

Network Uncertainty in Selfish Routing

Chryssis Georgiou, Theophanis Pavlides, Anna Philippou
Department of Computer Science
University of Cyprus
CY - 1678, Nicosia, Cyprus
{chryssis, phanosp, annap}@cs.ucy.ac.cy

Abstract

We study the problem of selfish routing in the presence of incomplete network information. Our model consists of a number of users who wish to route their traffic on a network of m parallel links with the objective of minimizing their latency. However, in doing so, they face the challenge of lack of precise information on the capacity of the network links. This uncertainty is modelled via a set of probability distributions over all the possibilities, one for each user. The resulting model is an amalgamation of the KP-model of [13] and the congestion games with user-specific functions of [17].

We embark on a study of Nash equilibria and the price of anarchy in this new model. In particular, we propose polynomial-time algorithms for computing some special cases of pure Nash equilibria and we show that negative results of [17], for the non-existence of pure Nash equilibria in the case of three users, do not apply to our model. Consequently, we propose an interesting open problem in this area, that of the existence of pure Nash equilibria in the general case of our model. Furthermore, we consider appropriate notions for the social cost and the price of anarchy and obtain upper bounds for the latter. With respect to fully mixed Nash equilibria, we propose a method to compute them and show that when they exist they are unique. Finally we prove that the fully mixed Nash equilibrium maximizes the social welfare.

1. Introduction

In their pioneering work Koutsoupias and Papadimitriou [13] introduce a non-cooperative *weighted congestion game* (named in the literature as the KP-model)

This work is partially supported by the IST Programs of the European Union under contract numbers IST-2004-001907 (DELIS) and IST-2005-015964 (AEOLUS).

where n selfish users wish to route their unsplitable traffic onto m parallel links from a source to a destination. In this class of games, each link has a certain capacity representing the rate at which the link processes traffic, and users have complete knowledge of the system's parameters such as the link capacities and the traffic of other users. Furthermore, users choose how to route their traffic based on a common payoff function, which essentially captures the delay to be experienced on each link. However, modern non-cooperative systems, such as computer networks and the Internet, which have motivated the study of games such as that of [13], present incomplete information on various aspects of their behavior. For example, it is often the case that network users have incomplete information regarding the link capacities. Such uncertainty may arise if the network links represent complex paths created by routers which are constructed differently on separate occasions according to the presence of congestion or link failures.

In this paper we introduce an extension of the KP-model that captures these more realistic network scenarios. We consider a model where the network links are associated with a number of different capacities and each user's uncertainty about the capacity of the links is modelled via a probability distribution over all the possibilities. We assume that users may have different sources of information regarding the network, therefore, we take their probability distributions to be distinct from one another. This gives rise to a model with user-specific payoff functions, where each user uses its distinct probability distribution to take decisions as to how to route its traffic.

We may see that our model subsumes the KP-model since, in the case of users assigning probability one to the same capacity for each link, the two models coincide. Moreover, our model turns out to be an instance of *weighted congestion games with user-specific functions* studied by Milchtaich in [17].

We are interested in algorithmic problems related to *Nash equilibria* for the routing game we consider, that is, steady states in the game where no user has an incentive to unilaterally change its strategy. For example, we are interested in deciding whether and when Nash equilibria exist in our model, and, if so, determine efficiently the users strategies that give rise to these equilibria. We are concerned with both pure and mixed Nash equilibria, that is equilibria where the strategy of each user is a single link, or a probability distribution over links, respectively. Furthermore, we study two notions for the *social cost* of the game and associated notions for the *price of anarchy* or *coordination ratio* [13] which captures the performance degradation in the game due to the lack of coordination among the users.

Prior Work. Congestion games were first introduced by Rosenthal [23] and studied extensively thereafter. Rosenthal showed that these games admit pure Nash equilibria by using the notion of *potential functions*. Subsequent related work (e.g. [3, 18]) characterized games that admit potential functions as *potential games*. The problem of computing pure Nash equilibria was studied for congestion games in [3] and for weighted congestion games in [6]. The *KP-Model* [13] and its Nash Equilibria were studied extensively in the last years; see, for example, [2, 4, 6, 12, 15, 16] and [5] for a survey. Feldmann et al. [4] and Gairing et al. [7] propose algorithms to transform any user strategy to a Nash equilibrium without increasing the maximum congestion. Fully mixed Nash equilibria for the KP-model were first studied in [16]. The fully mixed Nash equilibrium conjecture, stating that the fully mixed Nash equilibrium has the worst social cost among all Nash equilibria, was first formulated in [7] and it was verified in [14] for a social cost defined as the sum of the users latencies.

The notion of the price of anarchy was first introduced and studied in the KP-Model [22]. Subsequently, tight bounds were proposed for it in [2, 12] for identical links, in [2] for related links, and in [1] for congestion games with linear latency functions.

Gairing et al. [8] were the first to consider an extension of the KP-model with incomplete information. Their model considers a game of parallel links with incomplete information on the traffics of the users, which makes it complementary to our work. The payoff functions employed by the users, which are universal and not user specific, take into account probabilistic information on the user traffics. Based on the seminal work by Harsanyi [11], the authors show that their model always admits a pure Nash equilibrium and propose an

algorithm for computing such an equilibrium for some special cases. Also they show that the fully mixed Nash equilibrium maximizes the social cost for special cases of their model and that, in the general case, more than one fully mixed Nash equilibrium may exist. Finally, they show asymptotically-tight upper bounds on the coordination ratio.

Milchtaich [17] studied congestion games in which the payoff function associated with each user is not universal but user-specific. He shows that these games do not admit a pure Nash equilibrium in the general case, but are guaranteed to exhibit such equilibria in special cases, such as the case of unweighted users. Our work is closely related to [17] since our game is an instance of that model. Thus we inherit the positive results obtained therein. However, we show that the negative results of [17] do not necessarily apply for our model.

Contribution. The contributions of our work are summarized as follows:

- We present an interesting new model that captures the idea of the uncertainty of the network state by defining appropriate user-specific payoff functions.
- We propose polynomial-time algorithms for computing some special cases of pure Nash equilibria and we demonstrate that the counter-example presented in [17] showing that pure Nash equilibria do not exist in the general case does not apply in our model. Thus, we identify an interesting open problem in this area, that of existence of pure Nash equilibria in the general case.
- We identify and employ two different expressions for the social cost and the associated notions for the price of anarchy. We obtain upper bounds for the latter in the general case and for special instances.
- We compute the fully mixed Nash equilibrium and show that when it exists it is unique. Also we show that for certain instances of the game all fully mixed Nash equilibria assign all links to all users equiprobably.
- Finally, we prove that the fully mixed Nash equilibrium maximizes the social welfare.

Document Structure. The paper is organized as follows. In Section 2 we introduce our model. In Section 3 we consider questions of existence and computation of pure Nash equilibria for special cases of our model and we state our conjecture for the existence of pure Nash equilibria in the general case. In Section 4

we study fully mixed Nash equilibria and we give some upper bounds for the coordination ratio. Finally, some of the omitted proofs can be found in [9].

2. Model and Definitions

Network. For all $k \in \mathbb{N}$, denote $[k] = \{1, 2, \dots, k\}$. We consider a *network* consisting of a set of m *parallel links* $1, 2, \dots, m$, or simply *links*, from a source to a destination, and n *network users* $1, 2, \dots, n$, or simply *users*, who wish to route their traffic along a link from the source to the destination. We assume that $n > 1$ and $m > 1$. (Throughout, we will be using subscripts for users and superscripts for links.) We denote by $w_i > 0$ the traffic of network user $i \in [n]$. We define \mathbf{w} as the $n \times 1$ vector containing the traffics of all users.

States. In our model, we assume that there exists uncertainty regarding the capacity of the network links. Thus, we define a *state* to be an $m \times 1$ vector, $\langle c^1, c^2, \dots, c^m \rangle$ where, for all $\ell \in [m]$, $c^\ell > 0$ represents the capacity of link ℓ . The *state space* of the network, denoted by Φ , is defined as the set containing all the possible states the network may realize. We let ϕ range over Φ and we write c_ϕ^ℓ for the capacity of link ℓ according to state ϕ .

Beliefs. Each user, based on some private knowledge, may have a different *belief* regarding the capacity of the network links. We assume that this knowledge is *probabilistic* and it has the form of a probability distribution function over the set of all states. In general, we write $b \in \Delta(\Phi)$ to denote a belief probability distribution over all states, and b_i for the belief of user $i \in [n]$. Furthermore, we write $b(\phi)$ for the probability assigned to state ϕ by belief b . We define the *belief profile* \mathbf{B} to be the $n \times 1$ vector $\langle b_1, b_2, \dots, b_n \rangle$ containing the beliefs of all users.

Games. We consider the routing game $\mathbb{G} = (n, m, \mathbf{w}, \mathbf{B})$ where n is the number of users, m is the number of links, \mathbf{w} is a traffic vector and \mathbf{B} a belief profile. A special instance of this model is the KP-model [13]. This arises whenever, for some $\phi \in \Phi$, $b_i(\phi) = 1$ for all $i \in [n]$.

Strategies. For the remainder of the section let us fix a game $\mathbb{G} = (n, m, \mathbf{w}, \mathbf{B})$. A *pure strategy* for a user $i \in [n]$ is the selection of some link $\ell \in [m]$. A *pure strategies profile* is an n -tuple $\langle \ell_1, \ell_2, \dots, \ell_n \rangle \in [m]^n$ of pure actions, one for each user, where ℓ_i is the selection of user $i \in [n]$. A *mixed strategy* for a user $i \in [n]$ is

a probability distribution $\Delta([m])$ over pure strategies, that is, over the set of links. We denote the probability assigned by user $i \in [n]$ to link $\ell \in [m]$ by p_i^ℓ . A *mixed strategies profile* is an $n \times m$ probability matrix \mathbf{P} , where $P[i, \ell] = p_i^\ell$ is the probability that user i chooses link ℓ . The *support* of the mixed strategy for user i is the set $\{\ell \in [m] \mid p_i^\ell > 0\}$.

Latency. For a pure strategies profile $\sigma = \langle \ell_1, \ell_2, \dots, \ell_n \rangle$, the *latency cost* of user $i \in [n]$ in state ϕ , denoted by $\lambda_{i, \phi}(\sigma)$, is

$$\frac{\sum_{k: \ell_k = \ell_i} w_k}{c_\phi^{\ell_i}}.$$

On the other hand, the *expected latency cost* of user $i \in [n]$ with belief b_i over all states, denoted by $\lambda_{i, b_i}(\sigma)$ is

$$\sum_{\phi \in \Phi} b_i(\phi) \cdot \lambda_{i, \phi}(\sigma).$$

For a mixed strategies profile \mathbf{P} , denote W^ℓ the *expected traffic* on link $\ell \in [m]$, $W^\ell = \sum_{i=1}^n p_i^\ell \cdot w_i$. Denote \mathbf{W} as the $m \times 1$ matrix containing the expected traffics on each link. Furthermore, the *expected latency cost* for user $i \in [n]$ with belief b_i on link $\ell \in [m]$, denoted by $\lambda_{i, b_i}^\ell(\mathbf{P})$, is the expectation over all possible states and over all random choices of the remaining users, of the latency cost for user i when its traffic is assigned to link ℓ . Thus,

$$\begin{aligned} \lambda_{i, b_i}^\ell(\mathbf{P}) &= \sum_{\phi \in \Phi} b_i(\phi) \cdot \frac{w_i + \sum_{k=1, k \neq i}^n p_k^\ell w_k}{c_\phi^\ell} \\ &= \sum_{\phi \in \Phi} b_i(\phi) \cdot \frac{w_i - p_i^\ell w_i + \sum_{k=1}^n p_k^\ell w_k}{c_\phi^\ell} \\ &= \sum_{\phi \in \Phi} b_i(\phi) \cdot \frac{(1 - p_i^\ell) w_i + W^\ell}{c_\phi^\ell}. \end{aligned}$$

For each user $i \in [n]$, with belief b_i , the *minimum expected latency cost* $\lambda_{i, b_i}(\mathbf{P})$ is the minimum, over all links $\ell \in [m]$, of the expected latency cost for user i of belief b_i on link ℓ :

$$\lambda_{i, b_i}(\mathbf{P}) = \min_{\ell \in [m]} \lambda_{i, b_i}^\ell(\mathbf{P}) \quad (1)$$

When \mathbf{P} is inferred from the context we simply write $\lambda_{i, b_i}^\ell, \lambda_{i, b_i}$.

For simplicity, we employ the notation $c_i^\ell = \frac{1}{\sum_{\phi \in \Phi} \frac{b_i(\phi)}{c_\phi^\ell}}$. Now the expected latency cost of user i can be written as

$$\lambda_{i, b_i}^\ell = \frac{(1 - p_i^\ell) w_i + W^\ell}{c_i^\ell}.$$

Solving for p_i^ℓ we have that for every user $i \in [n]$ and link $\ell \in [m]$

$$p_i^\ell = \frac{W^\ell + w_i - c_i^\ell \lambda_{i,b}}{w_i} \quad (2)$$

Nash Equilibrium. The notion of a Nash equilibrium [20, 21] is defined for our model in the usual way. Specifically, a probability matrix \mathbf{P} is a *Nash equilibrium* (often abbreviated as *NE*), if, for all users $i \in [n]$ and for all links $\ell \in [m]$:

$$\lambda_{i,b_i}^\ell \begin{cases} = & \lambda_{i,b_i}, & \text{if } p_i^\ell > 0 \\ \geq & \lambda_{i,b_i}, & \text{if } p_i^\ell = 0 \end{cases}$$

Thus, each user assigns its traffic with positive probability only on links for which its expected latency cost is minimized. This implies that there is no incentive for a user to unilaterally deviate from its strategy to improve its expected latency cost. We refer to probabilities in a Nash equilibrium as *Nash probabilities*.

Social Cost and Coordination Ratio. Associated with a routing game $\mathbb{G} = (n, m, \mathbf{w}, \mathbf{B})$ and a mixed probability matrix \mathbf{P} is the *Social Cost* denoted by $SC(\mathbb{G}, \mathbf{P})$. Since every user's belief for the capacities of the network differs, there is no objective value for the latency of a link or for the exact congestion of the network. Thus, we are forced to depart from the standard definition for the social cost employed in the literature (the expected maximum congestion), and we consider two social cost definitions that take into account the subjective user beliefs as follows:

- the sum of their individual cost: $SC_1(\mathbb{G}, \mathbf{P}) = \sum_{i=1}^n \lambda_{i,b_i}(\mathbf{P})$
- the maximum of their individual cost: $SC_2(\mathbb{G}, \mathbf{P}) = \max_{i \in [n]} \lambda_{i,b_i}(\mathbf{P})$

Similarly, we give two definitions for the *Social Optimum*, or simply the *Optimum*, associated with a routing game $\mathbb{G} = (n, m, \mathbf{w}, \mathbf{B})$, denoted by $OPT(\mathbb{G})$, as the minimum over all pure assignments of

- the sum of their individual cost: $OPT_1(\mathbb{G}) = \min_{\sigma \in [m]^n} \sum_{i=1}^n \lambda_{i,b_i}(\sigma)$
- the maximum of their individual cost: $OPT_2(\mathbb{G}) = \min_{\sigma \in [m]^n} \max_{i \in [n]} \lambda_{i,b_i}(\sigma)$

The above notions give rise to two definitions for the *Coordination Ratio* for our model: $CR_i = \max_{\mathbb{G}, \mathbf{P}} \frac{SC_i(\mathbb{G}, \mathbf{P})}{OPT_i(\mathbb{G})}$, for $i \in \{1, 2\}$.

3. Pure Nash Equilibria

In this section we consider the problem of existence of pure Nash equilibria for our model. It is well known ([3, 23]) that any unweighted congestion game has a pure Nash equilibrium. Furthermore, in the KP-model pure NE exist and can be efficiently computed [6]. On the other hand, in [17] it is shown that weighted congestion games with user-specific functions do not always possess a pure NE. Specifically, a counter-example with three users and three resources (links) is given. For our model, a special case of the games of [17] and an extension of [22], we inherit the positive results of [17]. However, we show that the counter-example of [17] is not valid for our model, since it can be shown that for games with three users NE always exist. In this section, we present polynomial-time algorithms for computing pure Nash equilibria for a number of special cases and we conclude with the conjecture that pure NE exist in the general case in our model.

3.1 Polynomial-Time Algorithms for Special Cases

The case of $m = 2$ links

First we consider the case of an arbitrary number of users n and $m = 2$ links. In fact, our algorithm solves the more general problem of existence of pure NE in the case where the links have some initial traffic $\mathbf{t} = \langle t_1, t_2 \rangle$, where t_i is the initial traffic of link $i \in [m]$. First, we have a useful definition.

Definition 3.1 Consider a game $\mathbb{G} = (n, 2, \mathbf{w}, \mathbf{B})$ with initial traffic $\mathbf{t} = \langle t_1, t_2 \rangle$. We define the tolerance of user i for link j as the value α_i^j which satisfies

$$\frac{t_j + \alpha_i^j}{c_i^j} = \frac{t_{j \oplus 1} + T - \alpha_i^j + w_i}{c_i^{j \oplus 1}}$$

where $T = \sum_{i \in [n]} w_i$ and $a \oplus b = (a + b) \bmod 2$.

Thus, given a two-link game with an associated load T to be assigned on the two links, the tolerance of user i for a link j , α_i^j , is the maximum fragment of the load T the user can tolerate on link j while routing its traffic on it. This implies that, if link j has load α_i^j (and consequently link $j \oplus 1$ has load $T - \alpha_i^j$), user i has no incentive to change its strategy. We have the following lemma:

Lemma 3.2 Consider a strategy $\langle \ell_1, \dots, \ell_n \rangle$ in the game $\mathbb{G} = (n, 2, \mathbf{w}, \mathbf{B})$ and suppose $\ell_1 = 1$. Then, user 1 satisfies the NE condition

$$\frac{t_1 + \sum_{\ell_i=1} w_i}{c_1^1} \leq \frac{t_2 + \sum_{\ell_i=2} w_i + w_1}{c_1^2}$$

if and only if $\sum_{\ell_i=1} w_i \leq \alpha_1^1$.

Proof. First suppose that $\sum_{\ell_i=1} w_i \leq \alpha_1^1$. We have

$$\begin{aligned} \frac{t_1 + \sum_{\ell_i=1} w_i}{c_1^1} &\leq \frac{t_1 + \alpha_1^1}{c_1^1} = \frac{t_2 + \sum_{i \in [n]} w_i - \alpha_1^1 + w_1}{c_1^2} \\ &\leq \frac{t_2 + \sum_{i \in [n]} w_i - \sum_{\ell_i=1} w_i + w_1}{c_1^2} \\ &= \frac{t_2 + \sum_{\ell_i=2} w_i + w_1}{c_1^2} \end{aligned}$$

as required. To prove the other way round, suppose $\sum_{\ell_i=1} w_i > \alpha_1^1$. We have

$$\begin{aligned} \frac{t_1 + \sum_{\ell_i=1} w_i}{c_1^1} &> \frac{t_1 + \alpha_1^1}{c_1^1} = \frac{t_2 + \sum_{i \in [n]} w_i - \alpha_1^1 + w_1}{c_1^2} \\ &> \frac{t_2 + \sum_{i \in [n]} w_i - \sum_{\ell_i=1} w_i + w_1}{c_1^2} \\ &= \frac{t_2 + \sum_{\ell_i=2} w_i + w_1}{c_1^2} \end{aligned}$$

which completes the proof. \square

Figure 1 presents Algorithm A_{twolinks} which solves the problem in $O(n^2)$. It behaves greedily by selecting the user, k , with the highest tolerance over the two links, and it assigns k on the specific link, ℓ_k . It then proceeds to recursively construct an assignment for the remaining users in the same network, but where the initial load on link ℓ_k is increased by w_k . Hereafter, for $\sigma = \langle \ell_1, \dots, \ell_x \rangle$, we write $\sigma[k \mapsto \ell]$, for $\langle \ell_1, \dots, \ell_{k-1}, \ell, \ell_{k+1}, \dots, \ell_x \rangle$.

Theorem 3.3 *For any game $\mathbb{G} = (n, 2, \mathbf{w}, \mathbf{B})$, algorithm A_{twolinks} computes a pure Nash equilibrium in time $O(n^2)$.*

Proof. We will prove the theorem by induction on n . Clearly, for $n = 1$ the claim holds. Assume that the claim holds for any game with $n = \nu$. We will show that it holds for $n = \nu + 1$. Consider the execution of the algorithm with input a game \mathbb{G} with $n = \nu + 1$, and an initial traffic vector \mathbf{t} . The algorithm returns the strategy $\sigma_{\nu+1} = \sigma_\nu[k \mapsto \ell_k]$, where σ_ν is the strategy returned by the recursive call $A_{\text{twolinks}}(\mathbb{G}', \mathbf{t}')$. By the induction hypothesis, since \mathbb{G}' has $n = \nu$, σ_ν is a pure NE for \mathbb{G}_ν and \mathbf{t}' . We have to show that $\sigma_{\nu+1}$ is a pure Nash equilibrium for game $\mathbb{G}_{\nu+1}$ and initial vector \mathbf{t} . First note that, all users in \mathbb{G}_ν satisfy the NE condition for game $\mathbb{G}_{\nu+1}$ and initial vector \mathbf{t} . It remains to show that this assignment is also acceptable for user k . Let N_1, N_2 be the sets of users playing on links 1 and 2, respectively, in game σ_ν . Two cases exist:

Algorithm $A_{\text{twolinks}}(\mathbb{G}, \mathbf{t})$

INPUT: A game $\mathbb{G} = (n, 2, \mathbf{w}, \mathbf{B})$ and a traffic vector $\mathbf{t} = \langle t_1, t_2 \rangle$.

OUTPUT: A pure NE σ .

1. Let $\sigma = \langle \rangle$ and $T = \sum_{k \in [n]} w_k$.
2. Compute c_i^ℓ for every user $i \in [n]$ and every link $\ell \in [m]$.
3. For every user $u \in [n]$ do:
 - (a) For every link $j \in [m]$ do:

Compute $\alpha_u^j = \frac{c_u^1 c_u^2}{c_u^1 + c_u^2} \left(\frac{t_{j \oplus 1} + T + w_u}{c_u^{j \oplus 1}} - \frac{t_j}{c_u^j} \right)$
 - (b) Let $\ell_u = \max(\alpha_u^1, \alpha_u^2)$.
4. Let $k \in [n]$ be the user where $\alpha_k^{\ell_k} = \max_{u \in [n]} \alpha_u^{\ell_u}$, and let $\sigma = \langle \ell_k \rangle$ and $t_{\ell_k} = t_{\ell_k} + w_k$.
5. if $n = 1$ then return σ else
 - (a) Let $\mathbf{w}' = \langle w_1, \dots, w_{k-1}, w_{k+1}, \dots, w_n \rangle$, $\mathbf{B}' = \langle b_1, \dots, b_{k-1}, b_{k+1}, \dots, b_n \rangle$, $\mathbb{G}' = (n-1, 2, \mathbf{w}', \mathbf{B}')$ and let \mathbf{t}' be the new initial traffic vector.
 - (b) return $A_{\text{twolinks}}(\mathbb{G}', \mathbf{t}')[k \mapsto \ell_k]$.

Figure 1. Algorithm A_{twolinks}

- If $N_{\ell_k} = \emptyset$, then the claim follows trivially by observing that $\alpha_k^{\ell_k} \geq w_k$ and Lemma 3.2.
- On the other hand, if $N_{\ell_k} \neq \emptyset$ and $j \in N_{\ell_k}$

$$\frac{t_{\ell_k} + w_k + \sum_{u \in N_{\ell_k}} w_u}{c_j^{\ell_k}} \leq \frac{t_{\ell_k \oplus 1} + \sum_{u \in N_{\ell_k \oplus 1}} w_u + w_j}{c_j^{\ell_k \oplus 1}},$$

which implies, by Lemma 3.2, that $\alpha_j^{\ell_k} \geq \sum_{u \in N_{\ell_k}} w_u + w_k$. Since $\alpha_k^{\ell_k} \geq \alpha_j^{\ell_k}$, by Lemma 3.2,

$$\frac{t_{\ell_k} + w_k + \sum_{u \in N_{\ell_k}} w_u}{c_k^{\ell_k}} \leq \frac{t_{\ell_k \oplus 1} + \sum_{u \in N_{\ell_k \oplus 1}} w_u + w_k}{c_k^{\ell_k \oplus 1}}$$

which completes the proof that $\sigma_{\nu+1}$ is a pure Nash equilibrium.

Finally, we may see that the complexity of the above algorithm is given by the recursive equation $T(n) = T(n-1) + 8n$ whose solution is in $O(n^2)$. \square

The case of symmetric users

In this section we consider the case of *symmetric users*, that is, the case where all users have identical weights,

and we provide an $O(n^2m)$ algorithm for finding a pure NE for the model. Our algorithm, $A_{\text{symmetric}}$, shown in Figure 2, follows along the lines of the constructive proof of [17] for the same problem, for user-specific games. The contribution of our work is a simplification in the correctness proof.

We will be using the following definitions and notations.

- Given a strategy $\sigma = \langle \ell_1, \dots, \ell_n \rangle$, $\ell_i \in [m]$, we define the *state induced by the strategy* as $\mathbf{s} = \langle N_1, \dots, N_m \rangle$, where $N_i = \{j \in [n] \mid \ell_j = i\}$ is the set of users assigned to link i by σ .
- A user is a *defecting user* in a state \mathbf{s} if he does not satisfy the NE property in \mathbf{s} .
- We define the *game graph* of a game as the graph whose nodes are all possible states of the game and there exists an edge between states \mathbf{s} and \mathbf{s}' if $\mathbf{s} = \langle N_1, \dots, N_m \rangle$ and $\mathbf{s}' = \langle N_1, \dots, N_i - \{u\}, \dots, N_j \cup \{u\}, \dots, N_m \rangle$, where u is a defecting user in \mathbf{s} but not in \mathbf{s}' . We write $\mathbf{s} \xrightarrow{u} \mathbf{s}'$.

Algorithm $A_{\text{symmetric}}$

INPUT: A game $\mathbb{G} = (n, m, \langle w, \dots, w \rangle, \mathbf{B})$.

OUTPUT: A pure NE σ .

1. Let $|N_\ell| = 0$ for all $\ell \in [m]$ and $\sigma = \langle \rangle$.
2. Compute c_i^ℓ for every user $i \in [n]$ and every link $\ell \in [m]$.
3. For every user $i \in [n]$ do:
 - (a) Let $\ell \in [m]$ be a link such that $\frac{|N_\ell|+1}{c_i^\ell} \leq \frac{|N_j|+1}{c_i^j}$, $\forall j \neq \ell$
 - (b) Assign user i on link ℓ and set $|N_\ell| = |N_\ell| + 1$.
 - (c) while there exists user k with $\ell_k = \ell$ and $\ell' \in [m]$ such that $(\frac{|N_\ell|}{c_k^\ell} > \frac{|N_{\ell'}|+1}{c_k^{\ell'}})$ do:
 - move user k from link ℓ to link ℓ' , update σ and set
 - $|N_\ell| = |N_\ell| - 1$, $|N_{\ell'}| = |N_{\ell'}| + 1$ and $\ell = \ell'$

Figure 2. Algorithm $A_{\text{symmetric}}$

To prove the correctness of the algorithm we use the following lemma:

Lemma 3.4 Consider state $\langle N_1, \dots, N_m \rangle$ where $i \in N_j$ and suppose that user i satisfies the NE property, that is, for all $k \neq j$:

$$\frac{|N_j|}{c_i^j} \leq \frac{|N_k|+1}{c_i^k}.$$

Then, for any state $\langle L_1, \dots, L_m \rangle$ satisfying $|L_k| \geq |N_k|$ for $i \neq j$ and $|L_j| \leq |N_j|$, user i continues to satisfy the NE property.

Proof. In the new strategy we have for any $k \neq j$:

$$\frac{|L_j|}{c_i^j} \leq \frac{|N_j|}{c_i^j} \leq \frac{|N_k|+1}{c_i^k} \leq \frac{|L_k|+1}{c_i^k}$$

which establishes the claim. \square

Theorem 3.5 Given a game $\mathbb{G} = (n, m, \langle w, \dots, w \rangle, \mathbf{B})$, algorithm $A_{\text{symmetric}}$ computes a pure Nash equilibrium in time $O(n^2m)$.

Proof. Without loss of generality, we assume that $w_i = 1$, for all $i \in [n]$. We will prove the theorem by induction on n . Clearly, for $n = 1$ the claim holds. Assume that the claim holds for $n = \nu - 1$. We will show that it holds for $n = \nu$. By the induction hypothesis, at the end of the $(\nu - 1)^{\text{th}}$ iteration, $\nu - 1$ users have been assigned on links and the assignment induced for this game, say $\sigma_{\nu-1}$, is a Nash equilibrium. In the ν^{th} iteration, user ν assigns its traffic on link j which minimizes its latency (step 3(b)). Then one or more users may wish to deviate from link j to another link. To prove the claim we will show that $\sigma_{\nu-1}$ can be transformed into a NE in $O(\nu)$ moves (step 3(c)).

Let $\mathbf{s}_0 = \langle N_1, \dots, N_m \rangle$ be the state induced by $\sigma_{\nu-1}$. Suppose that, given this state, user P_ν chooses to route its traffic on link 1, giving rise to state

$$\mathbf{s}_1 = \langle N_1 \cup \{\nu\}, N_2, \dots, N_m \rangle.$$

Suppose that this placement results in a sequence of moves

$$\mathbf{s}_1 \xrightarrow{u_2} \mathbf{s}_2 \xrightarrow{u_3} \mathbf{s}_3 \xrightarrow{u_4} \dots$$

where $\mathbf{s}_i = \langle N_i^1, N_i^2, \dots, N_i^m \rangle$. We may prove that for all i , there exists j_i such that

1. $|N_i^{j_i}| = |N_{j_i}| + 1$, and $|N_i^j| = |N_j|$, for $j \neq j_i$,
2. the defecting user $u_{i+1} \in N_i^{j_i}$, and
3. all $u \in N_i^j$, $j \neq j_i$, satisfy the NE criterion.

We may prove this by induction on i . For the base case consider $i = 1$. Clearly, \mathbf{s}_1 satisfies property (1), with $j_1 = 1$. While all users in $\{\nu\} \cup N_2 \cup \dots \cup N_m$ can be seen to continue to satisfy the NE criterion in this new state, it is possible that some user $u \in N_1^1$ is no longer satisfied, that is the defecting user, if one exists, is some $u \in N_1^1$.

Suppose now that the claim holds for $i = k$ and consider $i = k + 1$. We observe, by the induction hypothesis, that, for j_{k+1} the new strategy of the defecting

player u_{k+1} , \mathbf{s}_{k+1} satisfies property (1). In addition, user u_{k+1} is satisfied in \mathbf{s}_{k+1} , and so are all users in $N_{k+1}^q \neq N_{k+1}^{j_{k+1}}$, which completes the induction.

Now consider users $u_1 = \nu, u_2, \dots$, in the execution above. We may see that, since user u_i satisfies the NE criterion in state \mathbf{s}_i , by Lemma 3.4, he will continue to satisfy it in every subsequent step. Thus, user u_i , will not defect in any of the moves following state \mathbf{s}_i , which implies that any user may defect at most once. In other words, the execution is finite and will converge to a NE in at most ν steps. This completes the proof that σ_ν is a pure NE.

To establish that the algorithm is in $O(n^2m)$ it is sufficient to note that in the i^{th} iteration of the algorithm we may have at most $i - 1$ defecting players, which amounts to a total of $O(n^2)$ defecting steps. The candidate users for defection may be identified in a single pass over all players, proceeding step(3(c)), in time $O(nm)$. This completes the proof. \square

The case of uniform user beliefs

We now turn to the *model of uniform user beliefs*, that is, games where each user believes all links to have equal capacity. We propose an algorithm that computes a pure NE for the model in the case where the links have some initial traffic $\mathbf{t} = \langle t_1, \dots, t_m \rangle$, where t_i is the initial traffic of link $i \in [m]$. Our algorithm, A_{uniform} , presented in Figure 3, is a slight modification of the algorithm of [6] (which in turn can be viewed as a variant of Graham's Longest Processing Time (LPT) algorithm [10]) to reflect the needs of our model. Essentially, the algorithm constructs a pure NE in a greedy fashion by processing the users in decreasing order of their weights, and, for each user k , it assigns the user on its preferred link ℓ_k and proceeds with the remaining users in the network where the initial load of link ℓ_k has increased by w_k .

Theorem 3.6 *Given a game $\mathbb{G} = (n, m, \mathbf{w}, \mathbf{B})$ under the model of uniform user beliefs, the algorithm A_{uniform} computes a pure Nash equilibrium in time $O(n(\log n + m))$.*

Proof. We will prove the theorem by induction on n . Clearly, for $n = 1$ the claim holds. Assume that the claim holds for any game with $n = \nu - 1$. We will show that it holds for $n = \nu$. By the induction hypothesis, at the end of the $(\nu - 1)^{\text{th}}$ iteration, $\nu - 1$ users have been assigned on links and the assignment induced for this game, say $\sigma_{\nu-1}$, is a Nash equilibrium. In the ν^{th} iteration, user ν , the user with the smallest weight amongst the ν users, assigns its traffic on link ℓ_ν which minimizes its latency (step 4(b)). We need to show that

Algorithm $A_{\text{uniform}}(\mathbb{G}, \mathbf{t})$

INPUT: A game $\mathbb{G} = (n, m, \mathbf{w}, \mathbf{B})$ and a traffic vector $\mathbf{t} = \langle t_1, \dots, t_m \rangle$.

OUTPUT: A pure NE σ .

1. Write $c_i = c_i^1$ ($= c_i^\ell$, for all ℓ) for every user $i \in [n]$.
2. Let $\sigma = \langle \rangle$.
3. Sort the users in decreasing order of weights.
4. For every user $i \in [n]$ do:
 - (a) Let ℓ_k be a link such that $\frac{w_k + t_{\ell_k}}{c_k} \leq \frac{w_k + t_j}{c_k}, \forall j \neq \ell_k$.
 - (b) Assign user k on link ℓ_k and update σ accordingly.
 - (c) Let $t_{\ell_k} = t_{\ell_k} + w_k$.

Figure 3. Algorithm A_{uniform}

the users in $[\nu - 1]$ remain satisfied after this placement. This clearly affects only the users on link ℓ_ν . Suppose $\sigma_{\nu-1} = \langle \ell_1, \dots, \ell_{\nu-1} \rangle$, and let $W_j = \sum_{\ell_u=j} w_u$, for each link j . Two cases exist:

- If $W_{\ell_\nu} = w_\nu$, that is ν is the only user on ℓ_ν , the claim follows.
- If $W_{\ell_\nu} > w_\nu$ and $k \in [\nu - 1]$ is a user with $\ell_k = \ell_\nu$, then, since user ν is in a Nash equilibrium,

$$\frac{W_{\ell_\nu} + t_{\ell_\nu}}{c_\nu} \leq \frac{w_\nu + W_j + t_j}{c_\nu}$$

for all $j \neq \ell_\nu$, which implies that $W_{\ell_\nu} + t_{\ell_\nu} \leq w_\nu + W_j + t_j$, for all $j \neq \ell_\nu$. Since $w_k \geq w_\nu$, $W_{\ell_\nu} + t_{\ell_\nu} \leq w_k + W_j + t_j$ and, consequently,

$$\frac{W_{\ell_\nu} + t_{\ell_\nu}}{c_k} \leq \frac{w_k + W_j + t_j}{c_k}$$

for all $j \neq \ell_\nu$. Thus, user k does not have an incentive to change strategy, which implies that the assignment σ_ν is a pure Nash equilibrium.

Finally, the complexity of the algorithm is in $O(n(\log n + m))$. \square

The case of $n = 3$

We have shown that any game with three users has a pure NE. The proof employs the notion of a *best-response cycle* which is a cycle in the game graph where defecting users move to the strategy that minimizes their latency. Specifically, the proof establishes in an exhaustive way that the game possesses no *best-response cycles*, which implies that a pure NE exists.

3.2 Existence of Pure Nash Equilibria (Conjecture)

The existence of pure Nash equilibria for this model in the general case remains open. Work for answering this question has been carried out in various directions. In particular, we have shown that the game as defined in this work is not an *exact potential game* [18] (see [9]) and therefore it does not admit an exact potential function. Further, our game is neither an *ordinal potential game*, since it has been shown that the state space of an instance of the game contains a cycle. This interesting observation is due to B. Monien [19]. Therefore, potential functions [18], a popular method for proving existence of NE, cannot be used for our model.

It turns out that the problem of existence of pure NE in our model is a non-trivial problem. Our efforts in applying graph-theoretic methods and inductive arguments have not been successful so far. The arguments end up failing mainly due to the arbitrary relation between the different user beliefs on the capacity of the network links (unlike the special cases presented before where beliefs are related or additional information is present). Naturally, and given the non-existence result on weighted congestion games with user specific payoff-functions [17], we attempted to disprove the existence of NE in our model. Typically, simple counter-examples considering a small number of resources (links) and users are used for such purposes (for example, in [17], the counter example involves 3 users and 3 resources, a scenario for which the existence has been proved in our model). This appears not to be the case for our model as simulations ran on numerous instances of the game (dealing with small number of users and links) suggest the existence of pure NE. Given the lack of a simple counter-example, the polynomial-time algorithms for special cases, and our intuition we conjecture that

Conjecture 3.7 *For any game $\mathbb{G} = (n, m, \mathbf{w}, \mathbf{B})$ there is at least one pure Nash equilibrium.*

4. Fully Mixed Nash Equilibria

In this section we compute the probabilities of the fully mixed Nash equilibria for our model, if one exists. Furthermore, in the case it exists, we illustrate its uniqueness, and that it maximizes the social cost. Finally we show bounds on the Price of Anarchy (Coordination Ratio) that hold in our model.

4.1 Existence and Uniqueness of Fully Mixed Nash Equilibria.

We begin with the following result.

Lemma 4.1 *For a game \mathbb{G} and a Nash Equilibrium \mathbf{P} in the fully mixed case for every user $i \in [n]$, the minimum expected latency cost function λ_{i,b_i} is,*

$$\lambda_{i,b_i} = \frac{1}{\sum_{j \in [m]} c_i^j} \left((m-1)w_i + \sum_{k \in [n]} w_k \right)$$

Proof. For every user $i \in [n]$ it holds that

$$\begin{aligned} \sum_{\ell \in [m]} p_i^\ell &= 1 \stackrel{(2)}{\Leftrightarrow} \sum_{\ell \in [m]} \left(\frac{W^\ell + w_i - c_i^\ell \lambda_{i,b_i}}{w_i} \right) = 1 \\ \Leftrightarrow \frac{1}{w_i} \sum_{\ell \in [m]} W^\ell + m - \frac{\lambda_{i,b_i}}{w_i} \sum_{\ell \in [m]} c_i^\ell &= 1 \\ \Leftrightarrow \frac{\lambda_{i,b_i}}{w_i} \sum_{\ell \in [m]} c_i^\ell &= \frac{1}{w_i} \sum_{\ell \in [m]} W^\ell + m - 1 \\ \Leftrightarrow \lambda_{i,b_i} \sum_{\ell \in [m]} c_i^\ell &= (m-1)w_i + \sum_{\ell \in [m]} W^\ell \\ \Leftrightarrow \lambda_{i,b_i} \sum_{\ell \in [m]} c_i^\ell &= (m-1)w_i + \sum_{\ell \in [m]} \sum_{k \in [n]} p_k^\ell w_k \\ &= (m-1)w_i + \sum_{k \in [n]} w_k \sum_{\ell \in [m]} p_k^\ell \\ \Leftrightarrow \lambda_{i,b_i} \sum_{\ell \in [m]} c_i^\ell &= (m-1)w_i + \sum_{k \in [n]} w_k \\ \Leftrightarrow \lambda_{i,b_i} &= \frac{1}{\sum_{\ell \in [m]} c_i^\ell} \left((m-1)w_i + \sum_{k \in [n]} w_k \right) \end{aligned}$$

□

Lemma 4.2 *For a game \mathbb{G} and a fully mixed Nash Equilibrium \mathbf{P} , for any $\ell \in [m]$ we have*

$$W^\ell = \frac{1}{n-1} \sum_{i=1}^n \left((m-1) \frac{c_i^\ell}{\sum_{j \in [m]} c_i^j} + \sum_{k=1}^n \frac{c_k^\ell}{\sum_{j \in [m]} c_k^j} \right) w_i.$$

Lemma 4.3 *Consider a game \mathbb{G} and a fully mixed Nash Equilibrium \mathbf{P} . Then, for all users $i \in [n]$ and links $\ell \in [m]$ we have*

$$\begin{aligned} p_i^\ell &= 1 - \frac{(m-1)c_i^\ell}{\sum_{j=1}^m c_i^j} - \frac{1}{n-1} \frac{1}{w_i} \left[\sum_{k=1}^n w_k \left(1 - \sum_{k=1}^n \frac{c_k^\ell}{\sum_{j=1}^m c_k^j} \right) \right. \\ &\quad \left. + \frac{(n-1)c_i^\ell}{\sum_{j=1}^m c_i^j} - (m-1) \sum_{k=1}^n \frac{c_k^\ell}{\sum_{j=1}^m c_k^j} w_k \right] \end{aligned}$$

Proof. It follows by substituting the results from Lemma 4.1 and Lemma 4.2 in equation (2). □

Remark 4.4 Lemma 4.3 holds for every probability matrix \mathbf{P} where $p_i^\ell \neq 0$ and $\sum_{\ell=1}^m p_i^\ell = 1$.

We now relate the probabilities calculated in Lemma 4.3 with Nash probabilities.

Lemma 4.5 If for every $i \in [n]$ and every $\ell \in [m]$

$$p_i^\ell = 1 - \frac{(m-1)c_i^\ell}{\sum_{j=1}^m c_i^j} - \frac{1}{n-1} \frac{1}{w_i} \left[\sum_{k=1}^n w_k \left(1 - \sum_{k=1}^n \frac{c_k^\ell}{\sum_{j=1}^m c_k^j} \right) + \frac{(n-1)c_i^\ell}{\sum_{j=1}^m c_i^j} \right] - (m-1) \sum_{k=1}^n \frac{c_k^\ell}{\sum_{j=1}^m c_k^j} w_k \in (0, 1),$$

then in the fully mixed case these probabilities are Nash probabilities.

By Lemma 4.3 and Lemma 4.5 we establish:

Theorem 4.6 (Existence and Uniqueness of Nash Equilibria) Consider the fully mixed case. Then for every user $i \in [n]$ and every link $\ell \in [m]$,

$$p_i^\ell = 1 - \frac{(m-1)c_i^\ell}{\sum_{j=1}^m c_i^j} - \frac{1}{n-1} \frac{1}{w_i} \left[\sum_{k=1}^n w_k \left(1 - \sum_{k=1}^n \frac{c_k^\ell}{\sum_{j=1}^m c_k^j} \right) + \frac{(n-1)c_i^\ell}{\sum_{j=1}^m c_i^j} \right] - (m-1) \sum_{k=1}^n \frac{c_k^\ell}{\sum_{j=1}^m c_k^j} w_k \in (0, 1)$$

if and only if there exist a Nash equilibrium which must be unique and the p_i^ℓ 's are its associated Nash probabilities.

Theorem 4.6 implies the following.

Corollary 4.7 The fully mixed Nash equilibrium when it exists can be calculated in $O(nm)$ time.

From Theorem 4.6 and Lemma 4.3 we get the following result for the model of uniform user beliefs.

Theorem 4.8 Under the model of uniform user beliefs, for any game \mathbb{G} , let \mathbf{F} be the fully mixed Nash equilibrium. Then for any user $i \in [n]$ and any link $\ell \in [m]$ $p_i^\ell = \frac{1}{m}$.

4.2 Worst Case Equilibrium and Price of Anarchy.

We show that the Fully Mixed Nash Equilibrium maximizes social cost. Since both social costs are based on the individual costs of every user $i \in [n]$, we first extend a known relation previously shown in related models [8, 14].

Lemma 4.9 Take any game \mathbb{G} , a (mixed) Nash Equilibria \mathbf{P} and the fully mixed Nash equilibrium \mathbf{F} , Then for any user $i \in [n]$, $\lambda_{i,b_i}(\mathbf{P}) \leq \lambda_{i,b_i}(\mathbf{F})$.

Corollary 4.10 In the case where the Fully Mixed Nash Equilibrium does not exist, then for any probability matrix \mathbf{F} that is derived from Remark 4.4, Lemma 4.9 still holds.

The following two theorems follow from Lemma 4.9 and the definition of the social costs.

Theorem 4.11 The fully mixed Nash equilibrium maximizes the social cost $SC_1(\mathbb{G}, \mathbf{F})$.

Theorem 4.12 The fully mixed Nash equilibrium maximizes the social cost $SC_2(\mathbb{G}, \mathbf{F})$

Corollary 4.10 and Theorems 4.8, 4.11 and 4.12 lead to the following two theorems (the first for the model of uniform user beliefs and the second for the general case):

Theorem 4.13 Take any game \mathbb{G} and any Nash equilibrium \mathbf{P} under the model of uniform user beliefs, then

$$(a) \frac{SC_1(\mathbb{G}, \mathbf{P})}{OPT_1(\mathbb{G})} \leq \left(\frac{c_{max}}{c_{min}} \right) \frac{m+n-1}{m},$$

$$(b) \frac{SC_2(\mathbb{G}, \mathbf{P})}{OPT_2(\mathbb{G})} \leq \left(\frac{c_{max}}{c_{min}} \right) \frac{m+n-1}{m},$$

where $c_{max} = \max_{i \in [n], \ell \in [m]} c_i^\ell$, and $c_{min} = \min_{i \in [n], \ell \in [m]} c_i^\ell$.

Theorem 4.14 Take any game \mathbb{G} and any Nash equilibrium \mathbf{P} , then

$$(a) \frac{SC_1(\mathbb{G}, \mathbf{P})}{OPT_1(\mathbb{G})} \leq \frac{(c_{max})^2}{c_{min}} \frac{1}{\sum_{j=1}^m c_{min}^j} (m+n-1),$$

$$(b) \frac{SC_2(\mathbb{G}, \mathbf{P})}{OPT_2(\mathbb{G})} \leq \frac{(c_{max})^2}{c_{min}} \frac{1}{\sum_{j=1}^m c_{min}^j} (m+n-1),$$

where $c_{max} = \max_{i \in [n], \ell \in [m]} c_i^\ell$, $c_{min} = \min_{i \in [n], \ell \in [m]} c_i^\ell$, and $c_{min}^\ell = \min_{i \in [n]} c_i^\ell$, $\ell \in [m]$.

It remains open to investigate the tightness of the above upper bounds. However, based on results introduced in related works [1, 8] we believe these bounds can be improved.

4.3. Conclusions

In this paper we introduce an extension of the KP-model where the network links may present a number of different capacities and each user's uncertainty about the capacity of the links is modeled via a probability distribution over all the possibilities. This gives rise to

a model with user-specific payoff functions, where each user uses its distinct probability distribution to take decisions as to how to route its traffic.

We embark on a study of Nash equilibria and the price of anarchy in this new model. In particular, we propose polynomial-time algorithms for computing some special cases of pure Nash equilibria and we propose an interesting open problem in this area, that of the existence of pure Nash equilibria in the general case of our model. Furthermore, we consider appropriate notions for the social cost and the price of anarchy and obtain upper bounds for the latter. With respect to fully mixed Nash equilibria, we propose a method to compute them and show that when they exist they are unique. Finally we prove that the fully mixed Nash equilibrium maximizes the social welfare.

Acknowledgments. We would like to thank Marios Mavronicolas and Burkhard Monien for several insightful discussions.

References

- [1] G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 67–73, 2005.
- [2] A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 413–420, 2002.
- [3] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The Complexity of Nash Equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 604–612, 2004.
- [4] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Nashification and the Coordination Ratio for a Selfish Routing Game. In *Proceedings of the 30th International Colloquium of Automata, Languages and Programming*, LNCS 2719, pages 514–526, 2003.
- [5] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish Routing in Non-Cooperative Networks: A Survey. In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science*, LNCS 2747, pages 21–45, 2003.
- [6] D. Fotakis, S. C. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In *Proceedings of the 29th International Colloquium*, LNCS 2380, pages 123–134, 2002.
- [7] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Extreme Nash Equilibria. In *Proceedings of the 8th Italian Conference of Theoretical Computer Science*, LNCS 2841, pages 1–20, 2003.
- [8] M. Gairing, B. Monien, and K. Tiemann. Selfish Routing with Incomplete Information. In *Proceedings of the 17th Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 203–212, 2005.
- [9] C. Georgiou, T. Pavlides, and A. Philippou. Network uncertainty in selfish routing. Technical report, DELIS – Dynamically Evolving, Large-Scale Information Systems, 2005. Submitted for publication.
- [10] R. L. Graham. Bounds on Multiprocessing Timing Anomalies. *SIAM Journal on Applied Mathematics*, 17:416–429, 1969.
- [11] J. C. Harsanyi. Games with Incomplete Information Played by Bayesian Players, I, II, III. *Management Science*, 14:159–182, 320–332, 468–502, 1967.
- [12] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computer Systems*, 36(6):683–693, 2003.
- [13] E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, LNCS 1563, pages 404–413, 1999.
- [14] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. In *Proceedings of the 21st International Symposium of Theoretical Aspects of Computer Science*, LNCS 2996, pages 547–558, 2004.
- [15] T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the Worst-Case Nash Equilibrium? In *Proceedings of the 28th International Symposium of Mathematical Foundations of Computer Science*, LNCS 2747, pages 551–561, 2003.
- [16] M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. In *Proceedings of the 33th Annual ACM Symposium on Theory of Computing*, pages 510–519, 2001.
- [17] I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
- [18] D. Mondener and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [19] B. Monien. Personal Communication. October, 2005.
- [20] J. F. Nash. Equilibrium Points in n -Person Games. *Proceedings of the National Academy of Sciences of the United States of America*, 36:48–49, 1950.
- [21] J. F. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [22] C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of the 33th Annual ACM Symposium on Theory of Computing*, pages 749–753, 2001.
- [23] R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.