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 her own traffic. In a *Nash equilibrium*, each user selfishly routes her traffic on those links that minimize her *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 pure Nash equilibrium, constructing a Nash equilibrium, constructing the pure Nash equilibria of minimum and maximum social cost, and computing the social cost of a given mixed Nash equilibrium. Our work provides a comprehensive collection of efficient algorithms, hardness 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.

* An extended abstract of this work appeared in the Proceedings of the 29th International Colloquium on Automata, Languages and Programming, Lecture Notes in Computer Science 2380, pp. 123–134, 2002. This work has been partially supported by the IST Program of the European Union under contract numbers IST-1999-14186 (ALCOM-FT), IST-2001-33116 (FLAGS), and IST-015964 (AEOLUS), 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. Most of this work was done while the author was a postdoctoral fellow at the Max Planck Institut für Informatik, 66123 Saarbrücken, Germany. Email: `fotakis@aegean.gr`

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

§ Department of Informatics, University of Athens, Greece. Email: `elias@di.uoa.gr`

¶ Department of Computer Science, University of Cyprus, P.O. Box 20537, Nicosia CY-1678, Cyprus. Currently visiting Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, Germany. Email: `mavronic@ucy.ac.cy`

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

1 Introduction

Nash equilibrium [35] is arguably the most important solution concept in Game Theory [36]. 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.

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 [25]. 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 her own *traffic*. For each link, a *capacity* specifies the rate at which the link processes traffic. In a Nash equilibrium, each user selfishly routes her 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 distinguish 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 are interested in algorithmic problems related to the computation of Nash equilibria for the selfish routing game we consider. More specifically, we seek to determine the computational complexity of the following algorithmic problems, assuming that all traffics and capacities are given:

- Decide whether there is a pure Nash equilibrium; if so, determine the corresponding users' strategies.
- Determine the users' strategies in a mixed Nash equilibrium.
- Determine the *best* and the *worst* Nash equilibria and their social cost.
- Given a mixed Nash equilibrium, compute its social cost.

Our study sometimes distinguishes between the cases of *uniform capacities*, where all capacities are equal, and of *arbitrary capacities*. We also distinguish between the cases of *identical traffics*, where all traffics are equal, and of *arbitrary traffics*.

1.1 Contribution

Pure Nash Equilibria. We start with pure Nash equilibria. By the linearity of the expected latency cost functions we consider, the celebrated result of Nash [35] on the existence of Nash equilibria 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 that the *lexicographically minimum* sorted vector of link latencies corresponds to a Nash equilibrium. The proof itself is *inefficient* (in the sense of Papadimitriou [38]) 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.

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 assigning each traffic in order to the link that currently minimizes its expected latency cost. The time complexity of the algorithm is $O(n(m + \log n))$.

We proceed to consider the related problems of determining either the *best* Nash equilibrium or the *worst* pure Nash equilibrium (with respect to social cost). By simple reductions from 3-PARTITION [19, Problem SP15], we show that both problems are $\mathcal{NP}$-complete in the strong sense even for links of identical capacity (Theorem 3.3 and Theorem 3.4).

Mixed Nash Equilibria. 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 exploited by a single user whose support "crosses" another user's support (Proposition 4.1). Using this property, we establish that if the capacities are uniform and there are only two users, the worst mixed Nash equilibrium (with respect to social cost) is the fully mixed Nash equilibrium (Theorem 4.2).

We continue to formulate an efficient and elegant algorithm for computing a mixed Nash equilibrium (Theorem 5.1) in $O(m \log m)$ time. More specifically, the algorithm computes the *generalized fully mixed* Nash equilibrium for the model of identical traffics. In the generalized fully mixed Nash, the links are partitioned into a set of empty links, which do not belong to the support of any user, and a set of used links, that comprise the common support of all users. This generalizes the *fully mixed* Nash equilibrium [29]. The algorithm incrementally constructs the common support of all users by throwing away links with small capacity till it converges to the common support of a fully mixed Nash equilibrium. We also establish the *uniqueness* of the generalized fully mixed Nash equilibrium (Theorem 5.4).

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 33.041 times the social cost of the generalized fully mixed Nash equilibrium (Theorem 6.1). This result is shown by: (i) establishing some interesting properties of the support and the expected latency of the links in an arbitrary Nash equilibrium, and (ii) 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 some kind of completeness property of the fully mixed Nash equilibrium. It appears that it sometimes suffices 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.

Computing the Social Cost of a Mixed Nash Equilibrium. We then study the computational complexity of computing the social cost of a given mixed Nash equilibrium. We have obtained both negative and positive results here. As for the bad news, we show that the problem is $\#\mathcal{P}$-complete [39] even for instances with 3 links of identical capacity (Theorem 7.1). The proof employs 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. [23, Theorem 2.1] for the $\#\mathcal{P}$-completeness of the latter problem). We show that this probability can be recovered by two calls to a (hypothetical) oracle returning the social cost of any given mixed Nash equilibrium.

On the positive side, we get around the established hardness of computing exactly the social cost by presenting a Fully Polynomial-Time Randomized Approximation Scheme (FPRAS) 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.

Discussion and Comparison to Previous Work. The selfish routing game considered in this paper was introduced by Koutsoupias and Papadimitriou [25] as a vehicle for the study of the price of selfishness for routing over non-cooperative networks, like the Internet. This game was subsequently studied in the work of Mavronicolas and Spirakis [29], 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. Some immediately later studies of the selfish routing game from the same point of view, that of performance, include the works by Koutsoupias *et al.* [24], and by Czumaj and Vöcking [3].

Our work mostly 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. Our polynomial-time algorithms for the computation of pure and mixed Nash equilibria (Theorem 3.2 and Theorem 5.1, respectively) are the first known polynomial-time algorithms for the problem. These algorithms may offer ideas for settling other tractable instances of either the same game or other games of a similar flavor (e.g., selfish routing over a larger network). To this end, we also feel that they will be handy results offering insights into structure and the properties of Nash equilibria, such as the existence and uniqueness of the generalized fully mixed Nash equilibrium (Theorem 5.1 and Theorem 5.4), and the structure of the mixed Nash equilibrium of maximum social cost (Theorem 4.2 and Theorem 6.1).

Issues of computational complexity for the computation of Nash equilibria in general games have been raised by Megiddo [30], Megiddo and Papadimitriou [31], and Papadimitriou [38]. The $\mathcal{NP}$-hardness of computing a Nash equilibrium with certain properties in a general bimatrix game has been established by Gilboa and Zemel [20] and by Conitzer and Sandholm [2]. In this context, the hardness results we have obtained (Theorem 3.3, Theorem 3.4, and Theorem 7.1) indicate that optimization and counting problems in Algorithmic Game Theory may be hard even when restricted to simple games such as the selfish routing game considered in our work.

2 Framework

Most of our definitions are patterned after those in [25, Sections 1 & 2] and [29, Section 2].

2.1 Mathematical Preliminaries and 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.

For any two $m \times 1$ vectors $\mathbf{x}$ and $\mathbf{y}$, we 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.

2.2 Model

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 users wishes to route a particular amount of traffic along a (non-fixed) link from the source to the destination. Assume throughout that $m \geq 2$ and $n \geq 2$. Throughout, we will be using subscripts for users and superscripts for links.

Denote $c^\ell > 0$ the capacity of each link $\ell \in [m]$, representing the rate at which the link processes traffic. We assume, without loss of generality, that the links are indexed in non-increasing capacity order, i.e, $c^1 \geq c^2 \geq \cdots \geq c^m$. For each $k \in [m]$, we let $C(k) = \sum_{\ell=1}^{k} c^\ell$ denote the total capacity of the fastest k links. For any $k \in [m]$, we say that the set $[k]$ consisting of the fastest k links is a *fast link set*. We distinguish between the case of *uniform capacities*, where all link capacities are equal, and the case of *arbitrary capacities*, where link capacities may vary arbitrarily. In the case of uniform capacities, we assume, without loss of generality, that all link capacities are equal to 1.

Denote as w_i the traffic of each user $i \in [n]$. Define the $n \times 1$ traffic vector $\mathbf{w}$ in the natural way. We distinguish between the case of *uniform traffics*, where all user traffics are equal, and the case of *arbitrary traffics*, where user traffics may vary arbitrarily. In the case of uniform traffics, we assume, without loss of generality, that all user traffics are equal to 1.

2.2.1 Strategies and Profiles

A *pure strategy* for user $i \in [n]$ is a specific link. A *mixed strategy* for user $i \in [n]$ is a probability distribution over pure strategies; thus, a mixed strategy is a probability distribution over the set of links. A *pure strategies profile* is an n-tuple $\langle \ell_1, \ell_2, \ldots, \ell_n \rangle \in [m]^n$; a *mixed strategies profile* is an $n \times m$ probability matrix $\mathbf{P}$ of nm probabilities p_i^j, $i \in [n]$ and $j \in [m]$, where p_i^j is the probability that user i assigns her traffic to link j. 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.

Given mixed strategies profile $\mathbf{P}$, the *support* of each user $i \in [n]$ in $\mathbf{P}$, denoted $support(i)$, is the set of links to which i assigns positive probability. Formally, $support(i) = \{\ell \in [m] : p_i^\ell > 0\}$. The support of $\mathbf{P}$ is the union of the supports of all users. For each link $\ell \in [m]$, the *view* of link ℓ in $\mathbf{P}$, denoted $view(\ell)$, is the set of users $i \in [n]$ that assign positive probability to ℓ. Formally, $view(\ell) = \{i \in [n] : p_i^\ell > 0\}$. A link ℓ is *solo* if $|view(\ell)| = 1$.

A mixed strategies profile $\mathbf{P}$ is *fully mixed* [29] if for all users $i \in [n]$ and links $j \in [m]$, $p_i^j > 0$. A mixed strategies profile $\mathbf{P}$ is *generalized fully mixed* if there exists a subset $L \subseteq [m]$ such that for each pair of a user $i \in [n]$ and a link $j \in [m]$, $p_i^j > 0$ if $j \in L$, and $p_i^j = 0$ if $j \notin L$. Thus, the fully mixed strategies profile is the special case of generalized fully mixed strategies profiles where $L = [m]$.

2.2.2 Cost Measures

The *latency* for traffic w through link ℓ equals w/c^ℓ. For a pure strategies profile $\langle \ell_1, \ell_2, \ldots, \ell_n \rangle$, the *latency cost for user* i, denoted λ_i, is $\sum_{k:\ell_k=\ell_i} w_k/c^{\ell_i}$; that is, the latency cost for user i is the latency of the link it chooses. For a mixed strategies profile $\mathbf{P}$, denote Λ^ℓ the *expected latency* of link ℓ; clearly, $\Lambda^\ell = \sum_{i=1}^n p_i^\ell w_i / c^\ell$.

For a mixed strategies profile $\mathbf{P}$, the *expected latency cost* for user i on link ℓ, denoted λ_i^ℓ, is the expectation, over all random choices of the remaining users, of the latency cost for user i had its traffic been assigned to link ℓ; thus,

$$\lambda_i^\ell = \frac{w_i + \sum_{k=1,k\neq i} p_k^\ell w_k}{c^\ell} = (1 - p_i^\ell)\frac{w_i}{c^\ell} + \Lambda^\ell$$

For each user $i \in [n]$, the *minimum expected latency cost*, denoted λ_i, is the minimum, over all links ℓ, of the expected latency cost for user i on link ℓ; thus, $\lambda_i = \min_{\ell\in[m]} \lambda_i^\ell$.

Associated with a traffic vector $\mathbf{w}$ and a mixed strategies profile $\mathbf{P}$ is the *social cost* [25, Section 2], denoted $\mathrm{SC}(\mathbf{w}, \mathbf{P})$, which is the expectation, over all random choices of the users, of the maximum (over all links) latency of traffic through a link; thus,

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

Note that $\mathrm{SC}(\mathbf{w}, \mathbf{P})$ reduces to the maximum latency through a link in the case of pure strategies.

On the other hand, the *social optimum* [25, Section 2] associated with a traffic vector $\mathbf{w}$, denoted $\mathrm{OPT}(\mathbf{w})$, is the least possible maximum (over all links) latency of traffic through a link; thus,

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

2.2.3 Nash Equilibrium

The strategies profile $\mathbf{P}$ is a *Nash equilibrium* if for all users $i \in [n]$ and all links $\ell \in [m]$, $\lambda_i^\ell = \lambda_i$ if $p_i^\ell > 0$, and $\lambda_i^\ell \geq \lambda_i$ if $p_i^\ell = 0$. Thus, each user assigns her traffic with positive probability only to links for which her 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 decrease her expected latency cost.

2.2.4 Properties of Nash Equilibria

We continue to state some known properties of Nash equilibria that will be used in our later proofs. The following result due to Mavronicolas and Spirakis [29, Lemma 6.1] provides a simple characterization of existence and uniqueness of fully mixed Nash equilibria in the case of identical traffics and arbitrary capacities.

Lemma 2.1 (Mavronicolas and Spirakis [29]) *Consider the fully mixed case for n users of identical traffic and m parallel links with arbitrary capacities. 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 [29, 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, in the case of identical traffics, all users incur the same minimum expected latency cost in a fully mixed equilibrium.

Lemma 2.2 (Mavronicolas and Spirakis [29]) *Consider a fully mixed Nash equilibrium for a selfish routing game on n users of identical traffic and m parallel links with arbitrary capacities. Then the minimum expected latency cost of any user i is $\lambda_i = (m + n - 1)/C(m)$.*

Finally, Mavronicolas and Spirakis [29, Lemma 5.1] show that in the case of uniform capacities, all links are equiprobable in the fully mixed Nash equilibrium.

Lemma 2.3 (Mavronicolas and Spirakis [29]) *Consider the case of uniform capacities. Then, there exists a unique fully mixed Nash equilibrium with associated Nash probabilities $p_i^\ell = 1/m$ for each user $i \in [n]$ and link $\ell \in [m]$.*

2.3 Algorithmic Problems

In this section, we formally define some algorithmic problems related to the social cost of Nash equilibria. The definitions are given in the style of Garey and Johnson [19].

BEST NASH EQUILIBRIUM
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 rational number $B > 0$.
QUESTION: Is there a Nash equilibrium $\mathbf{P}$ with $\mathrm{SC}(\mathbf{w}, \mathbf{P}) \leq B$?

WORST PURE NASH EQUILIBRIUM
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 rational number $B > 0$.
QUESTION: Is there a pure Nash equilibrium $\mathbf{P}$ with $\mathrm{SC}(\mathbf{w}, \mathbf{P}) \geq B$?

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 mixed strategies profile $\mathbf{P}$ that is a Nash equilibrium.
OUTPUT: The social cost of the Nash equilibrium $\mathbf{P}$.

To establish our $\mathcal{NP}$-hardness results, we will use the problem of 3-PARTITION, which is $\mathcal{NP}$-complete in the strong sense (see e.g. [19, Problem SP15 and Theorem 4.4]):

3-PARTITION
INSTANCE: A positive integer $m \geq 2$; a positive integer B; 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

We start with a preliminary result remarked by Kurt Mehlhorn.

Theorem 3.1 *There exists at least one pure Nash equilibrium.*

Proof. Every pure strategies profile $\mathbf{P}$ induces a sorted latency vector $\mathbf{\Lambda} = \langle \Lambda^1, \Lambda^2, \ldots, \Lambda^m \rangle$, such that $\Lambda^1 \geq \Lambda^2 \geq \cdots \geq \Lambda^m$, in the natural way (rearrangement of links may be necessary to guarantee that the latency vector is sorted). We argue that the pure strategies profile $\mathbf{P}_0$ corresponding to the lexicographically minimum sorted latency vector $\mathbf{\Lambda}_0 = \langle \Lambda_0^1, \ldots, \Lambda_0^m \rangle$ is (pure) Nash equilibrium.

To reach a contradiction, we assume that $\mathbf{P}_0$ is not a Nash equilibrium. Thus there exists a user $i \in [n]$ assigning her traffic to link $j \in [m]$, and a link $k \in [m]$ such that $\Lambda_0^k + \frac{w_i}{c^k} < \Lambda_0^j$. Therefore if user i moves her traffic from link j to link k, the latency of link k after the move remains smaller than the latency of link j in $\mathbf{P}_0$. The latency of link j after the move is also smaller than j's latency in $\mathbf{P}_0$. Hence the sorted latency vector after user i moving her traffic from link j to link k is lexicographically less that $\mathbf{\Lambda}_0$, a contradiction. □

We remark that the proof of Theorem 3.1 establishes that the lexicographically minimum traffic vector corresponds to a pure Nash equilibrium. Since there are exponentially many pure strategies profiles, Theorem 3.1 only provides an *inefficient* proof of existence of pure Nash equilibria (see also Papadimitriou [38]). An efficient algorithmic proof of existence of pure Nash equilibria is the subject of the following section.

3.2 Polynomial Computation

We show that a pure Nash equilibrium can be computed by a simply greedy algorithm.

Theorem 3.2 *A pure Nash equilibrium can be computed in $O(n(m + \log n))$ time.*

Proof. The algorithm Greedy Pure Nash Equilibrium (GPNE) in Figure 3.2 computes a pure Nash equilibrium in time $O(n(m + \log n))$. GPNE works in a greedy fashion; it considers each traffic in non-increasing order and assigns it to the link of minimum latency had the traffic been assigned to that link.

For each $k \in [n]$, let $\mathbf{P}(k) = \langle \ell_1, \ldots, \ell_k \rangle$ denote the restriction of the pure strategies profile computed by GPNE to the first k users with the largest traffics, and let $\Lambda^j(k) = \sum_{i \in [k]: \ell_i = j} w_i / c^j$ denote the latency of link $j \in [m]$ in $\mathbf{P}(k)$ (ignoring all users in $[n] \setminus [k]$). We argue inductively that for all $k \in [n]$, $\mathbf{P}(k)$ is a pure Nash equilibrium for the first k users, if the remaining users in $[n] \setminus [k]$ are ignored.

The claim holds for $k = 1$, since the first user assigns her traffic to the link ℓ_1 minimizing her latency cost. We inductively assume that for $k \geq 1$, $\mathbf{P}(k)$ is a pure Nash equilibrium for the users in $[k]$ and show that $\mathbf{P}(k+1)$ is a pure Nash equilibrium for the users in $[k+1]$.

Algorithm GPNE:

- For each link $j \in [m]$, $\Lambda^j := 0$;
- Sort the users in non-increasing order of their traffics so that $w_1 \geq w_2 \geq \ldots \geq w_n$;
- For each user $k := 1$ to n, do
  - $\ell_k := \arg\min_{j \in [m]} \{\Lambda^j + w_k / c^j\}$;
  - $\Lambda^{\ell_k} := \Lambda^{\ell_k} + w_k / c^{\ell_k}$;
- return $\langle \ell_1, \ldots, \ell_n \rangle$;

Figure 1: The algorithm Greedy Pure Nash Equilibrium (GPNE) in pseudocode.

Let ℓ_{k+1} be the link to which user $k+1$ assigns her traffic. For every link $q \neq \ell_{k+1}$, every user $i \in [k+1]$ assigning her traffic to link q, and every link $j \in [m]$,

$$\Lambda^q(k+1) = \Lambda^q(k) \leq \Lambda^j(k) + \frac{w_i}{c^j} \leq \Lambda^j(k+1) + \frac{w_i}{c^j}$$

The equality holds because the latency of link q does not change when user $k+1$ is considered, the first inequality follows from the inductive hypothesis that $\mathbf{P}(k)$ is a pure Nash equilibrium, and the last inequality holds because the latency of any link does not decrease when user $k+1$ is considered.

In addition, for every user $i \in [k+1]$ assigning her traffic to link ℓ_{k+1} and every link $j \neq \ell_{k+1}$,

$$\begin{aligned} \Lambda^{\ell_{k+1}}(k+1) &= \Lambda^{\ell_{k+1}}(k) + \frac{w_{k+1}}{c^{\ell_{k+1}}} \leq \Lambda^j(k) + \frac{w_{k+1}}{c^j} \\ &= \Lambda^j(k+1) + \frac{w_{k+1}}{c^j} \leq \Lambda^q(k+1) + \frac{w_i}{c^j} \end{aligned}$$

The equalities hold because the latency of link ℓ_{k+1} increases by $w_k / c^{\ell_{k+1}}$ and the latency of link $j \neq \ell_{k+1}$ does not change when user $k+1$ is considered. The first inequality follows by the choice of ℓ_{k+1} as a link of minimum $\Lambda^j(k) + w_{k+1}/c^j$ over all links $j \in [m]$. The last inequality holds because for all users $i \in [k+1]$, $w_i \geq w_{k+1}$.

We observe that none of the first $k+1$ users can decrease her latency cost by switching to another link, which concludes the induction. Hence $\langle \ell_1, \ldots, \ell_n \rangle$ is a pure Nash equilibrium.

As for the time complexity of GPNE, sorting the traffics in non-increasing order takes $O(n \log n)$ time, and initializing $\mathbf{P}$ and choosing the link ℓ_k for each user $k \in [n]$ requires $O(nm)$ time. Hence the time complexity of GPNE is $O(n(m + \log n))$. □

Burkhard Monien [32] remarked that GPNE can be regarded as a variant of Graham's Longest Processing Time [21] algorithm for assigning tasks to identical machines. Nevertheless, since in our case the links may have arbitrary capacities, GPNE chooses the link that minimizes the completion time of the task under consideration (i.e. the load of a machine prior to the assignment of the task plus the overhead of this task) instead of choosing the link that will first become idle.

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 social cost of the best Nash equilibrium and the worst pure Nash equilibrium.

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

Proof. By the definition of the social cost, a selfish routing game admits a Nash equilibrium of social cost at most B if and only if 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.

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 if and only if (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. $\square$

Theorem 3.4 WORST PURE NASH EQUILIBRIUM *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-PARTITIONN, 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 $2B$ if and only if (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, $J_1, \dots J_m$ comprises a YES-certificate for the corresponding instance of 3-PARTITION. □

4 Worst Mixed Nash Equilibria

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 are not the same and have non-empty intersection, i.e. their supports cross each other.

Proposition 4.1 *Let* $\mathbf{P}$ *be any Nash equilibrium in the case of uniform capacities. Then* $\mathbf{P}$ *induces no solo link included in the support of a user that crosses another user.*

Proof. Let us assume that $\mathbf{P}$ induces a solo link ℓ that is included in the support of user i that crosses another user. Thus, there exists another link $j \in support(i)$ and another user $k \in view(j)$, so that $p_k^j > 0$. By the definition of users' expected latency cost,

$$\lambda_i^j \geq w_i + p_k^j w_k > w_i = \lambda_i^j \ ,$$

which contradicts the assumption that $\mathbf{P}$ is a Nash equilibrium. □

We now use Proposition 4.1 to provide a syntactic characterization of the worst mixed Nash equilibrium in the case of uniform capacities.

Theorem 4.2 *Consider the case 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 $w_1 \geq w_2$ and consider any Nash equilibrium $\mathbf{P}$. Thus, the social cost of $\mathbf{P}$ is w_1, if w_1 and w_2 are assigned to different links, and $w_1 + w_2$, if w_1 and w_2 are assigned to the same link ℓ. Therefore,

$$\mathrm{SC}(\mathbf{w}, \mathbf{P}) = w_1 + w_2 \sum_{\ell=1}^{m} p_1^\ell p_2^\ell$$

We will show that $\mathrm{SC}(\mathbf{w}, \mathbf{P})$ 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.

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

   (a) Assume first that $support(1) \cap support(2) = \emptyset$. Then, it is impossible to have a link ℓ with both $p_1^\ell > 0$ and $p_2^\ell > 0$, and $\mathrm{SC}(\mathbf{w}, \mathbf{P}) = w_1$.

   (b) Assume now that $support(1) \cap support(2) \neq \emptyset$. We will show that in this case $\mathbf{P}$ is the fully mixed Nash equilibrium. Proposition 4.1 implies that $support(1) = support(2)$, since otherwise, there would a solo link and the two supports would cross. We will show that, in fact, $support(1) = support(2) = [m]$.

   To reach a contradiction, assume that there is some link $j \in [m] \setminus support(1)$. Then, the expected latency cost of user 1 on link j is equal to w_1, which is less than $w_1 + w_2 p_2^\ell$, its expected latency cost on any link $\ell \in support(2)$. This contradicts the assumption that $\mathbf{P}$ is a Nash equilibrium. It follows that $support(1) = support(2) = [m]$, so that $\mathbf{P}$ is a fully mixed Nash equilibrium. By Lemma 2.3, $p_1^\ell = p_2^\ell = 1/m$, for all $\ell \in [m]$, so that

$$\mathrm{SC}(\mathbf{w}, \mathbf{P}) = w_1 + \frac{w_2}{m}$$

The claim follows. □

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

Theorem 5.1 *In the case of identical traffics, a generalized fully mixed Nash equilibrium can be computed in $O(m \log m)$ time.*

Proof. We present the polynomial-time algorithm Generalized Fully Mixed (GFM) that computes the support of a generalized fully mixed equilibrium. We recall that 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 2.1 gives a simple characterization of the games that admit a fully mixed Nash equilibrium.

Proposition 5.2 (Monien [32]) *Consider a selfish routing game Γ on n users of identical traffic and m parallel links with arbitrary capacities. Then Γ admits a fully mixed Nash equilibrium, which must be unique, if and only if $c^m > C(m)/(m + n - 1)$.*

Algorithm GFM:

- Sort the links in non-increasing order of their capacities so that $c^1 \geq c^2 \geq \cdots \geq c^m$;
- $m' := m;\ \ C(m') := \sum_{\ell=1}^{m'} c^\ell$;
- while $m' \geq 0$ do
  - if $c^{m'} > C(m')/(m' + n - 1)$ then
    - return the fully mixed strategies profile on $[n]$ and $[m']$;
  - else
    - $C(m' - 1) := C(m') - c^{m'};\ \ m' := m' - 1$;

Figure 2: The algorithm Generalized Fully Mixed (GFM) in pseudocode.

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 $C(m)$. The claim now follows from Lemma 2.1. □

The algorithm 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. GFM outputs the fully mixed Nash equilibrium for the restriction of the routing game to $[m']$ and terminates. The fast link set $[m']$ returned by 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 $[m']$ returned by 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 on the set of all users $[n]$ and the fast link set $[m']$ returned by* GFM *remains a Nash equilibrium when the game is extended to the entire set of links.*

Proof. Let $\mathbf{P}$ be the fully mixed Nash equilibrium for the selfish routing game on $[n]$ and $[m']$ returned by GFM. 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 $[n]$ and $[m]$.

The claim is trivial if $m' = m$. Otherwise, Lemma 2.2 implies that the minimum expected latency cost of any user i is $\lambda_i = (m' + n - 1)/C(m')$. Since GFM does not include link $m' + 1$ in the support of $\mathbf{P}$, $c^{m'+1} \leq C(m' + 1)/(m' + n)$. Therefore,

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

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 $[n]$ and $[m]$. □

The time complexity of 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 2.1, since all users have the same strategy. Therefore, the time complexity of GFM is $O(m \log m)$. $\square$

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 parallel links with arbitrary capacities. The equilibrium returned by* GFM *is the unique generalized fully mixed Nash equilibrium of Γ.*

Proof. Let $\mathbf{P}$ be the generalized fully mixed Nash equilibrium returned by 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 2.2, 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 assumption 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 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 2.2, the minimum expected latency cost of any user i in $\mathbf{P}'$ is $\lambda'_i = (q + n - 1)/C(q)$. On the other hand,

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

Combining $\sum_{j=q+1}^{m'} c^j \geq (m' - q)c^{m'}$ with (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. $\square$

We conclude this section with a 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 parallel links with arbitrary capacities. 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 $c^{m'} > C(m')/(m' + n - 1)$. Then $[m']$ 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 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. $\square$

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 parallel links with arbitrary capacities, the social cost of the worst Nash equilibrium is at most* 33.041 *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 a simplifying assumption 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

We recall the basis of the natural logarithms $e = 2.718\ldots$; for each x, denote $\exp(x) = e^x$. We denote as $\overline{\mathbf{P}}$ the generalized fully mixed Nash equilibrium, and as $\mathbf{P}$ 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 $[m']$ be the support of the generalized fully mixed Nash equilibrium $\overline{\mathbf{P}}$. 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 2.1 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 $i \in [n]$, 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 $\mathbf{P}$ be any Nash equilibrium. For any $k \in [m]$, the minimum expected latency cost of any user i in $\mathbf{P}$ is $\lambda_i \leq (k + n - 1)/C(k)$.*

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

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

a contradiction. $\square$

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 actual latency of any solo link.

Corollary 6.5 *Let* $\mathbf{P}$ *be any Nash equilibrium. For any link* $\ell \in [m]$, $\Lambda^\ell \leq (m' + n - 1)/C(m')$, *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 $\Lambda^\ell \leq \lambda_i \leq (m' + n - 1)/C(m')$. $\square$

Corollary 6.6 *Let* $\mathbf{P}$ *be any Nash equilibrium. For any solo link* $\ell \in [m]$, *the actual 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 actual 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 $1/c^\ell = \lambda_i \leq (m' + n - 1)/C(m')$. $\square$

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 actual latency of a solo link is at most $2\,\mathrm{SC}(\mathbf{1}, \overline{\mathbf{P}})$ (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. $\square$

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* $\mathbf{P}$,

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

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,

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

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

Proposition 5.5 and Proposition 6.7 imply that we can ignore any link not in the support of the generalized fully mixed Nash equilibrium. Hence, we can assume, without loss of generality, that the support of the generalized fully mixed Nash equilibrium coincides with the entire set of links.

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 (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 (Proposition 5.5), but it may be included in the support of some other Nash equilibrium as a solo link (Proposition 6.7). However, the actual latency of a solo link is at most $\frac{m'+n-1}{C(m')}$ (Corollary 6.6). Since the social cost of the generalized fully mixed Nash equilibrium is at least $\frac{m'+n-1}{2C(m')}$ (Proposition 6.3), solo links cannot increase the maximum link latency above $2\,\mathrm{SC}(\mathbf{1}, \overline{\mathbf{P}})$. Therefore we can ignore any link not in 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.

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 technical claim 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 $x \geq 8/c^m$, $h_j(x) \leq f_j(\frac{7x}{16e^2})^{8e^2/7}$.*

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

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

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 2.1, 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. [4]), we obtain that

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

For the first inequality, we use that the random variables $\overline{X}^1, \ldots, \overline{X}^m$ are negatively associated (see e.g. [4, 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. More precisely, for every $j \in [m]$, the j-th term of $g(x)$ jumps from $f_j(n/c^j) = (\frac{\overline{\Lambda}^j c^j}{2en})^{2n}$ to 0 at n/c^j. Since $\sum_{\ell \in [m]} \overline{\Lambda}^\ell c^\ell = n$, each jump of $g(x)$ is at most[1] $(\frac{1}{2e})^{2n}$. Using $n \geq 2$, we obtain that each jump of $g(x)$ is less than 0.0012. Therefore, for any $\alpha \in (1, e)$, there is at least one point $x \in (x^*, n)$ such that $g(x) \in (\ln \alpha - 0.0012, \ln \alpha]$. In the following, we let μ_α denote the smallest such value:

$$\mu_\alpha = \arg\min\{x \in (x^*, n) : g(x) \in (\ln \alpha - 0.0012, \ln \alpha]\} \tag{8}$$

By the definition of μ_α, $g(\mu_\alpha) \geq \ln \alpha - 0.0012$. 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 $(1 - \frac{e^{0.0012}}{\alpha})\mu_\alpha^*$ on the social cost of the generalized fully mixed Nash equilibrium.

Lemma 6.12 *For any* $\alpha \in (1, e)$, $\mathrm{SC}(\mathbf{1}, \overline{\mathbf{P}}) \geq (1 - \frac{e^{0.0012}}{\alpha})\mu_\alpha^*$, *where* $\mu_\alpha^* = \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}, \mu_\alpha\}$.

Proof. In Proposition 6.3, we show that $\mathrm{SC}(\mathbf{1}, \overline{\mathbf{P}}) \geq \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}\}$. 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

$$\mathbb{P}\mathrm{r}[\overline{L}_{\max} \geq \mu_\alpha] \geq 1 - \exp(-g(\mu_\alpha)) \geq 1 - \frac{e^{0.0012}}{\alpha}$$

because $g(\mu_\alpha) \geq \ln \alpha - 0.0012$. Therefore, the social cost of the generalized fully mixed Nash equilibrium is at least $(1 - \frac{e^{0.0012}}{\alpha}) \max\{\overline{\Lambda}^1, \frac{m+n-1}{2C(m)}, \mu_\alpha\}$. □

[1] The upper bound on the jumps of $g(x)$ holds even if there are $k \geq 2$ links with the same capacity, since in the generalized fully mixed Nash equilibrium, all of them receive the same expected traffic, which does not exceed n/k.

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[2] 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 $x > (m+n-1)/C(m)$,

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

$\square$

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

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

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

[2] We use the following Chernoff bound (see e.g. [34, Theorem 4.1]) 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 $\delta > 0$,

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

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

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

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):

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

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 $y \geq 1$,

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

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 using that $\alpha \in (1, e)$. $\square$

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 $\alpha \in (1, e)$,

$$\mathrm{SC}(\mathbf{1}, \mathbf{P}) \leq \left(\frac{16e^2}{7} - \frac{2(\ln \alpha)^{8e^2/7}}{\ln \ln \alpha} \right) \frac{\alpha}{\alpha - e^{0.0012}} \mathrm{SC}(\mathbf{1}, \overline{\mathbf{P}})$$

Using $\alpha = 2.175$, we conclude that $\mathrm{SC}(\mathbf{1}, \mathbf{P}) \leq 33.041 \, \mathrm{SC}(\mathbf{1}, \overline{\mathbf{P}})$. □

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. [25, 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 $J = \{w_1, \dots, w_n\}$ 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. [37]). Therefore, given n independent random variables $Y_1(w_1, 1/2), \dots, Y_n(w_n, 1/2)$, 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 [23, 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 $Y_1(w_1, 1/2), \dots, Y_n(w_n, 1/2)$ 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

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

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 [34, 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

$$\mathrm{SC}_2 = C + 1 + 2 \sum_{B=C+2}^{\infty} \mathbb{P}\mathrm{r}[Y \geq B] \tag{11}$$

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

$$\mathrm{SC}_2 - \mathrm{SC}_1 = 1 - 2\,\mathbb{P}\mathrm{r}[Y \geq C+1] \tag{12}$$

Since w_i's and C are integers, $\mathbb{P}\mathrm{r}[Y \leq C] = 1 - \mathbb{P}\mathrm{r}[Y \geq C+1]$. Therefore, (12) implies that $\mathbb{P}\mathrm{r}[Y \leq C] = (\mathrm{SC}_2 - \mathrm{SC}_1 + 1)/2$. □

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 Fully Polynomial-Time Approximation Scheme (FPRAS) that approximates the social cost of any given mixed Nash equilibrium. We recall that an algorithm A is a *Fully Polynomial-Time Randomized Approximation Scheme* (FPRAS, [22]) for a counting problem Π, if for every instance x and for any error parameter $\varepsilon > 0$, $\mathbb{P}\mathrm{r}[|\mathsf{A}(x) - \Pi(x)| \leq \varepsilon\Pi(x)] \geq 3/4$, and the running time of A is polynomial in $|x|$ and ε^{-1}.

Theorem 7.2 *Consider the case of uniform capacities. Then there is 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 probabilities in $\mathbf{P}$. For each experiment $i \in [N]$, let $L^i_{\max}$ be the measured maximum link latency. Return 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 corresponding to the experiments' outcome. The expectation of $L_{\max}$ is equal to the social cost of $\mathbf{P}$ and the variance of $L_{\max}$ 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. [34, Theorem 3.3]), we obtain that for any $\varepsilon > 0$,

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

where the last inequality follows from $\mathrm{Var}[L_{\max}] \leq n^2 w^2_{\max}/N$ and $\mathrm{SC}(\mathbf{w},\mathbf{P}) \geq w_{\max}$. 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 $\mathrm{SC}(\mathbf{w},\mathbf{P})$ is at least $3/4$. □

8 Subsequent Work

The conference publication of this work in 2002 partially motivated the so-called *Fully Mixed Nash Equilibrium Conjecture* and left open numerous interesting questions concerning the existence and the computational complexity of pure Nash equilibria. A substantial body of research on these topics has followed the original publication of this work.

The result on the existence of pure Nash equilibria (Theorem 3.1) has been generalized in several directions and put into the more general perspective of potential functions (see e.g. [33]). Even-Dar *et al.* [6, Theorem 3.5] proved that the proof technique of Theorem 3.1 implies a generalized ordinal potential function for the more general case of unrelated parallel links. Recently Fotakis *et al.* [12, Theorem 3.1 and Remark 3.2] established the same result in the considerably more general setting of games among dynamic coalitions of users. In [10], Theorem 3.1 was generalized with respect to the network topology. More specifically, Fotakis *et al.* [10, Theorem 1] proved that a generalization of the selfish routing game considered in this work admits a weighted potential function, and thus a pure Nash equilibrium, for any network even if the users have different sources and destinations.

The computational complexity of pure Nash equilibria has also attracted considerable interest. Feldmann *et al.* [9, Theorem 1] introduced the notion of *Nashification* algorithms, which turn a given pure assignment into a pure Nash equilibrium of non-increased social cost, and presented the first Nashification algorithm with time complexity $O(m^2n)$. Subsequently, Gairing *et al.* [13] presented a polynomial-time Nashification algorithm for restricted parallel links of identical capacity, where each user can route her traffic only on a specific subset of allowed links. Fotakis *et al.* [11] considered series-parallel networks, a class of networks significantly more complex than parallel links, and proved that a pure Nash equilibrium can be computed by a simple greedy algorithm similar to GPNE if the users' best responses are symmetric. On the negative side, [8, 1] proved that the problem of computing a pure Nash equilibrium is $\mathcal{PLS}$-complete if the game involves a general network topology and users of identical traffic but with different sources and destinations.

The *Fully Mixed Nash Equilibrium Conjecture* asserts that the fully mixed Nash equilibrium, when it exists, maximizes the social cost. This natural conjecture and has been partially motivated by Theorem 4.2 and Theorem 6.1 and was explicitly formulated by Gairing *et al.* [16].

The Fully Mixed Nash Equilibrium Conjecture first studied in a systematic way by Lücking *et al.* [27]. The conjecture could be proved for several special cases of the problem (see e.g. [16, 17, 27]). For the special case of arbitrary users and identical links, Gairing *et al.* [17, Theorem 5.1] proved that if the number of links is sufficiently large and the number of users is equal the number of links, the social cost of any Nash equilibrium is at most twice the ratio between maximum and average traffic. The Fully Mixed Nash equilibrium conjecture was recently disproved by Fischer and Vöcking [7, Theorem 4] for the case of arbitrary users and uniform capacities. It is an interesting open problem whether the conjecture holds for the case of identical users and arbitrary capacities.

Conjectures motivated by and similar to the Fully Mixed Nash Equilibrium Conjecture were recently formulated and studied in an intensive way for several variants of the selfish routing game considered in this paper [5, 14, 15, 26, 28]. For an advocate of conjectures related to the fully mixed Nash equilibrium, we refer the reader to the recent survey [18].

Acknowledgments. We thank Kurt Mehlhorn for contributing Theorem 3.1. We thank Burkhard Monien for some extremely helpful discussions on our work and for contributing an improvement to the time complexity of the algorithm GFM (from $O(m^2)$ to $O(m \log m)$). We thank Peter Sanders and Berthold Vöcking for a fruitful discussion right after a seminar talk on Nash equilibria for selfish routing given by Paul Spirakis at the Max Planck Institut für Informatik. That discussion contributed some initial ideas for the design of the algorithm GPNE. Dimitris thanks Peter Sanders for helpful discussions about "balls and bins" experiments that contributed the initial idea of Theorem 6.1.

References

[1] H. Ackermann, H. Röglin, and B. Vöcking. On the Impact of Combinatorial Structure on Congestion Games. *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pp. 613-622, October 2006.

[2] V. Conitzer and T. Sandholm. Complexity Results about Nash Equilibria. *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, Georg Gottlob, Toby Walsh (Eds.), pp. 765–771, Morgan Kaufmann, August 2003.

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

[4] D. Dubhashi and D. Ranjan. Balls and Bins: A Study in Negative Dependence. *Random Structures and Algorithms*, Vol. 13, No. 2, pp. 99–124, 1998.

[5] R. Elsässer, M. Gairing, T. Lucking, M. Mavronicolas and B. Monien. A Simple Graph-Theoretic Model for Selfish Restricted Scheduling. *Proceedings of the 1st International Workshop on Internet and Network Economics*, X. Deng and Y. Ye (Eds.), Lecture Notes in Computer Science 3828, pp. 195–209, Springer-Verlag, December 2005.

[6] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence Time to Nash Equilibria. *Proceedings of the 30th International Colloquium on Automata, Languages and Programming*, J.C.M. Baeten, J.K. Lenstra, J. Parrow, G.J. Woeginger (Eds.), Volume 2719, Lecture Notes in Computer Science 2719, pp. 502–513, Springer-Verlag, June 2003. Full version to appear in *ACM Transactions on Algorithms*.

[7] S. Fischer and B. Vöcking. On the Structure and Complexity of Worst-Case Equilibria. *Theoretical Computer Science*, Vol. 378, pp. 165–174, 2007.

[8] A. Fabrikant, C. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pp. 604-612, June 2004.

[9] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Nashification and the Coordination ratio for a Selfish Routing Game. *Proceedings of the 30th International Colloquium on Automata, Languages and Programming*, J.C.M. Baeten, J.K. Lenstra, J. Parrow, G.J. Woeginger (Eds.), Lecture Notes in Computer Science 2719, pp. 514–526, Springer-Verlag, June 2003.

[10] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *Theoretical Computer Science*, Vol. 348, pp. 226–239, 2005.

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

[12] D. Fotakis, S. Kontogiannis, and P. Spirakis. Atomic Congestion Games among Coalitions. *Proceedings of the 33th International Colloquium on Automata, Languages and Programming*, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener (Eds.), Lecture Notes in Computer Science 4051, pp. 572–583, Springer-Verlag, July 2006. Full version to appear in *ACM Transactions on Algorithms*.

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

[14] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The Price of Anarchy for Polynomial Social Cost. *Theoretical Computer Science*, Vol. 369, Nos. 1-3, pp. 116–135, December 2006.

[15] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, J. Diaz, J. Karhumäki, A. Lepistö and D. Sannella (Eds.), Lecture Notes in Computer Science 3142, pp. 645–657, July 2004.

[16] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis. Extreme Nash Equilibria. *Proceedings of the 8th Italian Conference on Theoretical Computer Science*, C. Blundo and C. Laneve (Eds.), pp. 1–20, Lecture Notes in Computer Science 2841, Springer-Verlag, Berlin, October 2003.

[17] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis. Structure and Complexity of Extreme Nash Equilibria. *Theoretical Computer Science*, Special Issue "Game Theory Meets Theoretical Computer Science", M. Mavronicolas and S. Abramsky (Guest Eds.), Vol. 343, Nos. 1-2, pp. 133–157, October 2005.

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

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

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

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

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

[23] J. Kleinberg, Y. Rabani, and E. Tardos. Allocating Bandwidth for Bursty Connections. *SIAM Journal on Computing*, Vol. 30, No. 1, pp. 191–217, 2000.

[24] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, Vol. 36, No. 6, pp. 683–693, November 2003.

[25] 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, Lecture Notes in Computer Science 1563, Springer-Verlag, March 1999.

[26] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. *Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science*, V. Diekert and M. Habib (Eds.), Lecture Notes in Computer Science 2996, pp. 547-558, Springer-Verlag, March 2004. Full version accepted to *Theoretical Computer Science*.

[27] T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the Worst-Case Nash equilibrium? *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science*, B. Rovan and P. Vojtas (Eds.), Lecture Notes in Computer Science 2747, pp. 551–561, Springer-Verlag, August 2003.

[28] M. Mavronicolas, P. Panagopoulou and P. Spirakis. A Cost Mechanism for Fair Pricing of Resource Usage. *Proceedings of the 1st International Workshop on Internet and Network Economics*, X. Deng and Y. Ye (Eds.), Lecture Notes in Computer Science 3828, pp. 210–224, Springer-Verlag, December, 2005. Full version acceptred to *Algorithmica*.

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

[30] N. Megiddo. A Note on the Complexity of P-Matrix LCP and Computing an Equilibrium. *IBM Research Report RJ-6439*, IBM Almaden Research Center, 1988.

[31] 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.

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

[33] D. Monderer and L. Shapley. Potential Games. *Games and Economic Behavior*, Vol. 14, pp. 124–143, 1996.

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

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

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

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

[38] 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.

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