

Selfish Routing with Incomplete Information*

Martin Gairing[†]

Burkhard Monien[†]

Karsten Tiemann^{†,‡}

ABSTRACT

In his seminal work Harsanyi [13] introduced an elegant approach to study non-cooperative games with *incomplete information* where the players are uncertain about some parameters. To model such games he introduced the *Harsanyi transformation*, which converts a game with incomplete information to a strategic game where players may have different *types*. In the resulting *Bayesian game* players' uncertainty about each others types is described by a probability distribution over all possible *type profiles*.

In this work, we introduce a particular selfish routing game with incomplete information that we call *Bayesian routing game*. Here, n selfish *users* wish to assign their *traffic* to one of m *links*. Users do not know each others traffic. Following Harsanyi's approach, we introduce for each user a set of possible types.

This paper presents a comprehensive collection of results for the Bayesian routing game.

- We prove, with help of a potential function, that every Bayesian routing game possesses a pure Bayesian Nash equilibrium. For the model of identical links and independent type distribution we give a polynomial time algorithm to compute a pure Bayesian Nash equilibrium.
- We study structural properties of fully mixed Bayesian Nash equilibria for the model of identical links and show that they maximize individual cost. In general there exists more than one fully mixed Bayesian Nash equilibrium. We characterize the class of fully mixed

*This work has been partially supported by the DFG-SFB 376 and by the European Union within the 6th Framework Programme under contract 001907 (DELIS).

[†]Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany. Email: {gairing,bm,tiemann}@uni-paderborn.de

[‡]International Graduate School of Dynamic Intelligent Systems

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SPAA'05, July 18–20, 2005, Las Vegas, Nevada, USA.
Copyright 2005 ACM 1-58113-986-1/05/0007 ...\$5.00.

Bayesian Nash equilibria in the case of independent type distribution.

- We conclude with results on coordination ratio for the model of identical links for three social cost measures, that is, social cost as expected maximum congestion, sum of individual costs and maximum individual cost. For the latter two we are able to give (asymptotic) tight bounds using our results on fully mixed Bayesian Nash equilibria.

To the best of our knowledge this is the first time that mixed Bayesian Nash equilibria have been studied in conjunction with social cost.

Categories and Subject Descriptors

F.2.0 [Theory of Computation]: General

General Terms

Performance, Theory

Keywords

Bayesian Game, Incomplete Information, Selfish Routing, Nash Equilibria, Coordination Ratio

1. INTRODUCTION

Motivation and Framework. In recent years, motivated by non-cooperative systems like the Internet, combining ideas from game theory and theoretical computer science has become more and more attractive. In many of these large-scale non-cooperative systems users have only incomplete information about the system for several reasons. In his seminal work Harsanyi [13] introduced an elegant approach to study non-cooperative games with *incomplete information* where the players are uncertain about some parameters. To model such games he introduced the *Harsanyi transformation*, which converts a game with incomplete information to a strategic game where players may have different *types*. In the resulting *Bayesian game* players' uncertainty about each others type is described by a probability distribution over all possible *type profiles*. Using this probability distribution, players make their decisions according to the concept of *Bayesian decision theory*.

In this work, we introduce a particular selfish routing game with incomplete information that we call *Bayesian routing game*. Here, n selfish *users* wish to assign their *traffic* to one of m *links*. Each link has a certain *capacity* which

specifies the rate at which the link processes traffic. Users do not know each others traffic. Following Harsanyi's approach, we introduce for each user a set of possible types. We assume that these sets are finite. Each type of a user may have different traffic. Furthermore, we assume that there is a *common probability distribution* \mathbf{p} over the set of all possible type realizations. In general \mathbf{p} can be arbitrary, however sometimes we assume \mathbf{p} to be *independent*, that is \mathbf{p} can be expressed by n independent probability distributions, one for each user.

In a *pure strategy* a user chooses for each of its types a particular link whereas in a *mixed strategy* a user uses a probability distribution over all his possible pure strategies. A *strategy profile* specifies a strategy for each of the users. Users choose strategies in order to minimize their *individual cost* which is defined as the expected congestion. Note, that due to the Bayesian model the individual cost in a pure strategy profile is given by the expectation over the type distribution \mathbf{p} . In the case of mixed strategy profiles, the expectation is taken over the strategies of the users and the type distribution \mathbf{p} . Users do not cooperate with each other nor they adhere to a global objective function, the so called *social cost*. A stable state in which no user has an incentive to unilaterally change its strategy is called a *Bayesian Nash equilibrium*. In our study we distinguish between *pure* and *mixed* Bayesian Nash equilibria. Of special interest to our work are *fully mixed* Bayesian Nash equilibria, where each user assigns each of its types to each link with strictly positive probability.

If users are completely informed about each others traffic, that is each user has only a single type, then we are in the setting of a simple model for selfish routing (called KP-model) which was introduced in the pioneering work of Koutsoupias and Papadimitriou [16]. In this setting Bayesian Nash equilibria become Nash equilibria. As in [16], we use the *coordination ratio* or *price of anarchy* as a measure of the maximum performance degradation due to the selfish behavior of the users. The coordination ratio can be defined with respect to different social cost measures.

Contribution. Due to the new dimension that the incomplete information introduces to the routing game, solving problems for the Bayesian routing game requires new techniques. In this paper, we present a comprehensive collection of results for our Bayesian routing game. We partition our results into three major parts:

(1) Existence and computational complexity of pure Bayesian Nash equilibria:

We prove that every Bayesian routing game possesses a pure Bayesian Nash equilibrium. In particular we introduce a potential function and show that its value decreases whenever a user unilaterally changes its strategy to improve its individual cost. This result can also be generalized to a larger class of games, called *weighted Bayesian congestion games*. For the model of identical links and independent type distributions we show that a pure Bayesian Nash equilibrium can be computed in polynomial time. For the model of Bayesian routing games with related links and for the model of identical links and arbitrary type distribution the complexity of determining a pure Bayesian Nash equilibrium remains open.

(2) Properties of fully mixed Bayesian Nash equilibria:

We show that in the model of identical links the individual cost of each user is maximized in a fully mixed Bayesian Nash equilibrium. This also implies that a user has the same individual cost in any fully mixed Bayesian Nash equilibrium. We define a certain fully mixed Bayesian Nash equilibrium that always exists. We show that in general there might exist more than one fully mixed Bayesian Nash equilibrium and we study their structural properties. Finally, we show the dimension of the space of fully mixed Bayesian Nash equilibria in the case of independent type distributions.

(3) Results on Coordination ratio:

We conclude our paper with results on the coordination ratio for three different social cost measures in the model of identical links.

- The expected maximum congestion on a link is a social cost measure that expresses the social welfare of the system. Here, we are able to show lower and upper bounds on the coordination ratio for different special cases. The exact coordination ratio for Bayesian routing games remains open, even for identical links.
- A social cost measure that describes average user welfare is the sum of the individual costs. In this setting, we show that for identical links, each fully mixed Bayesian Nash equilibrium has maximum social cost. Using this fact we proof an upper bound of $\frac{m+n-1}{m}$ on the coordination ratio in case of identical links. This bound is asymptotically tight, even for KP-games.
- We also study social cost as maximum individual cost. For identical links, we show asymptotically tight upper bounds on coordination ratio for Bayesian routing games and KP-games of $\frac{m+n-1}{m}$ and $2 - \frac{1}{m}$, respectively.

To the best of our knowledge this is the first time that mixed Bayesian Nash equilibria are studied in conjunction with social cost.

Related Work. The class of *congestion games* was introduced by Rosenthal [26] and extensively studied afterwards (see e.g. [3, 4, 8, 21, 22, 27]). In Rosenthal's model each player has complete information and as its strategy a subset of resources. Resource utility functions can be arbitrary but they only depend on the number of players sharing the same resource. Rosenthal showed that such games always admit a pure Nash equilibrium using a potential function. Subsequent papers [3, 22, 27] characterize games that possess a potential function as potential games and show their relation to congestion games. The complexity of computing pure Nash equilibria in congestion games was studied by Fabrikant et al. [3]. Milchtaich [21] considers weighted congestion games with player specific payoff functions and shows that these games do not admit a pure Nash equilibrium in general. Fotakis et al. [8] consider weighted congestion games and proved the existence of pure Nash equilibria, if resources have linear cost functions.

The *KP-model* [16] for routing selfish users on parallel links, and its Nash equilibria, were studied extensively in the last years; see, for example, [2, 5, 7, 11, 15, 18], and [6] for a recent survey. Graham's LPT scheduling algorithm

[12] computes a pure Nash equilibrium in the KP-model [7]. Algorithms to transform any assignment to a Nash equilibrium with non-increased maximum congestion have been presented in [5, 9, 11].

The *coordination ratio*, also known as *price of anarchy* [24], was first introduced and studied by Koutsoupias and Papadimitriou [16]. For social cost defined as the expected maximum congestion, there exist *tight* bounds of $\Theta\left(\frac{\log m}{\log \log m}\right)$ for identical links [2, 15] and $\Theta\left(\frac{\log m}{\log \log \log m}\right)$ [2] for related links. The fully mixed Nash equilibrium conjecture, which states that the fully mixed Nash equilibrium has worst social cost among all Nash equilibria, was motivated by some results in [20], explicitly formulated in [11] and further studied in [18]. For social cost defined as the sum of individual costs the conjecture holds [10, 17].

A framework for studying competitive situations where the players have incomplete information was developed by Harsanyi. The Nobel prize-winner introduced Bayesian games in his pioneering work [13, 14]. Facchini et al. [4] consider Bayesian congestion models with players of identical weight but players have incomplete information about each others preferences. Beier et al. [1] focus on a service provider congestion game with incomplete information. For an introduction to Bayesian games we refer to [19, 23].

Road Map. The rest of this paper is organized as follows. Section 2 presents an exact definition of the Bayesian routing games considered in this paper. The existence and computation of pure Bayesian Nash equilibria is studied in Section 3. Some interesting structural properties of fully mixed Bayesian Nash equilibria are presented in Section 4. Section 5 gives a thorough study of the coordination ratio for three definitions of social cost.

2. NOTATION

Instance. For all $k \in \mathbb{N}$ denote $[k] = \{1, \dots, k\}$. We consider the following *Bayesian routing game* $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$. Each of n users $1, 2, \dots, n$ wishes to assign a particular amount of traffic to one of m links $1, 2, \dots, m$. Denote $\mathbf{c} = (c_1, \dots, c_m)$, where c_j is the *capacity* of link $j \in [m]$. In the model of *identical links* all capacities are equal to 1. In this case we write $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$. Link capacities may vary arbitrary in the model of *related links*. For each user $i \in [n]$ there is a set of possible types T_i . Denote $T = T_1 \times \dots \times T_n$ as the set of all possible type profiles. For each type $t \in T_i, i \in [n]$, denote by $w(t)$ the *traffic* of type t .

There is a *common probability distribution* (common for all users) $\mathbf{p} = (p(t_1, \dots, t_n))_{(t_1, \dots, t_n) \in T}$ over the set of possible type profiles T , thus $\mathbf{p} : T \mapsto [0, 1]$. Denote by $p(i, t)$ the probability that user i is of type t , that is, $p(i, t) = \sum_{(t_1, \dots, t_n) \in T, t_i = t} p(t_1, \dots, t_n)$. We say that \mathbf{p} is *independent* if $p(t_1, \dots, t_n) = \prod_{i \in [n]} p(i, t_i)$ for all $(t_1, \dots, t_n) \in T$, otherwise \mathbf{p} is *correlated*. Following the concept of conditional probability $p(t_1, \dots, t_n | t_k = t) = \frac{p(t_1, \dots, t_n)}{p(k, t)}$, that is, the probability of a type profile (t_1, \dots, t_n) given that $t_k = t$ is the probability of type profile (t_1, \dots, t_n) divided by the probability that player k is of type t ; this is known as Bayes' Theorem. Note, that $p(t_1, \dots, t_n | t_k = t) = p(t_1, \dots, t_{k-1}, t_{k+1}, \dots, t_n | t_k = t)$. Denote by $E(i) = \sum_{(t_1, \dots, t_n) \in T} p(t_1, \dots, t_n) \cdot w(t_i)$ the expected traffic of user $i \in [n]$. Furthermore, define the expected total traffic as $E = \sum_{i \in [n]} E(i)$.

A special instance of Bayesian routing games in which each user only has a single type is a *KP-game* (see [16]). For a KP-game we write $\Gamma_{\text{KP}} = (n, m, \mathbf{c}, T, 1)$. The set T contains only one type vector t that is used with probability 1.

Strategies and Strategy Profiles. A *pure strategy* σ_i for user $i \in [n]$ is a mapping of the set of possible types T_i to the set of links $[m]$; thus, $\sigma_i : T_i \mapsto [m]$. Denote Σ_i as the set of all possible pure strategies for user $i \in [n]$ and denote $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$. A *mixed strategy* $Q_i = (q(i, \sigma_i))_{\sigma_i \in \Sigma_i}$ for user $i \in [n]$ is a probability distribution over Σ_i , where $q(i, \sigma_i)$ denotes the probability that user i chooses the pure strategy σ_i . A *pure strategy profile* \mathbf{L} is represented by an n -tuple $(\sigma_1, \dots, \sigma_n) \in \Sigma$. Call \mathbf{L} *normal* if $\sigma_i(t) = \sigma_i(t')$ for all $t, t' \in T_i$ and for all $i \in [n]$. A *mixed strategy profile* $\mathbf{Q} = (Q_1, \dots, Q_n)$ is represented by an n -tuple of mixed strategies. The *support* of a mixed strategy for user $i \in [n]$, denoted $\text{support}(i)$, is the set of links to which user i assigns at least one type $t \in T_i$ with positive probability, that is,

$$\text{support}(i) = \{j \in [m] \mid \exists \sigma_i \in \Sigma_i, \exists t \in T_i \text{ with } q(i, \sigma_i) > 0 \text{ and } \sigma_i(t) = j\}.$$

Similarly, the support of any type $t \in T_i$ of user $i \in [n]$ is defined by

$$\text{support}(t) = \{j \in [m] \mid \exists \sigma_i \in \Sigma_i \text{ with } q(i, \sigma_i) > 0 \wedge \sigma_i(t) = j\}.$$

Call a strategy profile \mathbf{F} *fully mixed* if $\text{support}(t_i) = [m]$ for all $t_i \in T_i, i \in [n]$.

System and Individual Cost. Fix any probability distribution \mathbf{p} over the set of possible type profiles T and a mixed strategy profile \mathbf{Q} . The *expected load* on link $j \in [m]$, denote $\delta_j(\mathbf{Q}, \mathbf{p})$, is defined by

$$\delta_j(\mathbf{Q}, \mathbf{p}) = \sum_{(\sigma_1, \dots, \sigma_n) \in \Sigma} \prod_{i \in [n]} q(i, \sigma_i) \sum_{(t_1, \dots, t_n) \in T} p(t_1, \dots, t_n) \sum_{\substack{i \in [n], \\ \sigma_i(t_i) = j}} w(t_i).$$

In the same way, denote $\delta_j^{-k}(\mathbf{Q}, (\mathbf{p} | t_k = t))$ as the expected load of all users $i \in [n], i \neq k$, on link $j \in [m]$ given that $t_k = t$; thus,

$$\delta_j^{-k}(\mathbf{Q}, (\mathbf{p} | t_k = t)) = \sum_{(\sigma_1, \dots, \sigma_n) \in \Sigma} \prod_{i \in [n]} q(i, \sigma_i) \sum_{\substack{(t_1, \dots, t_n) \in T, \\ t_k = t}} p(t_1, \dots, t_n) \sum_{\substack{i \in [n], i \neq k \\ \sigma_i(t_i) = j}} w(t_i).$$

Denote $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ as the individual cost of user $i \in [n]$, if user i is of type t ; then

$$v_{(i,t)}(\mathbf{Q}, \mathbf{p}) = \sum_{j \in [m]} \sum_{\substack{\sigma_i \in \Sigma_i \\ \sigma_i(t) = j}} q(i, \sigma_i) \frac{\delta_j^{-i}(\mathbf{Q}, (\mathbf{p} | t_i = t)) + w(t)}{c_j}.$$

Note, that $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ does not depend on the other types $t' \in T_i \setminus \{t\}$ of user i .

Moreover, denote $u_i(\mathbf{Q}, \mathbf{p})$ as the *individual cost* of user i ; then

$$u_i(\mathbf{Q}, \mathbf{p}) = \sum_{t \in T_i} p(i, t) \cdot v_{(i,t)}(\mathbf{Q}, \mathbf{p}).$$

A strategy profile \mathbf{Q} is a Bayesian Nash equilibrium, if and only if in \mathbf{Q} no user has an incentive to deviate from its current (mixed) strategy, that is, no user can decrease its *individual cost* if the other users stick to their strategies.

More formally, $\mathbf{Q} = (Q_1, \dots, Q_n)$ is a Bayesian Nash equilibrium if and only if

$$u_i(\mathbf{Q}, \mathbf{p}) \leq u_i(\mathbf{Q}', \mathbf{p})$$

for all $\mathbf{Q}' = (Q_1, \dots, Q'_i, \dots, Q_n)$ and for all $i \in [n]$. Moreover, since $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ does not depend on the other types $t' \in T_i \setminus \{t\}$ of user i , the above condition is equivalent to

$$v_{(i,t)}(\mathbf{Q}, \mathbf{p}) \leq v_{(i,t)}(\mathbf{Q}', \mathbf{p})$$

for all $\mathbf{Q}' = (Q_1, \dots, Q'_i, \dots, Q_n)$ and for all $i \in [n], t \in T_i$.

Social Cost and Coordination Ratio. Associated with a Bayesian routing game $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ and strategy profile \mathbf{Q} is the *social cost* as a measure of social welfare. We consider three social cost definitions:

- the expectation of the expected maximum congestion,

$$\begin{aligned} \text{SC}_{\text{MSP}}(\mathbf{Q}, \Gamma) &= \sum_{(\sigma_1, \dots, \sigma_n) \in \Sigma} \prod_{i \in [n]} q(i, \sigma_i) \sum_{(t_1, \dots, t_n) \in T} p(t_1, \dots, t_n) \\ &\quad \cdot \max_{j \in [m]} \left\{ \frac{\sum_{i \in [n], \sigma_i(t_i)=j} w(t_i)}{c_j} \right\}, \end{aligned}$$

- the sum of individual costs,

$$\text{SC}_{\text{SUM}}(\mathbf{Q}, \Gamma) = \sum_{i \in [n]} u_i(\mathbf{Q}, \mathbf{p}),$$

- the maximum individual cost,

$$\text{SC}_{\text{MAX}}(\mathbf{Q}, \Gamma) = \max_{i \in [n]} u_i(\mathbf{Q}, \mathbf{p}).$$

Let $*$ \in $\{\text{MSP}, \text{SUM}, \text{MAX}\}$. Denote the corresponding optimum social cost by $\text{OPT}_*(\Gamma) = \min_{\mathbf{Q}} \text{SC}_*(\mathbf{Q}, \Gamma)$. The *coordination ratio* CR_* , is the maximum value, over all instances Γ and Bayesian Nash equilibria \mathbf{Q} , of the ratio $\frac{\text{SC}_*(\mathbf{Q}, \Gamma)}{\text{OPT}_*(\Gamma)}$.

Type agent representation. Following Harsanyi's transformation we can also consider each type t of a user i as an independent *type agent* (i, t) . In this setting, a mixed strategy profile is defined by $\mathbf{R} = ((r(i, t, j))_{j \in [m]})_{i \in [n], t \in T_i}$, where $r(i, t, j)$ is the probability that type agent (i, t) chooses link j . From \mathbf{R} we can compute a mixed strategy profile $\mathbf{Q} = ((q(i, \sigma_i)_{\sigma_i \in \Sigma_i})_{i \in [n]})$ by defining

$$q(i, \sigma_i) = \prod_{t \in T_i} r(i, t, \sigma_i(t))$$

for all $i \in [n]$ and all $\sigma_i \in \Sigma_i$. Note, that for pure strategy profiles both representations are isomorphic. A Bayesian Nash equilibrium \mathbf{R} in type agent representation has the property that no type agent (i, t) can unilaterally improve its individual cost, thus,

$$v_{(i,t)}(\mathbf{R}, \mathbf{p}) \leq v_{(i,t)}(\mathbf{R}', \mathbf{p})$$

for all strategy profiles \mathbf{R}' that result from \mathbf{R} when only type agent (i, t) changes strategy.

DEFINITION 2.1. *Given a pure strategy profile \mathbf{R} in type agent representation a selfish step of a type agent (i, t) , $i \in [n]$, $t \in T_i$, is a strategy change of this type agent which improves its individual cost. This means that*

$$v_{(i,t)}(\mathbf{R}, \mathbf{p}) > v_{(i,t)}(\mathbf{R}', \mathbf{p}),$$

where \mathbf{R}' is the strategy profile that result from \mathbf{R} when only type agent (i, t) changes strategy.

One could have used also the type agent representation for defining pure and mixed strategies. We feel that our notation leads to more intuitive formulations in this complex environment.

3. EXISTENCE AND COMPUTATION OF PURE BAYESIAN NASH EQUILIBRIA

In this section we study the existence and the computational complexity of pure Bayesian Nash equilibria. We first show that there is always a pure Bayesian Nash equilibrium in any Bayesian routing game (Theorem 3.1). This result can be generalized to a more general class of games, that we call weighted Bayesian congestion games (Theorem 3.2). We close with a polynomial time algorithm that computes a pure Bayesian Nash equilibrium for a Bayesian routing game with identical links and independent type distribution (Theorem 3.3).

THEOREM 3.1. *Every Bayesian routing game $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ possesses a pure Bayesian Nash equilibrium.*

PROOF. For any pure strategy profile $\mathbf{L} = (\sigma_1, \dots, \sigma_n)$, we define the following potential function:

$$\begin{aligned} \Phi(\mathbf{L}) &= \sum_{i \in [n]} \sum_{t \in T_i} p(i, t) \cdot w(t) \cdot \left(v_{(i,t)}(\mathbf{L}, \mathbf{p}) + \frac{w(t)}{c_{\sigma_i(t)}} \right) \\ &= \sum_{i \in [n]} \sum_{t \in T_i} p(i, t) \cdot w(t) \cdot \frac{\delta_{\sigma_i(t)}^{-i}(\mathbf{L}, (\mathbf{p}|_{t_i = t})) + 2w(t)}{c_{\sigma_i(t)}} \end{aligned}$$

Note, that the potential function sums up over all type agents. Consider a selfish step of type agent (r, t^*) from link k to link j . Define $\mathbf{L}' = (\sigma'_1, \dots, \sigma'_n)$ as the assignment resulting from \mathbf{L} after this selfish step. By definition of selfish step,

$$\begin{aligned} v_{(r,t^*)}(\mathbf{L}, \mathbf{p}) &= \frac{\delta_k^{-r}(\mathbf{L}, (\mathbf{p}|_{t_r = t^*})) + w(t^*)}{c_k} \\ &> \frac{\delta_j^{-r}(\mathbf{L}', (\mathbf{p}|_{t_r = t^*})) + w(t^*)}{c_j} \quad (1) \\ &= v_{(r,t^*)}(\mathbf{L}', \mathbf{p}) \end{aligned}$$

Moreover, $v_{(i,t)}(\mathbf{L}, \mathbf{p}) = v_{(i,t)}(\mathbf{L}', \mathbf{p})$ for all $i \in [n] \setminus \{r\}, t \in T_i, \sigma_i(t) \in [m] \setminus \{j, k\}$ and for all $t \in T_r \setminus \{t^*\}, \sigma_r(t) \in [m] \setminus \{j, k\}$. Now consider the change $\Delta(\Phi)$ due to this selfish step. It is $\Delta(\Phi) = \Phi(\mathbf{L}') - \Phi(\mathbf{L}) = \Delta_1(\Phi) + \Delta_2(\Phi) + \Delta_3(\Phi)$, where $\Delta_1(\Phi)$ is the change in Φ that directly belongs to type agent (r, t^*) , $\Delta_2(\Phi)$ is the change in Φ of type agents assigned to j and $\Delta_3(\Phi)$ is the change in Φ of type agents assigned to k . We have,

$$\begin{aligned} \Delta_1(\Phi) &= p(r, t^*) \cdot w(t^*) \cdot \left(\frac{\delta_j^{-r}(\mathbf{L}, (\mathbf{p}|_{t_r = t^*})) + 2w(t^*)}{c_j} \right. \\ &\quad \left. - \frac{\delta_k^{-r}(\mathbf{L}, (\mathbf{p}|_{t_r = t^*})) + 2w(t^*)}{c_k} \right), \end{aligned}$$

and

$$\begin{aligned}
\Delta_2(\Phi) &= \sum_{\substack{i \in [n], \\ i \neq r}} \sum_{\substack{t \in T_i, \\ \sigma_i(t)=j}} p(i, t) \cdot w(t) \cdot \frac{1}{c_j} \\
&\quad \cdot \left[\delta_j^{-i}(\mathbf{L}', (p|_{t_i=t})) - \delta_j^{-i}(\mathbf{L}, (p|_{t_i=t})) \right] \\
&= \sum_{\substack{i \in [n], \\ i \neq r}} \sum_{\substack{t \in T_i, \\ \sigma_i(t)=j}} p(i, t) \cdot w(t) \cdot \frac{1}{c_j} \\
&\quad \cdot \sum_{\substack{(t_1, \dots, t_n) \in T, \\ t_i=t}} p(t_1, \dots, t_n | t_i=t) \\
&\quad \cdot \left[\sum_{\substack{s \in [n], s \neq i \\ \sigma'_s(t_s)=j}} w(t_s) - \sum_{\substack{s \in [n], s \neq i \\ \sigma_s(t_s)=j}} w(t_s) \right] \\
&= \sum_{\substack{i \in [n], \\ i \neq r}} \sum_{\substack{t \in T_i, \\ \sigma_i(t)=j}} p(i, t) \cdot w(t) \cdot \frac{1}{c_j} \\
&\quad \cdot \sum_{\substack{(t_1, \dots, t_n) \in T, \\ t_i=t, t_r=t^*}} p(t_1, \dots, t_n | t_i=t) \cdot w(t^*) \\
&= \sum_{\substack{i \in [n], \\ i \neq r}} \sum_{\substack{t \in T_i, \\ \sigma_i(t)=j}} w(t) \cdot \frac{1}{c_j} \sum_{\substack{(t_1, \dots, t_n) \in T, \\ t_i=t, t_r=t^*}} p(t_1, \dots, t_n) \cdot w(t^*) \\
&= \frac{w(t^*)}{c_j} \sum_{\substack{i \in [n], \\ i \neq r}} \sum_{\substack{t \in T_i, \\ \sigma_i(t)=j}} \sum_{\substack{(t_1, \dots, t_n) \in T, \\ t_i=t, t_r=t^*}} p(t_1, \dots, t_n) \cdot w(t).
\end{aligned}$$

Hence, by law of conditional probability,

$$\begin{aligned}
\Delta_2(\Phi) &= \frac{p(r, t^*) \cdot w(t^*)}{c_j} \sum_{\substack{i \in [n], \\ i \neq r}} \sum_{\substack{t \in T_i, \\ \sigma_i(t)=j}} \\
&\quad \cdot \sum_{\substack{(t_1, \dots, t_n) \in T, \\ t_i=t, t_r=t^*}} p(t_1, \dots, t_n | t_r=t^*) \cdot w(t) \\
&= \frac{p(r, t^*) \cdot w(t^*)}{c_j} \\
&\quad \cdot \sum_{\substack{i \in [n], \\ i \neq r}} \sum_{\substack{(t_1, \dots, t_n) \in T, \\ \sigma_i(t_i)=j, t_r=t^*}} p(t_1, \dots, t_n | t_r=t^*) \cdot w(t_i) \\
&= \frac{p(r, t^*) \cdot w(t^*)}{c_j} \\
&\quad \cdot \sum_{\substack{(t_1, \dots, t_n) \in T, \\ t_r=t^*}} p(t_1, \dots, t_n | t_r=t^*) \sum_{\substack{i \in [n], i \neq r \\ \sigma_i(t_i)=j}} w(t_i) \\
&= p(r, t^*) \cdot w(t^*) \cdot \frac{\delta_j^{-r}(\mathbf{L}, (\mathbf{p}|_{t_r=t^*}))}{c_j}.
\end{aligned}$$

In the same way

$$\begin{aligned}
\Delta_3(\Phi) &= - \sum_{\substack{i \in [n], \\ i \neq r}} \sum_{\substack{t \in T_i, \\ \sigma_i(t)=k}} p(i, t) \cdot w(t) \cdot \frac{1}{c_k} \\
&\quad \cdot \left[\delta_j^{-i}(\mathbf{L}', (p|_{t_i=t})) - \delta_j^{-i}(\mathbf{L}, (p|_{t_i=t})) \right] \\
&= -p(r, t^*) \cdot w(t^*) \cdot \frac{\delta_j^{-r}(\mathbf{L}, (\mathbf{p}|_{t_r=t^*}))}{c_k}
\end{aligned}$$

We get,

$$\begin{aligned}
\Delta(\Phi) &= \Delta_1(\Phi) + \Delta_2(\Phi) + \Delta_3(\Phi) \\
&= 2 \cdot p(r, t^*) \cdot w(t^*) \cdot \left(\frac{\delta_j^{-r}(\mathbf{L}, (\mathbf{p}|_{t_r=t^*})) + w(t^*)}{c_j} \right. \\
&\quad \left. - \frac{\delta_k^{-r}(\mathbf{L}, (\mathbf{p}|_{t_r=t^*})) + w(t^*)}{c_k} \right) \\
&< 0,
\end{aligned}$$

where the last inequality follows from Equation (1). Thus, any selfish step decreases the value of the potential function $\Phi(\mathbf{L})$. Since the number of possible assignments is finite, the claim follows. \square

A generalization of the Bayesian routing game considered in this paper is a *weighted Bayesian congestion game with linear cost functions*. In a congestion game [26] each user $i \in [n]$ can assign its traffic to a subset s_i of the resources out of a given set $S_i \subseteq 2^{[m]}$ of subsets of resources. The cost function of resource $e \in [m]$ is given by an arbitrary non-decreasing linear cost function $f_e(x) = a_e x + b_e$. A pure strategy profile \mathbf{L} is then defined by $\mathbf{L} = (\sigma_1, \dots, \sigma_n)$ with $\sigma_i : T_i \mapsto S_i$ for all $i \in [n]$. We can generalize Theorem 3.1 to this setting.

THEOREM 3.2. *Every weighted Bayesian congestion game Γ_C with linear cost functions possesses a pure Bayesian Nash equilibrium.*

PROOF. For a pure strategy profile \mathbf{L} define the following potential function

$$\begin{aligned}
\Phi_C(\mathbf{L}) &= \sum_{i \in [n]} \sum_{t \in T_i} \sum_{e \in \sigma_i(t)} p(i, t) \cdot w(t) \\
&\quad \cdot \left[f_e(\delta_e^{-i}(\mathbf{L}, (\mathbf{p}|_{t_i=t})) + w(t)) + f_e(w(t)) \right].
\end{aligned}$$

Similar to the proof of Theorem 3.1 we can show that any selfish step decreases the value of Φ_C . \square

This generalizes a result of Fotakis et al. [8] to the Bayesian setting. In particular our potential function reduces to their potential function if each user has only a single type.

We now turn to the model of identical links and show how a pure Bayesian Nash equilibrium can be computed in polynomial time if the type distribution is independent.

THEOREM 3.3. *Let Γ be a Bayesian routing game with independent type distribution and identical links. It is possible to compute a (normal) pure Bayesian Nash equilibrium for Γ in polynomial time.*

PROOF. Given $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ calculate for each user $i \in [n]$ its expected traffic $E(i)$. Use these expected traffics to construct a KP-game $\Gamma_{KP} = (n, m, \mathbf{1}, \{(t_1, \dots, t_n)\}, \mathbf{1})$ where $w(t_i) = E(i)$ for all $i \in [n]$. Calculate a pure Nash equilibrium $\alpha : [n] \mapsto [m]$ for Γ_{KP} in polynomial time by assigning the users in order of non-increasing user traffics to minimum load links (see [7]). Now define $\mathbf{L} = (\sigma_1, \dots, \sigma_n)$ and set $\sigma_i(t) = \alpha(i)$ for all (i, t) , $i \in [n]$, $t \in T_i$. It remains to show that \mathbf{L} is a Bayesian Nash equilibrium for Γ .

Suppose by way of contradiction that α is a pure Nash equilibrium for Γ_{KP} but that \mathbf{L} is not a Bayesian Nash equilibrium for Γ . The latter means that there is a type agent (i, t^*) of a user $i \in [n]$ who can improve by moving from

$\sigma_i(t^*) = \alpha(i)$ to another link $l \neq \sigma_i(t^*)$. By construction of \mathbf{L} , we get

$$\begin{aligned} \sum_{\substack{k \in [n], k \neq i, \\ \alpha(k) = \alpha(i)}} E(k) &= \sum_{\substack{k \in [n], \\ k \neq i}} \sum_{\substack{t \in T_k, \\ \sigma_k(t) = \sigma_i(t^*)}} p(k, t) \cdot w(t) \\ &> \sum_{\substack{k \in [n], \\ k \neq i}} \sum_{\substack{t \in T_k, \\ \sigma_k(t) = t}} p(k, t) \cdot w(t) \\ &= \sum_{\substack{k \in [n], k \neq i, \\ \alpha(k) = t}} E(k). \end{aligned}$$

Therefore, in Γ_{KP} , user i can decrease its individual cost by switching from link $\alpha(i)$ to link l . This is a contradiction to the initial assumption that α is a Nash equilibrium for Γ_{KP} . \square

The algorithm sketched in the proof of Theorem 3.3 cannot be used to compute pure Bayesian Nash equilibria for games with related links or for games with correlated type distribution. One reason is that it always computes a *normal* Bayesian Nash equilibrium. It is possible to construct a game with related links and independent type distribution for which a normal Bayesian Nash equilibrium does not exist. We now present a Bayesian routing game Γ with correlated type distribution and identical links for that a normal Bayesian Nash equilibrium does not exist. Set $\Gamma = (3, 2, \mathbf{1}, T_1 \times T_2 \times T_3, \mathbf{p})$ where the type sets are $T_1 = \{t_1, t'_1\}$, $T_2 = \{t_2, t'_2\}$ and $T_3 = \{t_3, t'_3\}$. The types are of traffic $w(t_1) = w(t'_1) = 1$, $w(t_2) = w(t_3) = 10$ and $w(t'_2) = w(t'_3) = 0$. The correlated distribution \mathbf{p} is given by $p(t_1, t'_2, t_3) = p(t'_1, t_2, t'_3) = \frac{1}{20}$ and $p(t_1, t_2, t_3) = \frac{9}{10}$. Observe that there is no normal equilibrium assigning the users 2 and 3 to the same link. If they are on different links the type t_1 has an incentive to deviate from the link user 3 is assigned to whereas type t'_1 has an incentive to deviate from the link user 2 is assigned to.

4. PROPERTIES OF FULLY MIXED BAYESIAN NASH EQUILIBRIA

In this section we study properties of fully mixed Bayesian Nash equilibria for identical links. We start by showing that the individual costs of the users are maximized in a fully mixed Bayesian Nash equilibrium and we give a simple expression for the individual cost.

THEOREM 4.1. *Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ be a Bayesian routing game with identical links, let \mathbf{F} be a fully mixed Bayesian Nash equilibrium and let \mathbf{Q} be any Bayesian Nash equilibrium for Γ . Then for all $i \in [n]$,*

- (a) $u_i(\mathbf{Q}, \mathbf{p}) \leq u_i(\mathbf{F}, \mathbf{p})$, and
- (b) $u_i(\mathbf{F}, \mathbf{p}) = E(i) + \frac{1}{m} \sum_{s \in [n], s \neq i} E(s)$.

PROOF. For any pair of users $i, s \in [n]$ and any type $t \in T_i$ define

$$E(s|t_i = t) = \sum_{\substack{(t_1, \dots, t_n) \in T \\ t_i = t}} p(t_1, \dots, t_n | t_i = t) w(t_s)$$

as the expected load of user s , given that user i has type t .

We first prove that $u_i(\mathbf{Q}, \mathbf{p}) \leq u_i(\mathbf{F}, \mathbf{p})$ for all $i \in [n]$. Let \mathbf{Q} be a Bayesian Nash equilibrium, let (i, t) be any type agent of user i and let $k \in \text{support}(t)$. Then

$$v_{(i,t)}(\mathbf{Q}, \mathbf{p}) = w(t) + \delta_k^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)),$$

and

$$\delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \geq \delta_k^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t))$$

for all $j \in [m]$. Since

$$\sum_{j \in [m]} \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) = \sum_{s \in [n], s \neq i} E(s|t_i = t),$$

it follows that $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ is maximized if

$$\delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) = \delta_k^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t))$$

for all $j \in [m]$. This holds for any fully mixed Nash equilibrium \mathbf{F} , so $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ is maximized in \mathbf{F} for all $i \in [n], t \in T_i$. Thus,

$$u_i(\mathbf{Q}, \mathbf{p}) = \sum_{t \in T_i} p(i, t) v_{(i,t)}(\mathbf{Q}, \mathbf{p}) \leq u_i(\mathbf{F}, \mathbf{p})$$

for all $i \in [n]$. This proves the first part of the lemma.

Now, we show that $u_i(\mathbf{F}, \mathbf{p}) = E(i) + \frac{1}{m} \sum_{s \in [n], s \neq i} E(s)$ holds for all $i \in [n]$. It is

$$\begin{aligned} u_i(\mathbf{F}, \mathbf{p}) &= \sum_{t \in T_i} p(i, t) \cdot v_{(i,t)}(\mathbf{F}, \mathbf{p}) \\ &= \sum_{t \in T_i} p(i, t) \left[w(t) + \delta_k^{-i}(\mathbf{F}, (\mathbf{p}|t_i = t)) \right], \end{aligned}$$

for all $k \in [m]$. Thus

$$\begin{aligned} u_i(\mathbf{F}, \mathbf{p}) &= E(i) + \sum_{t \in T_i} p(i, t) \frac{\sum_{s \in [n], s \neq i} E(s|t_i = t)}{m} \\ &= E(i) + \frac{1}{m} \sum_{s \in [n], s \neq i} \sum_{t \in T_i} p(i, t) E(s|t_i = t) \\ &= E(i) + \frac{1}{m} \sum_{s \in [n], s \neq i} E(s), \end{aligned}$$

as needed. \square

We proceed by defining a fully mixed strategy profile \mathbf{F}^* .

DEFINITION 4.1. *Define the standard fully mixed strategy profile \mathbf{F}^* as the fully mixed strategy profile that assigns every type agent to every link with probability $\frac{1}{m}$.*

It is easy to see that for any Bayesian routing game Γ with identical links the standard fully mixed strategy profile is a Bayesian Nash equilibrium. This fact was for the special case of KP-games stated in [20]. We call a fully mixed strategy profile \mathbf{F}^* *standard fully mixed Bayesian Nash equilibrium*.

There are Bayesian routing games for which the standard fully mixed Bayesian Nash equilibrium is the unique fully mixed Bayesian Nash equilibrium. For instance, let Γ be a Bayesian routing game with correlated type distribution \mathbf{p} . Let \mathbf{p} be such that for each $t \in T_i, i \in [n]$ there is at most one type profile $(t_1, \dots, t, \dots, t_n) \in T$ such that $p(t_1, \dots, t, \dots, t_n) > 0$. Such a game defines a set of disjoint KP-games. It was shown in [20] that the fully mixed Nash equilibrium for a KP-game is unique and has probabilities $\frac{1}{m}$ for all users and links. This implies that for Γ only the standard fully mixed Bayesian Nash equilibrium exists.

In general there exists more than one fully mixed Bayesian Nash equilibrium. In the remainder of this section we study

the structure of fully mixed Bayesian Nash equilibria for Bayesian routing games with independent type distribution and identical links. Lemma 4.2 presents an exact characterization of fully mixed Bayesian Nash equilibria in type agent representation. In particular Lemma 4.2 shows, how the traffic must be distributed to the links. Lemma 4.3 uses this result to show the dimension of the space of fully mixed Bayesian Nash equilibria.

LEMMA 4.2. *Let Γ be a Bayesian routing game with independent type distribution and identical links. Let \mathbf{R} be a fully mixed strategy profile for Γ in type agent representation. Then \mathbf{R} is a fully mixed Bayesian Nash equilibrium if and only if $\frac{1}{m} \cdot E(i) = \sum_{t \in T_i} r(i, t, j) \cdot p(i, t) \cdot w(t)$ for all $i \in [n]$ and $j \in [m]$.*

LEMMA 4.3. *Let Γ be a Bayesian routing game with independent type distribution and identical links. Define τ as the total number of type agents. The dimension of the space of fully mixed Bayesian Nash equilibria for Γ in type agent representation is $(\tau - n)(m - 1)$.*

PROOF. Let \mathbf{R} be a fully mixed Bayesian Nash equilibrium in type agent representation. Define $\tau_i = |T_i|$ for all $i \in [n]$. Then $\tau = \sum_{i \in [n]} \tau_i$. By Lemma 4.2, we know that \mathbf{R} is a fully mixed Bayesian Nash equilibrium if and only if $\sum_{t \in T_i} r(i, t, j) \cdot p(i, t) \cdot w(t) = \frac{1}{m} \cdot E(i)$ for all $i \in [n]$ and $j \in [m]$.

For each user $i \in [n]$ the equations

$$\begin{aligned} (1) \quad & r(i, t, j) > 0 && \forall i \in [n], \forall t \in T_i, \forall j \in [m] \\ (2) \quad & \sum_{j=1}^m r(i, t, j) = 1 && \forall i \in [n], \forall t \in T_i \\ (3) \quad & \sum_{t \in T_i} r(i, t, j) \cdot p(i, t) \cdot w(t) = \frac{E(i)}{m} && \forall i \in [n] \end{aligned}$$

have a solution space of dimension $(\tau_i - 1)(m - 1)$. Summing up over all $i \in [n]$ proves the assumption. \square

5. SOCIAL COST AND COORDINATION RATIO

5.1 Social Cost as Expected Maximum Congestion

In this section we study social cost as the expected maximum congestion on any link, which is a measure for the welfare of the system. For the special case of KP-games this social cost measure was introduced in [16] and asymptotic tight bounds on coordination ratio were given in [2, 15]. Their techniques use Chernoff bounds to show that for identical links the quotient between the expected maximum load and the maximum expected load on a link is at most $\mathcal{O}\left(\frac{\log m}{\log \log m}\right)$. These techniques cannot be applied for our Bayesian routing game as the following Lemma shows.

LEMMA 5.1. *For every $\epsilon > 0$, there exists a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with identical links and independent type distribution, and a pure Bayesian Nash equilibrium \mathbf{L} which has optimum social cost, such that*

$$\frac{\text{OPT}_{\text{MSP}}(\Gamma)}{\delta_j(\mathbf{L}, \mathbf{p})} \geq \frac{m}{1 + \epsilon}, \text{ for all } j \in [m].$$

PROOF. Define $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with independent type distribution, where $n = m$, $T_i = \{t_i, t'_i\}$, $w(t_i) = 0$, $w(t'_i) = a$, $p(i, t_i) = 1 - \frac{1}{a}$ and $p(i, t'_i) = \frac{1}{a}$ for all $i \in [n]$. Let \mathbf{L} be the pure Bayesian Nash equilibrium that maps both types of user i to link i . It is easy to see that $\text{OPT}_{\text{MSP}}(\Gamma) = \text{SC}_{\text{MSP}}(\mathbf{L}, \Gamma)$. Clearly, $\delta_j(\mathbf{L}, \mathbf{p}) = 1$ for all $j \in [m]$. On the other hand, $\text{OPT}_{\text{MSP}}(\Gamma) = (1 - (1 - \frac{1}{a})^m) a$. Now,

$$\lim_{a \rightarrow \infty} \left(1 - (1 - \frac{1}{a})^m\right) a = m.$$

The claim follows. \square

We now turn our attention to the standard fully mixed Bayesian Nash equilibrium on identical links. For KP-games this is the only fully mixed equilibrium. Gairing et al. [11] conjectured that the fully mixed Nash equilibrium has worst social cost in the KP-model. Thus, the standard fully mixed Bayesian equilibrium is a candidate for a worst case example. The next lemma shows, that the coordination ratio of the standard fully mixed Nash equilibrium does not increase if we have incomplete information.

THEOREM 5.2. *Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ be a Bayesian routing game with identical links and let \mathbf{F}^* be the standard fully mixed Bayesian Nash equilibrium. Then*

$$\frac{\text{SC}_{\text{MSP}}(\mathbf{F}^*, \Gamma)}{\text{OPT}_{\text{MSP}}(\Gamma)} = \mathcal{O}\left(\frac{\log m}{\log \log m}\right).$$

PROOF. Consider an arbitrary type profile $t = (t_1, \dots, t_n)$. Given t we define the game $\Gamma_{\text{KP}}(t) = (n, m, \mathbf{1}, \{(t_1, \dots, t_n)\}, \mathbf{1})$. We consider the fully mixed Nash equilibrium \mathbf{Q}^* for $\Gamma_{\text{KP}}(t)$. In \mathbf{Q}^* each user is assigned to each link with probability $1/m$ (see [20]). According to [2] and [20] it holds that $\frac{\text{SC}_{\text{MSP}}(\mathbf{Q}^*, \Gamma_{\text{KP}}(t))}{\text{OPT}_{\text{MSP}}(\Gamma_{\text{KP}}(t))} = \mathcal{O}\left(\frac{\log m}{\log \log m}\right)$. Thus,

$$\begin{aligned} \text{SC}_{\text{MSP}}(\mathbf{F}^*, \Gamma) &= \sum_{t \in T} p(t) \cdot \text{SC}_{\text{MSP}}(\mathbf{Q}^*, \Gamma_{\text{KP}}(t)) \\ &= \sum_{t \in T} p(t) \cdot \text{OPT}_{\text{MSP}}(\Gamma_{\text{KP}}(t)) \cdot \mathcal{O}\left(\frac{\log m}{\log \log m}\right) \\ &\leq \text{OPT}_{\text{MSP}}(\Gamma) \cdot \mathcal{O}\left(\frac{\log m}{\log \log m}\right). \end{aligned}$$

\square

Since in general there is more than one fully mixed Bayesian Nash equilibrium, the natural question arises, whether they have all the same social cost. As we see now, this is not the case.

LEMMA 5.3. *There exists a Bayesian routing game Γ with identical links and a fully mixed Bayesian Nash equilibrium \mathbf{F} such that*

$$\text{SC}_{\text{MSP}}(\mathbf{F}, \Gamma) > \text{SC}_{\text{MSP}}(\mathbf{F}^*, \Gamma).$$

PROOF. Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with $n = 2$, $m = 3$ and $T_i = \{t_i, t'_i\}$ with $w(t_i) = 2$, $w(t'_i) = 1$ for all $i \in \{1, 2\}$. Furthermore, let $p(i, t_i) = p(i, t'_i) = \frac{1}{2}$ for all $i \in \{1, 2\}$. \mathbf{F}^* assigns each type to each link with a probability of $\frac{1}{3}$. A simple calculation shows that $\text{SC}_{\text{MSP}}(\mathbf{F}^*, \Gamma) = \frac{13}{6}$. The fully mixed Bayesian Nash equilibrium \mathbf{F}_δ , $0 < \delta < \frac{1}{4}$, assigns all types of traffic 1 to link 1 with a probability of $1 - 4\delta$, to link

2 with a probability of 2δ and to link 3 with a probability of 2δ . Each type of traffic 2 is assigned to link 1 with a probability of 2δ , to link 2 with a probability of $\frac{1}{2} - \delta$ and to link 3 with a probability of $\frac{1}{2} - \delta$. Now, for each $\epsilon > 0$ there is a $\delta > 0$ such that $\text{SC}_{\text{MSP}}(\mathbf{F}_\delta, \Gamma) > \frac{9}{4} - \epsilon$. \square

It is well known (see [19]) that mixed Nash equilibria in games with complete information can be viewed as pure Bayesian Nash equilibria in a Bayesian game, where for each player all its types are identical. The following definition and theorem applies this result to Bayesian routing games.

DEFINITION 5.1. *A KP-like game is a Bayesian routing game with an independent type distribution such that $w(t) = w(t')$ for all types $t, t' \in T_i, i \in [n]$.*

THEOREM 5.4. *Let $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ be any KP-like game and let \mathbf{L} be a pure Bayesian Nash equilibrium for Γ . Then*

- (a) $\frac{\text{SC}_{\text{MSP}}(\mathbf{L}, \Gamma)}{\text{OPT}_{\text{MSP}}(\Gamma)} = \mathcal{O}\left(\frac{\log m}{\log \log m}\right)$ if Γ has identical links,
(b) $\frac{\text{SC}_{\text{MSP}}(\mathbf{L}, \Gamma)}{\text{OPT}_{\text{MSP}}(\Gamma)} = \mathcal{O}\left(\frac{\log m}{\log \log \log m}\right)$ if Γ has related links,

and both bounds are asymptotically tight.

PROOF. The proof uses a well known construction (see [19]). The upper bounds follow from the corresponding upper bounds on the coordination ratio for KP-games [2, 15]. \square

We close this section by giving a lower bound on coordination ratio for normal pure Bayesian Nash equilibria.

THEOREM 5.5. *There exists a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with identical links and a normal pure Bayesian Nash equilibrium \mathbf{L} such that*

$$\frac{\text{SC}_{\text{MSP}}(\mathbf{L}, \Gamma)}{\text{OPT}_{\text{MSP}}(\Gamma)} = \Omega\left(\frac{\log m}{\log \log m}\right).$$

PROOF. Let $m \in \mathbb{N}$ be a number such that $\sqrt{m} \in \mathbb{N}$. We consider a game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with identical links and an independent type distribution \mathbf{p} . There are two classes of users \mathcal{U}_1 and \mathcal{U}_2 . In class \mathcal{U}_1 we have m users each having two types: one of traffic 1 with probability $\frac{1}{\sqrt{m}}$ and one of traffic 0 with probability $1 - \frac{1}{\sqrt{m}}$. In class \mathcal{U}_2 there are $(m - \sqrt{m})\sqrt{m}$ users each with a single type of traffic $\frac{1}{\sqrt{m}}$.

Consider the pure strategy profile that assigns to each link one user from \mathcal{U}_1 and $\sqrt{m} - 1$ users from \mathcal{U}_2 . By analyzing the social cost of this vector we get $\text{OPT}_{\text{MSP}}(\Gamma) \leq 2$. Now consider the normal Bayesian Nash equilibrium \mathbf{L} where \sqrt{m} users from \mathcal{U}_1 are assigned to each link $j \in [\sqrt{m}]$ and \sqrt{m} users from \mathcal{U}_2 to each of the remaining $m - \sqrt{m}$ links.

To show a lower bound on $\text{SC}_{\text{MSP}}(\mathbf{L}, \Gamma)$ we consider any link $j \in [\sqrt{m}]$. The load, say X_j , on link j is defined by a sum of \sqrt{m} independent random variables with mean 1. Thus X_j is Poisson distributed for $\sqrt{m} \rightarrow \infty$ (see e.g. [25]). It holds $P(X_j \geq k) > P(X_j = k) = \frac{1}{e \cdot k!}$ and therefore $P(X_j < k) \leq 1 - \frac{1}{e \cdot k!}$. Now,

$$P(X_1 < k \wedge \dots \wedge X_{\sqrt{m}} < k) \leq \left(1 - \frac{1}{e \cdot k!}\right)^{\sqrt{m}} \leq e^{-\frac{1}{e \cdot k!} \cdot \sqrt{m}}$$

and $e^{-\frac{1}{e \cdot k!} \cdot \sqrt{m}} \leq \frac{1}{m}$ for $k \in \Omega\left(\frac{\log m}{\log \log m}\right)$. Thus,

$$\text{SC}_{\text{MSP}}(\mathbf{L}, \Gamma) \geq (1 - \frac{1}{m}) \cdot k = \Omega\left(\frac{\log m}{\log \log m}\right).$$

\square

5.2 Social Cost as Sum of Individual Costs

In this section, we study the coordination ratio for social cost as the sum of individual cost, which is a measure of average user welfare. In Theorem 5.6, we show that here fully mixed Bayesian Nash equilibria have worst social cost. This result is then used to prove an asymptotic tight bound on coordination ratio (Theorem 5.7).

The next theorem follows immediately from Theorem 4.1 and the definition of social cost.

THEOREM 5.6. *Let Γ be a Bayesian routing game with identical links, let \mathbf{F} be a fully mixed Bayesian Nash equilibrium and let \mathbf{Q} be any Bayesian Nash equilibrium for Γ . Then $\text{SC}_{\text{SUM}}(\mathbf{Q}, \Gamma) \leq \text{SC}_{\text{SUM}}(\mathbf{F}, \Gamma)$.*

THEOREM 5.7. *Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ be a Bayesian routing game with identical links and let \mathbf{Q} be any Bayesian Nash equilibrium. Then*

$$\frac{\text{SC}_{\text{SUM}}(\mathbf{Q}, \Gamma)}{\text{OPT}_{\text{SUM}}(\Gamma)} \leq \frac{m + n - 1}{m},$$

and this bound is tight up to a factor of $(1 + \epsilon)$ for any $\epsilon > 0$, even if Γ is a KP-game.

PROOF. By Theorem 5.6 it suffices to show the claim for a fully mixed Nash equilibrium \mathbf{F} . By Theorem 4.1 we have $u_i(\mathbf{F}, \mathbf{p}) = E(i) + \frac{1}{m} \sum_{s \in [n], s \neq i} E(s)$. Now

$$\begin{aligned} \text{SC}_{\text{SUM}}(\mathbf{F}, \Gamma) &= \sum_{i \in [n]} u_i(\mathbf{F}, \mathbf{p}) \\ &= \sum_{i \in [n]} \left[E(i) + \frac{1}{m} \sum_{s \in [n], s \neq i} E(s) \right] \\ &= E + \frac{n-1}{m} E \\ &= \frac{m+n-1}{m} E. \end{aligned}$$

On the other hand, since $u_i(\mathbf{Q}, \mathbf{p}) \geq E(i)$ for any user $i \in [n]$, any strategy profile \mathbf{Q} and any probability measure \mathbf{p} , we have

$$\text{OPT}_{\text{SUM}}(\Gamma) \geq \sum_{i \in [n]} E(i) = E.$$

The upper bound follows.

We now prove that this upper bound is tight. To do so, we show that for any $\epsilon > 0$, any number of users n and any number of links $m \geq 2$, there is a game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ such that $\text{OPT}_{\text{SUM}}(\Gamma) \leq (1 + \epsilon) \cdot E$. This suffices since $\text{SC}_{\text{SUM}}(\mathbf{F}^*, \Gamma) = \frac{m+n-1}{m} E$. Let Γ be a Bayesian routing game where every user $i \in [n]$ has only a single type t_i . First, notice that if $n \leq m$ then we can assign each user to a separate link which yields $\text{OPT}_{\text{SUM}}(\Gamma) = E$. Now, let $n > m$ and define the following game Γ : There are two sets of users $\mathcal{U}_1, \mathcal{U}_2$. The set \mathcal{U}_1 consists of $n - m + 1$ users with $w(t_i) = 1$ for all $i \in \mathcal{U}_1$ and \mathcal{U}_2 consists of $m - 1$ users with $w(t_i) = k$ for all $i \in \mathcal{U}_2$. Let \mathbf{L} be the pure strategy profile, that assigns all users from \mathcal{U}_1 to link m and each of the $m - 1$ users from \mathcal{U}_2 separately to a link from $[m - 1]$. It is

$$E = n - m + 1 + (m - 1)k,$$

and thus

$$\begin{aligned} \text{OPT}_{\text{SUM}}(\Gamma) &\leq \text{SC}_{\text{SUM}}(\mathbf{L}, \Gamma) = (n - m + 1)^2 + (m - 1)k \\ &= \frac{(n - m + 1)^2 + (m - 1)k}{n - m + 1 + (m - 1)k} E \\ &= \left(1 + \frac{(n - m)(n - m + 1)}{n + (m - 1)(k - 1)}\right) \cdot E. \end{aligned}$$

Clearly, for any $\varepsilon > 0$ there exists a k such that $\frac{(n - m)(n - m + 1)}{n + (m - 1)(k - 1)} \leq \varepsilon$, which completes the tightness proof. \square

5.3 Social Cost as Maximum Individual Cost

Now, we analyze the coordination ratio for social cost as the maximum of individual cost. Theorem 5.8 shows, that all fully mixed Bayesian Nash equilibria have worst social cost. In Theorem 5.9, we apply this result to proof an asymptotic tight bound on coordination ratio for Bayesian routing games and KP-games.

The next theorem follows immediately from Theorem 4.1 and the definition of social cost.

THEOREM 5.8. *Let Γ be a Bayesian routing game with identical links, let \mathbf{F} be a fully mixed Bayesian Nash equilibrium and let \mathbf{Q} be any Bayesian Nash equilibrium for Γ . Then $\text{SC}_{\text{MAX}}(\mathbf{Q}, \Gamma) \leq \text{SC}_{\text{MAX}}(\mathbf{F}, \Gamma)$.*

THEOREM 5.9. *Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ be a Bayesian routing game with identical links and let \mathbf{Q} be any Bayesian Nash equilibrium. Then*

$$\begin{aligned} (a) \quad \frac{\text{SC}_{\text{MAX}}(\mathbf{Q}, \Gamma)}{\text{OPT}_{\text{MAX}}(\Gamma)} &\leq \frac{m + n - 1}{m} \text{ and} \\ (b) \quad \frac{\text{SC}_{\text{MAX}}(\mathbf{Q}, \Gamma)}{\text{OPT}_{\text{MAX}}(\Gamma)} &\leq 2 - \frac{1}{m}, \text{ if } \Gamma \text{ is a KP-game.} \end{aligned}$$

The bound from (a) is tight up to a factor of $(1 + \varepsilon)$ for any $\varepsilon > 0$ and the bound from (b) is tight.

PROOF. We first show an upper bound on $\text{SC}_{\text{MAX}}(\mathbf{Q}, \Gamma)$ that holds for both cases. Then we consider case (a) and (b) separately. In each case we proof lower bounds on $\text{OPT}_{\text{MAX}}(\Gamma)$ and show tightness.

Let \mathbf{F} be a fully mixed Bayesian Nash equilibrium for Γ . By Theorem 4.1,

$$u_i(\mathbf{Q}, \mathbf{p}) \leq u_i(\mathbf{F}, \mathbf{p}) = E(i) + \frac{1}{m} \sum_{s \neq i} E(s) = \frac{E}{m} + \frac{m - 1}{m} E(i), \quad (2)$$

for any user $i \in [n]$.

Case (a):

Upper bound: Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$. For any strategy profile \mathbf{Q}' and any user $i \in [n]$, it holds that $u_i(\mathbf{Q}', \mathbf{p}) \geq E(i)$ and therefore $\sum_{i \in [n]} u_i(\mathbf{Q}', \mathbf{p}) \geq E$. This implies

$$\text{OPT}_{\text{MAX}}(\Gamma) \geq \frac{E}{n}. \quad (3)$$

Clearly $\text{OPT}_{\text{MAX}}(\Gamma) \geq E(i)$ for all $i \in [n]$. For any user $i \in [n]$ we get by Equation 2, that

$$\begin{aligned} u_i(\mathbf{Q}, \mathbf{p}) &\leq \frac{E}{m} + \frac{m - 1}{m} \cdot \text{OPT}_{\text{MAX}}(\Gamma) \\ &\stackrel{(3)}{\leq} \frac{n}{m} \cdot \text{OPT}_{\text{MAX}}(\Gamma) + \frac{m - 1}{m} \cdot \text{OPT}_{\text{MAX}}(\Gamma) \\ &= \frac{m + n - 1}{m} \cdot \text{OPT}_{\text{MAX}}(\Gamma). \end{aligned}$$

The upper bound in (a) follows.

Lower bound: We now proof that this upper bound is tight up to a factor of $(1 + \varepsilon)$. Consider for arbitrary $k, a, r \in \mathbb{N}$ the game $\Gamma_{k,a,r} = (n, m, \mathbf{1}, T, \mathbf{p})$ with independent type distribution and $n = k \cdot (m - 1)$. Each user $i \in [n]$ has 2 types $T_i = \{t_i, t'_i\}$ with traffics $w(t_i) = 1$, $w(t'_i) = a \cdot r$ and probabilities $p(i, t_i) = 1 - \frac{1}{a}$, $p(i, t'_i) = \frac{1}{a}$. It is $E(i) = r + 1 - \frac{1}{a}$ for all users $i \in [n]$.

Define a pure strategy profile \mathbf{L} that assigns all types t'_i , $i \in [n]$, of traffic 1 to link m . The types t_i , $i \in [n]$, are evenly distributed among the links in $[m - 1]$, that is, \mathbf{L} assigns exactly k of these types to each link in $[m - 1]$. Now for any user i and any $\varepsilon > 0$ there is a sufficient large a , such that

$$\begin{aligned} u_i(\mathbf{L}, \mathbf{p}) &= \left(1 - \frac{1}{a}\right) \cdot \left(1 + (k - 1) \cdot \left(1 - \frac{1}{a}\right)\right) \\ &\quad + \frac{1}{a} \cdot ((n - 1)r + r \cdot a) \\ &\leq \frac{k + r}{1 + \varepsilon}. \end{aligned}$$

On the other hand for any fully mixed Bayesian Nash equilibrium \mathbf{F} and any user $i \in [n]$, we have by Theorem 4.1

$$\begin{aligned} u_i(\mathbf{F}, \mathbf{p}) &= \left(1 + \frac{n - 1}{m}\right) \cdot E(i) \\ &= \frac{m + n - 1}{m} \cdot \left(r + 1 - \frac{1}{a}\right), \end{aligned}$$

and thus for sufficient large r

$$\begin{aligned} \frac{u_i(\mathbf{F}, \mathbf{p})}{u_i(\mathbf{L}, \mathbf{p})} &\geq \frac{(r + 1 - \frac{1}{a})(1 + \varepsilon)}{(k + r)} \cdot \frac{m + n - 1}{m} \\ &\geq \frac{m + n - 1}{m} \cdot (1 + \varepsilon). \end{aligned}$$

This proves that the bound in (a) is tight up to a factor of $(1 + \varepsilon)$.

Case (b):

Upper bound: Let $\Gamma_{\text{KP}} = (n, m, \mathbf{1}, \{(t_1, \dots, t_n)\}, 1)$ be any KP-game. Here $E(i) = w(t_i)$ for all $i \in [n]$. Clearly, $\text{OPT}_{\text{MAX}}(\Gamma_{\text{KP}}) \geq E(i)$ for all $i \in [n]$ and $\text{OPT}_{\text{MAX}}(\Gamma_{\text{KP}}) \geq \frac{E}{m}$. We get with Equation 2, that

$$\begin{aligned} u_i(\mathbf{Q}, \mathbf{p}) &\leq \frac{E}{m} + \frac{m - 1}{m} E(i) \\ &\leq \text{OPT}_{\text{MAX}}(\Gamma_{\text{KP}}) + \frac{m - 1}{m} \text{OPT}_{\text{MAX}}(\Gamma_{\text{KP}}) \\ &= \left(2 - \frac{1}{m}\right) \text{OPT}_{\text{MAX}}(\Gamma_{\text{KP}}). \end{aligned}$$

The upper bound in (b) follows.

Lower bound: To see tightness, consider any KP-game $\Gamma_{\text{KP}} = (n, m, \mathbf{1}, \{(t_1, \dots, t_n)\}, 1)$ with $n = m$ users and $w(t_1) = \dots = w(t_n) = 1$. It is $\text{OPT}_{\text{MAX}}(\Gamma_{\text{KP}}) = 1$. Now, for the fully mixed Nash equilibrium \mathbf{F} and any user $i \in [n]$, by Equation 2,

$$u_i(\mathbf{F}, \mathbf{p}) = \frac{E}{m} + \frac{m - 1}{m} E(i) = \left(2 - \frac{1}{m}\right) \cdot \text{OPT}_{\text{MAX}}(\Gamma_{\text{KP}}).$$

This proves that the lower bound in (b) is tight. \square

Acknowledgment. We would like to thank Thomas Lücking and Marios Mavronicolas for many fruitful discussions and helpful comments.

6. REFERENCES

- [1] R. Beier, A. Czumaj, P. Krysta, and B. Vöcking. Computing Equilibria for Congestion Games with (Im)perfect Information. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'04)*, pages 746–755, 2004.
- [2] A. Czumaj and B. Vöcking. Tight Bounds for Worst-Case Equilibria. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*, pages 413–420, 2002. Also accepted to *Journal of Algorithms* as Special Issue of SODA'02.
- [3] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC'04)*, pages 604–612, 2004.
- [4] G. Facchini, F. van Megan, P. Borm, and S. Tijs. Congestion Models and Weighted Bayesian Potential Games. *Theory and Decision*, 42:193–206, 1997.
- [5] 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 (ICALP'03)*, LNCS 2719, pages 514–526, 2003.
- [6] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish Routing in Non-Cooperative Networks: A Survey. *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS'03)*, LNCS 2747, pages 21–45, 2003.
- [7] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. *Proceedings of the 29th International Colloquium on Automata, Languages, and Programming (ICALP'02)*, LNCS 2380, pages 123–134, 2002.
- [8] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming (ICALP'04)*, LNCS 3142, pages 593–605, 2004.
- [9] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing Nash Equilibria for Scheduling on Restricted Parallel Links. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC'04)*, pages 613–622, 2004.
- [10] 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 (ICALP'04)*, LNCS 3142, pages 645–657, 2004.
- [11] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Extreme Nash Equilibria. *Proceedings of the 8th Italian Conference on Theoretical Computer Science (ICTCS'03)*, LNCS 2841, pages 1–20, 2003. Also accepted to *Theoretical Computer Science*, Special Issue on *Game Theory Meets Theoretical Computer Science*.
- [12] R. L. Graham. Bounds on Multiprocessing Timing Anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.
- [13] J. C. Harsanyi. Games with Incomplete Information Played by Bayesian Players, I, II, III. *Management Science*, 14:159–182, 320–332, 468–502, 1967.
- [14] J. C. Harsanyi. Games with Randomly Disturbed Payoffs. *International Journal on Game Theory*, 21:1–23, 1973.
- [15] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.
- [16] E. Koutsoupias and C. Papadimitriou. Worst-Case Equilibria. *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, LNCS 1563, pages 404–413, 1999.
- [17] 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 (STACS'04)*, LNCS 2996, pages 547–558, 2004.
- [18] 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 (MFCS'03)*, LNCS 2747, pages 551–561, 2003.
- [19] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [20] M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 510–519, 2001.
- [21] I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
- [22] D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [23] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.
- [24] C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 749–753, 2001.
- [25] A. Papoulis. *Probability, Random variables, and Stochastic processes*. McGraw-Hill Book company, 1984.
- [26] R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [27] M. Voorneveld, P Borm, F. van Megan, S. Tijs, and G. Facchini. Congestion Games and Potentials Reconsidered. *International Game Theory Review*, 1:283–299, 1999.