}/N$ of the measured values."

Let $L_{\max}$ be the random variable denoting the outcome of the algorithm. $L_{\max}$ is the mean of N identically distributed independent random variables. Its expectation is equal to the social cost of $\mathbf{P}$ and its variance is bounded is at most $n^2 w^2_{\max}/N$, where $w_{\max}$ denotes the maximum traffic in $\mathbf{w}$. Applying Chebyshev's inequality (see e.g. [23]), we obtain that for any $\varepsilon > 0$,

$$\mathrm{I\!Pr}[|L_{\max} - \mathsf{SC}(\mathbf{w},\mathbf{P})| \geq \varepsilon\,\mathsf{SC}(\mathbf{w},\mathbf{P})] \leq \frac{\mathrm{Var}[L_{\max}]}{\varepsilon^2\,\mathsf{SC}^2(\mathbf{w},\mathbf{P})} \leq \frac{n^2}{\varepsilon^2 N} \ ,$$

where the last inequality follows from $\mathrm{Var}[L_{\max}] \leq n^2 w^2_{\max}/N$ and $\mathsf{SC}(\mathbf{w},\mathbf{P}) \geq w_{\max}$. Therefore, for all $\varepsilon > 0$ and $N \geq 4n^2/\varepsilon^2$, the probability that the outcome of the algorithm is within a factor of $(1 \pm \varepsilon)$ from $\mathsf{SC}(\mathbf{w},\mathbf{P})$ is at least $3/4$. ■

8 Conclusion and Open Problems

We have presented a comprehensive collection of algorithmic, hardness and structural results for the computation of Nash equilibria for a specific game that models selfish routing over a set of parallel links. Our work leaves open numerous interesting questions that are directly related to our results. We list a few of them here.

- What is the complexity of computing the supports of a pure Nash equilibrium? Theorem 3.2 shows that it is $O(n \lg n + nm) = O(n \max\{\lg n, m\})$. Can this be further improved?

- Recall that the $\mathcal{NP}$-hardness proof (Theorem 3.4) for WORST NASH EQUILIBRIUM SUPPORTS uses different traffics, while it assumes uniform capacities. What happens in the model of identical traffics and arbitrary capacities? Does the problem remain $\mathcal{NP}$-hard?

- Consider the *specific* pure Nash equilibria that are computed by the algorithm that is implicit in the proof of Theorem 3.1 and the algorithm A_{pure} in the proof of Theorem 3.2. It would be interesting to study how well these specific pure Nash equilibria approximate the worst one (in terms of social cost).

- What is the complexity of computing the supports of a generalized fully mixed Nash equilibrium? Theorem 5.1 shows that it is $O(m \lg m)$ in the case where all traffics are identical. Can this be further improved? Nothing is known about the general case, where traffics are not necessarily identical.

- What is the complexity of computing the supports of a generalized fully mixed Nash equilibrium? Theorem 5.1 shows that it is $O(m^2)$ in the case where all traffics are identical. Can this be further improved? Nothing is known about the general case, where traffics are not necessarily identical. Is the general problem still in $\mathcal{P}$?

- It is tempting to conjecture that Theorem 4.2 holds for all values of $n \geq 2$. To prove this, it may require to obtain an improved understanding of the combinatorial structure of fully mixed Nash equilibria.

Besides these directly related open problems, we feel that the most significant extension of our work would be to study other specific games and classify their instances according to the computational complexity of computing the Nash equilibria of the game. We hope that our work provides an initial solid ground for such studies.

Acknowledgments. We thank Kurt Mehlhorn for contributing Theorem 3.1 to our work. We thank Burkhard Monien for some extremely helpful discussions on our work, and for contributing to our work an improvement to the time complexity (from $\Theta(m^2)$ to $\Theta(m \lg m)$) of Algorithm A_{gfm} (Theorem 5.1). We thank Peter Sanders Berthold Vöcking for a fruitful discussion right after a seminar talk on Nash equilibria for selfish routing given by Paul Spirakis at Max-Planck-Institut für Informatik. That discussion contributed some initial ideas for the design of Algorithm A_{pure} (Theorem 3.2). Dimitris thanks Peter Sanders for helpful discussions on "balls and bins" experiments that contributed the initial idea of the proof of Theorem 6.1. We also thank Rainer Feldman for helpful comments on an earlier version of this paper. Finally, we thank Costas Busch for bringing to our attention the allegoric movie *A Beautiful Mind*.

References

[1] A. Czumaj and B. Vöcking, "Tight Bounds for Worst-Case Equilibria," *Proceedings of the 13th Annual ACM Symposium on Discrete Algorithms,* January 2002, to appear.

[2] X. Deng, C. Papadimitriou and S. Safra, "On the Complexity of Equilibria," *Proceedings of the 34th Annual ACM Symposium on Theory of Computing,* to appear.

[3] D. Dubhashi and D. Ranjan. Balls and Bins: A Study in Negative Dependence. *Random Structures and Algorithms*, 13(2):99–124, 1998.

[4] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of $\mathcal{NP}$-Completeness*. W.H. Freeman, 1979.

[5] I. Gilboa and E. Zemel, "Nash and Correlated Equilibria: Some Complexity Considerations," *Games and Economic Behavior,* Vol. 1, pp. 80–93, 1989.

[6] R. L. Graham, "Bounds on Multiprocessing Timing Anomalies," *SIAM Journal on Applied Mathematics,* Vol. 17, pp. 416–429, 1969.

[7] J. Y. Halpern, "A Computer Scientist Looks at Game Theory," *Games and Economic Behavior,* to appear, 2002.

[8] S. Kakutani, "A Generalization of Brouwer's Fixed Point Theorem," *Duke Mathematical Journal,* Vol. 8, pp. 457–459, 1941.

[9] R. M. Karp and M. Luby, "Monte Carlo Approximation Algorithms for Enumeration and Reliability Provblems," *Proceedings of the 24th Annual IEEE Symposium on Foundations of Computer Science,* pp. 56–64, May 1983.

[10] R. M. Karp, M. Luby and N. Madras, "Monte Carlo Approximation Algorithms for Enumeration Provblems," *Journal of Algorithms,* Vol. 10, pp. 429–448, 1989.

[11] J. Kleinberg, Y. Rabani, and E. Tardos. Allocating Bandwidth for Bursty Connections. *SIAM J. Comput.*, 30(1):191–217, 2000.

[12] D. Koller and N. Megiddo, "Finding Mixed Strategies with Small Supports in Extensive Form Games," *International Journal of Game Theory,* Vol. 25, pp. 73–92, 1996.

[13] D. Koller, N. Megiddo and B. von Stengel, "Fast Algorithms for Finding Randomized Stategies in Game Trees," *Proceedings of the 26th Annual ACM Symposium on Theory of Computing,* pp. 750–759, May 1994.

[14] Y. Korilis, A. Lazar and A. Orda, "Architecting Noncooperative Networks," *IEEE Journal on Selected Areas in Communications,* Vol. 13, No. 7, pp. 1241–1251, 1995.

[15] E. Koutsoupias, M. Mavronicolas and P. Spirakis, "Approximate Equilibria and Ball Fusion," *Proceedings of the 9th International Colloquium on Structural Information and Communication Complexity*, June 2002.

[16] E. Koutsoupias and C. H. Papadimitriou, "Worst-case Equilibria," *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, G. Meinel* and *S. Tison* eds., pp. 404–413, Vol. 1563, Lecture Notes in Computer Science, Springer-Verlag, Trier, Germany, March 1999.

[17] C. E. Lemke and J. T. Howson, "Equilibrium Points of Bimatrix Games," *Journal of the Society for Industrial and Applied Mathematics,* Vol. 12, pp. 413–423, 1964.

[18] M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. *Algorithmica*, 48:91–126, 2007.

[19] R. D. McKelvey and A. McLennan, "Computation of Equilibria in Finite Games," in *Handbook of Computational Economics,* H. Amman, D. Kendrick and J. Rust eds., pp. 87–142, 1996.

[20] N. Megiddo, "A Note on the Complexity of P-Matrix LCP and Computing an Equilibrium," Research Report RJ6439, IBM Almaden Research Center, San Jose, CA95120, 1988.

[21] N. Megiddo and C. H. Papadimitriou, “On Total Functions, Existence Theorems, and Computational Complexity,” *Theoretical Computer Science,* Vol. 81, No. 2, pp. 317–324, 1991.

[22] B. Monien, Personal Communication, April 2002.

[23] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, 1995.

[24] R. B. Myerson, *Game Theory,* Harvard University Press, Cambridge, Massachusetts, 1991.

[25] J. F. Nash, “Non-cooperative Games,” *Annals of Mathematics,* Vol. 54, No. 2, pp. 286–295, 1951.

[26] J. F. Nash, S. Nasar (ed.) and H. Kuhn (ed.), *The Essential John Nash,* Princeton University Press, 2001.

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

[28] C. H. Papadimitriou, *Computational Complexity,* Addison-Wesley, 1993.

[29] C. H. Papadimitriou, ”On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence,” *Journal of Computer and System Sciences,* Vol. 48, No. 3, pp. 498–532, June 1994.

[30] C. H. Papadimitriou, ”Algorithms, Games and the Internet,” *Proceedings of the 28th International Colloquium on Automata, Languages and Programming,* F. Orejas, P. G. Spirakis and J. von Leeuwen eds., pp. 1–3, Vol. 2076, Lecture Notes in Computer Science, Springer-Verlag, July 2001.

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

[32] J. Rosenmuller, ”On a Generalization of the Lemke-Howson Algorithm to Non-cooperative N-Person Games,” *SIAM Journal of Applied Mathematics,* Vol. 21, pp. 73–79, 1971.

[33] H. E. Scarf, *The Computation of Economic Equilibria,* Yale University Press, New Haven, Connecticut, 1973.

[34] L. S. Shapley, ”A Note on the Lemke-Howson Algorithm,” *Mathematical Programming Study,* Vol. 1, pp. 175–189, 1974.

[35] B. von Stengel, ”Computing Equlibria for Two-Person Games,” in *Handbook of Game Theory,* Vol. 3, R. J. Aumann and S. Hart eds., North Holland, Amsterdam, 1998.

[36] L. G. Valiant, ”The Complexity of Computing the Permanent,” *Theoretical Computer Science,* Vol. 8, pp. 189–201, 1979.

[37] Web page `http://www.cs.berkeley.edu/ christos/games/cs294.html` for the course *CS294: Seminar on Algorithmic Aspects of Game Theory,* taught by C. H. Papadimitriou at University of California at Berkeley, Spring 2001.

[38] Web page `http://www.abeautifulmind.com` for the movie *A Beautiful Mind.*

[39] R. Wilson, ”Computing Equilibria of n-Person Games,” *SIAM Journal of Applied Mathematics,* Vol. 21, pp. 80–87, 1971.