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

Dimitris Fotakis[1], Spyros Kontogiannis[1], Elias Koutsoupias[2], Marios Mavronicolas[3], and Paul Spirakis[4]

[1] Max Planck Institut für Informatik, 66123 Saarbrücken, Germany, fotakis@mpi-sb.mpg.de, spyros@mpi-sb.mpg.de
[2] Department of Computer Science, University of California at Los Angeles, & Department of Informatics, University of Athens, Greece, elias@cs.ucla.edu.
[3] Department of Computer Science, University of Cyprus, PO Box 20537, Nicosia CY-1678, Cyprus, mavronic@ucy.ac.cy.
[4] Computer Engineering and Informatics Dept., University of Patras, & Computer Technology Institute, PO Box 1122, 261 10 Patras, Greece, spirakis@cti.gr.

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 with given support characteristics, constructing the *worst* Nash equilibrium (the one with *maximum* social cost), constructing the *best* Nash equilibrium (the one with *minimum* 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 $\mathcal{NP}$-hardness 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.

1 Introduction

Nash equilibrium [14] is arguably the most important solution concept in Game Theory [15]. It may be viewed to represent a steady state of the play of a strategic

* 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.

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. [10, 19]); in fact, it is arguably one of the few, most important algorithmic problems for which no polynomial-time algorithms are known. Indeed, Papadimitriou [18, p. 1] actively advocates the problem of computing Nash equilibria as one of the most significant open problems in Theoretical Computer Science 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 [7], 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.

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.*

Contribution. We start with pure Nash equilibria. By the linearity of the expected latency cost functions we consider, a *mixed,* but not necessarily pure, Nash equilibrium always exists. The first result (Theorem 1), remarked by Kurt Mehlhorn, establishes that a pure Nash equilibrium always exists. To this end, we continue to present an efficient, yet simple algorithm (Theorem 2) that computes a pure Nash equilibrium.

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 both are $\mathcal{NP}$-hard (Theorems 3 and 4).

We now turn to mixed Nash equilibria. Our first major result here is an efficient and elegant algorithm for computing a mixed Nash equilibrium (Theorem 5). More specifically, the algorithm computes a *generalized fully mixed* Nash equilibrium; this is a generalization of *fully mixed* Nash equilibria [9].

We continue to 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 6). In close relation, we have attempted to obtain an analog of this result for the model of arbitrary capacities. We establish that *any* mixed Nash equilibrium, in particular the worst one, incurs a social cost that does not exceed 49.02 times the social cost of the fully mixed Nash equilibrium (Theorem 7). Theorems 6 and 7 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 Theorems 6 and 7) to obtain bounds for the general case.

We then shift gears to study the computational complexity of NASH EQUILIBRIUM SOCIAL COST. We have obtained both negative and positive results here. We first show that the problem is $\#\mathcal{P}$-complete (see, e.g., [16]) in general for the case of mixed Nash equilibria (Theorem 8). On the positive side, we get around the established hardness of computing *exacly* the social cost of any mixed Nash equilibrium by presenting a fully polynomial, randomized approximation scheme for computing the social cost of any given mixed Nash equilibrium to any required degree of approximation (Theorem 9).

We point out that the polynomial algorithms we have presented for the computation of pure and mixed Nash equilibria (Theorems 2 and 5, 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, the hardness results we have obtained (Theorems 3, 4, and 8) 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.

Related Work. The selfish routing game considered in this paper was first introduced by Koutsoupias and Papadimitriou [7] 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 [9], 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.* [6], 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 [10, 19]. All known such algorithms incur exponential running time, with the seminal algorithm of Lemke and Howson [8] being the prime example. Issues of computational complexity for the computation of Nash equilibria in general games have been raised by Megiddo [11], Megiddo and Papadimitriou [12], and Papadimitriou [17]. The $\mathcal{NP}$-hardness of computing a Nash equilibrium of a *general* bimatrix game with maximum payoff has been established by Gilboa and Zemel [3]. 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.

2 Framework

Most of our definitions are patterned after those in [7, Sections 1 & 2] and [9, Section 2].

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

For a probability matrix $\mathbf{P}$, define *indicator variables* $I_i^\ell \in \{0, 1\}$, $i \in [n]$ and $\ell \in [m]$, such that $I_i^\ell = 1$ if and only if $p_i^\ell > 0$. Thus, the support of the mixed strategy for user $i \in [n]$ is the set $\{\ell \in [m] \mid I_i^\ell = 1\}$. 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* [9] if $|view(\ell)| = 1$; thus, there is exactly one user, denoted $s(\ell)$, that considers a solo link ℓ.

Syntactic Classes of Mixed Strategies. By a *syntactic class* of mixed strategies, we mean a class of mixed strategies with common support characteristics. A mixed strategy profile $\mathbf{P}$ is *fully mixed* [9] if for all users $i \in [n]$ and links $j \in [m]$, $I_i^j = 1$. Throughout, we will be considering a pure strategy profile as a special case of a mixed strategy profile. in which all (mixed) strategies are pure. We proceed to define two new variations of fully mixed strategy profiles. A mixed strategy profile $\mathbf{P}$ is *generalized fully mixed* if there exists a subset $\mathsf{Links} \subseteq [m]$ such that for each pair of a user $i \in [n]$, and a link $j \in [m]$, $I_i^j = 1$ if $j \in \mathsf{Links}$ and 0 if $j \notin \mathsf{Links}$. Thus, the fully mixed strategy profile is the special case of generalized fully mixed strategy profiles where $\mathsf{Links} = [m]$.

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 strategy profile $\langle \ell_1, \ell_2, \ldots, \ell_n \rangle$, the *latency cost for user* $i \in [n]$, 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 strategy 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 strategy profile $\mathbf{P}$, the *expected latency cost* for user $i \in [n]$ 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 i had its traffic been assigned to link ℓ; thus, $\lambda_i^\ell = \frac{w_i + \sum_{k=1,k\neq i} p_k^\ell w_k}{c^\ell} = \frac{(1-p_i^\ell)w_i + W^\ell}{c^\ell}$. For each user $i \in [n]$, the *minimum expected latency cost,* denoted λ_i, is the minimum, over all links $\ell \in [m]$, of the expected latency cost for user i 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 strategy profile $\mathbf{P}$ is the *social cost* [7, 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, $\mathsf{SC}(\mathbf{w}, \mathbf{P}) = \sum_{\langle \ell_1, \ell_2, \ldots, \ell_n \rangle \in [m]^n} \left(\prod_{k=1}^n p_k^{\ell_k} \cdot \max_{\ell \in [m]} \frac{\sum_{k:\ell_k=\ell} w_k}{c^\ell} \right)$. 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* [7, 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 strategy profile $\mathbf{P}$, $\mathsf{OPT}(\mathbf{w})$ refers to the *optimum* pure strategy profile.

Nash Equilibria. We are interested in a special class of mixed strategies called Nash equilibria [14] that we describe below. Formally, the probability matrix $\mathbf{P}$ is a *Nash equilibrium* [7, Section 2] if for all users $i \in [n]$ and links $\ell \in [m]$, $\lambda_i^\ell = \lambda_i$ if $I_i^\ell = 1$, and $\lambda_i^\ell > \lambda_i$ if $I_i^\ell = 0$. 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.

For each link $\ell \in [m]$, denote $\widetilde{c^\ell} = c^\ell / (\sum_{j=1}^{n} c^j)$, the *normalized capacity* of link ℓ. The following result due to Mavronicolas and Spirakis [9, Theorem 14] provides necessary and sufficient conditions for the existence (and uniqueness) of Nash equilibria in the case of fully mixed strategies, assuming that all traffics are identical.

Lemma 1 (Mavronicolas and Spirakis [9]). *Consider the case of fully mixed strategy profiles, under the model of arbitrary capacities. Assume that all traffics are identical. Then, for all links $\ell \in [m]$, $\widetilde{c^\ell} \in \left(\frac{1}{m+n-1}, \frac{n}{m+n-1}\right)$ if and only if there exists a Nash equilibrium, which must be unique.*

We remark that although, apparently, Lemma 1 determines a collection of $2m$ necessary and sufficient conditions (m pairs with two conditions per pair) for a fully mixed Nash equilibrium, the fact that all normalized capacities sum to 1 implies that each pair reduces to one condition (say the one establishing the lower bound for c^ℓ, $\ell \in [m]$. Furthermore, all m conditions hold if (and only if) the one for $\min_{\ell \in [m]} c^\ell$ holds. Thus, Lemma 1 establishes that existence of a fully mixed Nash equilibrium can be decided in $\Theta(m)$ time by finding the minimum capacity c^{ℓ_0} and checking whether or not the corresponding normalized capacity $\widetilde{c^{\ell_0}}$ satisfies $\widetilde{c^{\ell_0}} > \frac{1}{m+n-1}$. (This observation is due to B. Monien [13].)
Algorithmic Problems. We now formally define several algorithmic problems related to Nash equilibria. A typical instance is defined by: 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.

In NASH EQUILIBRIUM SUPPORTS, we want to compute 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.

In BEST NASH EQUILIBRIUM SUPPORTS, we seek the user supports corresponding to the Nash equilibrium with the *minimum* social cost for the given system of users and links.

In WORST NASH EQUILIBRIUM SUPPORTS, we seek the user supports defining the Nash equilibrium with the *maximum* social cost for the given system of users and links.

NASH EQUILIBRIUM SOCIAL COST is a problem of a somehow counting nature. In addition to the user traffics and the link capacities, an instance is defined by a *Nash equilibrium* $\mathbf{P}$ for the system of the users and the links, and we want to compute the social cost of the Nash equilibrium $\mathbf{P}$.

3 Pure Nash Equilibria

We start with a preliminary result remarked by Kurt Mehlhorn.

Theorem 1. *There exists at least one pure Nash equilibrium.*

Proof sketch. Consider the universe of pure strategy 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 $\mathbf{\Lambda}_0$ and assume that it corresponds to a pure strategy 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 $i \in [n]$ assigned by $\mathbf{P}_0$ to link $j \in [m]$, and a link $\kappa \in [m]$ such that $\Lambda^j > \Lambda^\kappa + \frac{w_i}{c^\kappa}$. Construct now from $\mathbf{P}_0$ a pure strategy profile $\widehat{\mathbf{P}_0}$ which is identical to $\mathbf{P}_0$ except that user i 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, $\widehat{\Lambda^j} = \Lambda^j - \frac{w_i}{c^j} \; < \; \Lambda^j$, while by construction and assumption, $\widehat{\Lambda^\kappa} = \Lambda^\kappa + \frac{w_i}{c^\kappa} \; < \; \Lambda^j$. Since $\mathbf{\Lambda}_0$ is sorted in non-increasing order and $\Lambda^\kappa + \frac{w_i}{c^\kappa} < \Lambda^j$, Λ^j precedes Λ^κ in $\mathbf{\Lambda}_0$. Clearly, all entries preceding Λ^j in $\mathbf{\Lambda}_0$ remain unchanged in $\widehat{\mathbf{\Lambda}_0}$. Consider now the j-th entry of $\widehat{\mathbf{\Lambda}_0}$. There are three possibilities. The j-th entry of $\widehat{\mathbf{\Lambda}_0}$ is either $\widehat{\Lambda^j}$, or $\widehat{\Lambda^\kappa}$, or some entry of $\mathbf{\Lambda}_0$ that followed Λ^j in $\mathbf{\Lambda}_0$ and remained unchanged in $\widehat{\mathbf{\Lambda}_0}$. We obtain a contradiction in all possible cases. □

We remark that the proof of Theorem 1 establishes that the lexicographically minimum expected traffic vector represents a (pure) Nash equilibrium. Since there are exponentially many pure strategy profiles and that many expected traffic vectors, Theorem 1 only provides an *inefficient* proof of existence of pure Nash equilibria (cf. Papadimitriou [17]).

Computing a Pure Nash Equilibrium. We show:

Theorem 2. NASH EQUILIBRIUM SUPPORTS *is in $\mathcal{P}$ when restricted to pure equilibria.*

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

(This nice observation is due to B. Monien [13].) We remark that A_{pure} can be viewed as a variant of Graham's Longest Processing Time (LPT [4]) 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.

Computing the Supports of the Best or Worst Pure Nash Equilibria. We show:

Theorem 3. BEST NASH EQUILIBRIUM SUPPORTS *is $\mathcal{NP}$-hard.*

Proof sketch. Reduction from BIN PACKING (see, e.g., [16]). □

Theorem 4. WORST NASH EQUILIBRIUM SUPPORTS *is $\mathcal{NP}$-hard when restricted to pure equilibria.*

Proof sketch. Reduction from BIN PACKING (see, e.g., [16]). □

4 Mixed Nash Equilibria

We present a polynomial upper bound on the complexity of computing a mixed Nash equilibrium for the case where all traffics are identical. We show:

Theorem 5. *Assume that all traffics are identical. Then,* NASH EQUILIBRIUM SUPPORTS *is in $\mathcal{P}$ when it asks for the supports of a mixed equilibrium.*

Proof sketch. We present a polynomial-time algorithm A_{gfm} that computes the supports of a generalized fully mixed Nash equilibrium. We start with an informal description of A_{gfm}. In a preprocessing step, A_{gfm} sorts all capacities and computes all normalized capacities. Roughly speaking, A_{gfm} considers all subsets of fast links, starting with the set of all links; for each subset, it checks whether there exists a Nash equilibrium for the system of all users and the links in the subset, by using Lemma 1 (and the discussion following it). The algorithm A_{gfm} stops when it finds one; else, it drops the slowest link in the subset and continues recursively. Assume wlog that $c^1 \geq c^2 \geq \ldots \geq c^m$. For any integer m', where $1 \leq m' \leq m$, call a set of links $\{\ell_1, \ldots, \ell_{m'}\}$ a *fast link set*. So, we observe that A_{gfm} examines all generalized fully mixed strategy profiles for a system of all users and a fast link set. Hence, to establish correctness for A_{gfm}, we need to show that at least one of the generalized fully mixed strategy profiles for a system of all users and a fast link set is a Nash equilibrium. We show this by induction on m. □

We note that the preprocessing step of A_{gfm} takes $\Theta(m \lg m) + \Theta(m) = \Theta(m \lg m)$ time. Next, the initial step of A_{gfm} (which considers all links) checks the validity of a single condition (by the discussion following Lemma 1). After this, the loop is executed at most $m-1$ times. For $1 \leq m' \leq m-1$, the m'-th execution checks the validity of a single condition (by the discussion following Lemma 1) and the validity of an additional condition (from the definition of Nash equilibrium). Thus, the time complexity of A_{gfm} is at most $\Theta(m \lg m) + \sum_{1 \leq m' \leq m-1} 2 = \Theta(m \lg m)$.

A Characterization of the Worst Mixed Nash Equilibrium. We first prove a structural property of mixed Nash equilibria, which we then use to provide a syntactic characterization of the worst mixed Nash equilibrium under the model of uniform capacities.

For the following proposition, recall the concepts of solo link and view. In addition, let us say that a user *crosses* another user if their supports cross each other, i.e. their supports are neither disjoint nor the same.

Proposition 1. *In any Nash equilibrium* $\mathbf{P}$ *under the model of uniform capacities,* $\mathbf{P}$ *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 $\ell_0 \in \mathit{support}(s(\ell))$ and a user $i_0 \in \mathit{view}(\ell_0)$, so that $p_{i_0}^{\ell_0} > 0$. Therefore, $\lambda_{s(\ell)}^{\ell_0} = w_{s(\ell)} + p_{i_0}^{\ell_0} w_{i_0} > w_{s(\ell)} = \lambda_{s(\ell)}^{\ell}$, which contradicts the hypothesis that $\mathbf{P}$ is a Nash equilibrium. □

Theorem 6. *Consider the model of uniform capacities and assume that* $n = 2$*. Then, the worst Nash equilibrium is the fully mixed Nash equilibrium.*

Proof sketch. Assume wlog that $w_1 \geq w_2$, and consider any Nash equilibrium $\mathbf{P}$. If $\mathbf{P}$ is pure, we observe that it is not possible for both users to have the same pure strategy. This implies that the social cost of any pure equilibrium is $\max\{w_1, w_2\} = w_1$. If $\mathbf{P}$ is a mixed equilibrium, Proposition 1 implies that there are only two cases to consider: either $\mathit{support}(1) \cap \mathit{support}(2) = \emptyset$ or $\mathit{support}(1) = \mathit{support}(2)$. In the former case, the social cost is w_1. In the latter case, the only possible Nash equilibrium is the fully mixed one having social cost $\mathsf{SC}(\mathbf{w},\mathbf{P}) = w_1 + w_2 \sum_{\ell \in [m]} p_1^\ell p_2^\ell = w_1 + w_2 \cdot \frac{1}{m}$. □

Worst Versus Fully Mixed Nash Equilibria. We show:

Theorem 7. *Consider the model of identical traffics. Then, the social cost of the worst mixed Nash equilibrium is at most* 49.02 *times the social cost of any generalized fully mixed Nash equilibrium.*

Proof sketch. Assume that $c^1 \geq c^2 \geq \ldots \geq c^m$. Let $C_{\mathrm{tot}} = \sum_{\ell=1}^{m} c^\ell$. Wlog, we can assume that the minimum link capacity is 1 (i.e., $c^m = 1$) and that all users have a unit amount of traffic (i.e. for all $i \in [n]$, $w_i = w = 1$). We use $\mathbf{1}$ to denote the corresponding traffic vector. It can be easily verified that it suffices to show the following:

Lemma 2. *Let* $c^m = 1$*,* $c^1 \geq 2$*, and* $n \geq \min\{3 \ln m, C_{\mathrm{tot}} - m + 2\}$*. In addition, let* $\overline{\mathbf{P}}$ *be the generalized fully mixed strategy profile computed by the algorithm* A_{gfm}*, and let* $\mathbf{P}$ *be a strategy profile corresponding to an arbitrary Nash equilibrium. Then,* $49.02\, \mathsf{SC}(\mathbf{1}, \overline{\mathbf{P}}) \geq \mathsf{SC}(\mathbf{1}, \mathbf{P})$*.*

Proof sketch. To distinguish between the expected latencies of the generalized fully mixed Nash equilibrium defined by $\overline{\mathbf{P}}$ and the expected latencies of the Nash equilibrium defined by $\mathbf{P}$, throughout our proof, we use $\overline{\Lambda^1}, \ldots, \overline{\Lambda^m}$ to denote the expected link latencies in the generalized fully mixed equilibrium computed by the algorithm A_{gfm}, and $\Lambda^1, \ldots, \Lambda^m$ to denote the expected link latencies in $\mathbf{P}$. In addition, we use $\overline{\Lambda_{\max}}$ and $\Lambda_{\max}$ to denote the maximum link latency in the $\overline{\mathbf{P}}$ and $\mathbf{P}$, respectively. We first show some bounds on the expected link latencies for any Nash equilibrium. The first bound states that, for any non-solo link $\ell \in [m]$, the expected latency of ℓ in any Nash equilibrium is bounded from above by a small constant factor times the expected latency of the same link in the generalized fully mixed equilibrium computed by A_{gfm}, i.e. $\Lambda^\ell \leq 2\left(1 + \frac{1}{|\mathit{view}(\ell)|-1}\right) \overline{\Lambda^\ell}$, From now on, the analysis focuses on mixed strategy profiles $\mathbf{P}$. Then, we upper bound the probability that the maximum latency of the generalized fully mixed equilibrium does not exceed a given number μ by proving that for any $\mu \in \left[\overline{\Lambda^1}, \frac{n}{c^1}\right]$, $\Pr[\overline{\Lambda_{\max}} < \mu] \leq 4\exp\left(-\sum_{j=1}^{m}\left(\frac{\overline{\Lambda^j}}{2\mu}\right)^{\mu c^j}\right)$. We also bound from above the probability that the maximum latency of the equilibrium $\mathbf{P}$ is greater than or equal to a given number μ. In particular, we prove that for any $\mu > \frac{n}{C_{\text{tot}}} + \frac{m-1}{C_{\text{tot}}}$, $\Pr[\Lambda_{\max} \geq \mu] \leq \sum_{j=1}^{m}\left(\frac{e\Lambda^j}{\mu}\right)^{\mu c^j}$. So, we combine these to show that there exists a number $\mu^* \geq \overline{\Lambda^1}$ such that $\mathsf{SC}(\mathbf{1}, \overline{\mathbf{P}}) \geq \frac{\mu^*}{3}$ and $6e(1+0.0018625)\mu^* \geq \mathsf{SC}(\mathbf{1}, \mathbf{P})$. □

The proof is now complete. □

Computing the Social Cost of a Mixed Nash Equilibrium. We show:

Theorem 8. NASH EQUILIBRIUM SOCIAL COST *is $\#\mathcal{P}$-complete when restricted to mixed equilibria.*

Proof sketch. First of all, we remark that given a set of integer weights $J = \{w_1, \ldots, w_n\}$ and an integer $C \geq \frac{w_1+\cdots+w_n}{2}$, it is $\#\mathcal{P}$-complete to count the number of subsets of J with total weight at most C, since this corresponds to counting the number of feasible solutions of a KNAPSACK instance (see, e.g., [16]). Therefore, given n Bernoulli random variables Y_i, each taking an integer value w_i with probability $\frac{1}{2}$ and 0 otherwise, and an integer C as above, it is $\#\mathcal{P}$-complete to compute the probability that $Y = \sum_{i=1}^{n} Y_i$ exceeds C, i.e. $\Pr(Y \leq C)$. We show that two calls to a (hypothetical) oracle computing the social cost of a given mixed Nash equilibrium suffice to compute the above probability.

Given the random variables Y_i, we consider three identical capacity links (denoted as links 0, 1, 2 respectively) and $n+1$ users, where the user 0 has traffic C and the user i, $i \in [n]$, has traffic w_i. Since $C \geq \frac{w_1+\cdots+w_n}{2}$, if user 0 chooses link 0 with certainty (i.e. $p_0^0 = 1$) and each of the remaining users chooses link 1 or 2 with probability $\frac{1}{2}$ (i.e. $p_i^1 = p_i^2 = \frac{1}{2}$), this mixed strategy profile corresponds to a Nash equilibrium. In addition, since w_i's are integers, the social cost equals $\mathsf{SC}_1 = C + 2\sum_{B=C+1}^{\infty}\Pr(Y \geq B)$. If we increase the traffic of user 0 to $C+1$, the social cost becomes $\mathsf{SC}_2 = C + 1 + 2\sum_{B=C+2}^{\infty}\Pr(Y \geq B)$. Therefore, $2\Pr(Y \geq C+1) = 1 + \mathsf{SC}_1 - \mathsf{SC}_2$, and since C and w_i's are integers, $\Pr(Y \leq C) = 1 - \Pr(Y \geq C+1) = \frac{1-\mathsf{SC}_1+\mathsf{SC}_2}{2}$. □

Approximating the Social Cost of a Mixed Nash Equilibrium. We show:

Theorem 9. *Consider the model of uniform capacities. Then, there exists a fully polynomial, randomized approximation scheme for* NASH EQUILIBRIUM SOCIAL COST.

Proof sketch. The idea of the scheme is to define an efficiently samplable random variable Λ which accurately estimates the social cost of the given Nash equilibrium P on a (given) traffic vector $\mathbf{w}$. For this, we design the following experiment, where N is a fixed parameter 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. Define Λ to be the random variable representing the maximum latency (over all links); for each experiment E_i, $1 \leq i \leq N$, denote Λ_i the measured value for Λ. Output the mean $\frac{\sum_{r=1}^{N} \Lambda_r}{N}$ of the measured values." By the Strong Law of Large Numbers (see, e.g., [5, Section 7.5]), it follows that $\left| \frac{\sum_{r=1}^{N} \Lambda_r}{N} - \mathsf{SC}(\mathbf{w}, \mathbf{P}) \right| \leq \varepsilon\, \mathsf{SC}(\mathbf{w}, \mathbf{P})$, for any constant $\varepsilon > 0$ provided that $N \geq \frac{1}{\varepsilon}\ \mathsf{SC}(\mathbf{w}, \mathbf{P})$. By the results in [1,6], $\mathsf{SC}(\mathbf{w}, \mathbf{P}) = O\left(\frac{\lg n}{\lg \lg n}\right) \cdot \mathsf{OPT}(\mathbf{w})$. Since $\mathsf{OPT}(\mathbf{w}) \leq \sum_{i=1}^{n} w_i$, it follows that $\mathsf{SC}(\mathbf{w}, \mathbf{P}) = O\left(\frac{\lg n}{\lg \lg n}\right) \cdot \sum_{i=1}^{n} w_i$. It suffices to take N to be $\frac{1}{\varepsilon}$ times this upper bound on $\mathsf{SC}(\mathbf{w}, \mathbf{P})$. □

5 Open Problems

Our work leaves open numerous interesting questions that are directly related to our results. We list a few of them here: What is the time complexity of computing the supports of a pure Nash equilibrium? Theorem 2 shows that it is $O(n \lg n + nm) = O(n \max\{\lg n, m\})$. Can this be further improved? Consider the *specific* pure Nash equilibria that are computed by the algorithm that is implicit in the proof of Theorem 1 and the algorithm A_{pure} in the proof of Theorem 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 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.

It is tempting to conjecture that Theorem 6 holds for all values of $n \geq 2$. In addition, we conjecture that the generalized fully mixed strategy is actually the worst-case Nash equilibrium for identical traffics and capacitated links (Theorem 7 proves that it is already within constant factor from the worst case social cost).

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.

Some additional results on the combinatorial structure and the computational complexity of Nash equilibria for the selfish routing game considered in this paper were obtained recently in a follow-up work by Burkhard Monien [13].

Acknowledgments. We thank Kurt Mehlhorn for contributing Theorem 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). We thank 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 f̀ur Informatik.

That discussion contributed some initial ideas for the design of Algorithm A_{pure} (Theorem 2). 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.
2. X. Deng, C. Papadimitriou and S. Safra, "On the Complexity of Equilibria", *Proceedings of the 34th Annual ACM Symposium on Theory of Computing,* May 2002.
3. I. Gilboa and E. Zemel, "Nash and Correlated Equilibria: Some Complexity Considerations", *Games and Economic Behavior,* Vol. 1, pp. 80–93, 1989.
4. R. L. Graham, "Bounds on Multiprocessing Timing Anomalies," *SIAM Journal on Applied Mathematics,* Vol. 17, pp. 416–429, 1969.
5. G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes,* Oxford Science Publications, Second Edition, 1992.
6. 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.
7. E. Koutsoupias and C. H. Papadimitriou, "Worst-case Equilibria," *Proceedings of the 16th Symposium on Theoretical Aspects of Computer Science,* LNCS 1563, pp. 404–413, 1999.
8. 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.
9. M. Mavronicolas and P. Spirakis, "The Price of Selfish Routing," *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing,* pp. 510–519, 2001.
10. 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.
11. 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.
12. 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.
13. B. Monien, Personal Communication, April 2002.
14. J. F. Nash, "Non-cooperative Games," *Annals of Mathematics,* Vol. 54, No. 2, pp. 286–295, 1951.
15. M. J. Osborne and A. Rubinstein, *A Course in Game Theory,* MIT Press, 1994.
16. C. H. Papadimitriou, *Computational Complexity,* Addison-Wesley, 1994.
17. 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.
18. C. H. Papadimitriou, "Algorithms, Games and the Internet," *Proceedings of the 28th International Colloquium on Automata, Languages and Programming*, LNCS 2076, pp. 1–3, 2001.
19. 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, 1998.