Selfish Routing with Incomplete Information*

Martin Gairing Burkhard Monien Karsten Tiemann†

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

(July 6, 2006)

*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). A preliminary version of this paper appeared in the Proceedings of the 17th ACM Symposium on Parallelism in Algorithms and Architectures, pp. 203–212, July 2005.

†International Graduate School of Dynamic Intelligent Systems, University of Paderborn, Germany

Abstract

In his seminal work, Harsanyi [19] introduced an elegant approach to study non-cooperative games with *incomplete information*. In our work, we use this approach to define a new selfish routing game with incomplete information that we call *Bayesian routing game*. Here, each of n selfish *users* wishes to assign its *traffic* to one of m parallel *links*. However, users do not know each other's traffic. Following Harsanyi's approach, we introduce, for each user, a set of possible *types*. In our model, each type of a user corresponds to some traffic and the players' uncertainty about each other's traffic is described by a probability distribution over all possible *type profiles*.

We present a comprehensive collection of results about our Bayesian routing game. Our main findings are as follows:

- Using a *potential function*, we prove that every Bayesian routing game has a *pure Bayesian Nash equilibrium*. More precisely, we show this existence for a more general class of games that we call *weighted Bayesian congestion games*. For Bayesian routing games with *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 case of identical links and show that they maximize *Individual Cost*. In general, there is more than one fully mixed Bayesian Nash equilibrium. We characterize fully mixed Bayesian Nash equilibria for the case of independent type distribution.
- We conclude with bounds on *Coordination Ratio* for the case of identical links and for three different Social Cost measures: *Expected Maximum Latency*, *Sum of Individual Costs* and *Maximum Individual Cost*. For the latter two, we are able to give (asymptotically) tight bounds using the properties of fully mixed Bayesian Nash equilibria we proved.

1 Introduction

1.1 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 honored work, Harsanyi [19] introduced an elegant approach to studying non-cooperative games with *incomplete information*, where the players are uncertain about some parameters. To model such games, Harsanyi introduced the *Harsanyi transformation*, which converts a (strategic) game with incomplete information to a strategic game where players have different *types*. The type of a player represents its private information that is not common knowledge to all players. In the resulting *Bayesian game*, each player's uncertainty about each other's type is described by a probability distribution over all possible *type profiles*. Using this probability distribution, players make their decisions according to *Bayesian Decision Theory* [4]. In Bayesian Decision Theory probabilities are used as a measure of the degree of belief a person has in some proposition.

In this work, we introduce a particular selfish routing game with incomplete information that we call *Bayesian routing game*. Here, each of n selfish *users* wishes to assign its *traffic* to one of m *links*. Each link has a certain *capacity*, which specifies the rate at which the link processes traffic. In the case of *identical links*, all links have equal capacity. Link capacities vary arbitrary, in the case of *related links*. The *latency* of a link is the total traffic on the link divided by the capacity of the link. Users do not know each other's traffic. Following Harsanyi's approach, we introduce for each user a set of possible *types*. We assume that all *type sets* are finite. Each type of a user corresponds to some traffic. Furthermore, we assume that there is a joint probability distribution $\mathbf{p}$, called *type distribution*, over the set of all possible type realizations. In general, $\mathbf{p}$ can be arbitrary; however, sometimes we assume $\mathbf{p}$ to be *independent* – in that case, $\mathbf{p}$ is expressed as the product of n independent probability distributions, one for each user type set.

In a *pure strategy*, a user chooses for each of its types a particular link; so, a pure strategy is a function from the type set of a user to the set of links. 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 latency experienced by the user. Note that due to the Bayesian model, the Individual Cost in a pure strategy profile is given by the expectation over the type distribution $\mathbf{p}$. For mixed strategy profiles, the expectation is taken over both the type distribution **p** and the mixed strategies of the users.

The users neither cooperate with each other nor adhere to a global objective function, the so called *Social Cost* [22]. 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 strictly positive probability to each of its pure strategies.

If each user has only a single type, so that users are completely informed about each other's traffic, then we are in the setting of the selfish routing game with complete information introduced in the pioneering work of Koutsoupias and Papadimitriou [22]. We call it a *complete information routing game*. In this setting, Bayesian Nash equilibria become Nash equilibria. As in [22], we use the *Coordination Ratio*, or *Price of Anarchy* [30], 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.

As a generalization of our Bayesian routing games we also introduce *weighted Bayesian congestion games*. In a weighted Bayesian congestion game, the strategy set of each player is a subset of the power set of given resources. Weighted Bayesian congestion games provide us with a general framework for modeling any kind of non-cooperative resource sharing problem where the users do not know each other's traffic. A congestion game, as introduced by Rosenthal [31], is a weighted Bayesian congestion game where each user has only a single type of traffic one.

1.2 Contribution

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

(1) Existence and computational complexity of pure Bayesian Nash equilibria:
Our existence result applies for the class of *weighted Bayesian congestion games*. We define a new potential function that we use to prove that every weighted Bayesian congestion game possesses a pure Bayesian Nash equilibrium.

For the case of Bayesian routing games, identical links and independent type distributions, we show that a pure Bayesian Nash equilibrium can be computed in polynomial time. This computation is based on Graham's LPT scheduling algorithm [18]. For the case of related links and independent type distribution, and also for the case of identical links and arbitrary type distribution, the complexity of computing a pure Bayesian Nash equilibrium remains open.

(2) Properties of fully mixed Bayesian Nash equilibria:
We show that for the case 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 determine the dimension of the space of fully mixed Bayesian Nash equilibria for the case of independent type distributions.

(3) Bounds on Coordination Ratio:
We conclude with bounds on the Coordination Ratio for three different Social Cost measures and for the case of identical links.

- The *Expected Maximum Latency* 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 value of Coordination Ratio for this Social Cost measure remains open, even for the case of identical links.
- A Social Cost measure that describes average user welfare is the *Sum of Individual Costs*. In this setting, it follows that for the case of identical links, each fully mixed Bayesian Nash equilibrium has maximum Social Cost. Using this fact, we prove an upper bound of $\frac{m+n-1}{m}$ on the Coordination Ratio for the case of identical links. We prove that this bound is asymptotically tight, already for complete information routing games.
- We also study Social Cost as *Maximum Individual Cost*. For identical links, we show asymptotically tight upper bounds on Coordination Ratio of $\frac{m+n-1}{m}$ for Bayesian routing games and of $2 - \frac{1}{m}$ for complete information routing games.

To the best of our knowledge, this is the first time that mixed Bayesian Nash equilibria are studied in combination with Social Cost.

1.3 Related Work

Congestion games were introduced by Rosenthal [31] and extensively studied afterwards (see, e.g., [1, 5, 7, 8, 13, 27, 28, 32] and [16] for a recent survey). In Rosenthal's model, each player has complete information and each pure strategy of it is a subset of *resources*. Resource *cost functions* can be arbitrary, but they only depend on the number of players sharing the same resource. Rosenthal used a *potential function* to show that such games always admit a pure Nash equilibrium. Subsequent papers [28, 32] characterize games that possess a potential function as *potential games* and show their relation to congestion games. The complexity of computing pure Nash equilibria for congestion games was studied by Fabrikant *et al.* [7]. Milchtaich [27] considers *weighted congestion games* with *player-specific* payoff functions; he also shows that these games do not admit a pure Nash equilibrium in general. Fotakis *et al.* [13] considered weighted congestion games and proved the existence of pure Nash equilibria for the case where resources have linear cost functions.

Complete information routing games on parallel links, and their Nash equilibria, were studied extensively in the last few years; see, for example, [6, 9, 12, 15, 21, 22, 24] and [10] for a survey. Graham's LPT scheduling algorithm [18] computes a pure Nash equilibrium in this setting [12].

Harsanyi developed in his pioneering work [19, 20] a framework for studying competitive situations where the players have incomplete information. For an introduction to these so-called Bayesian games, we refer to [25, 29]. Facchini *et al.* [8] considered Bayesian congestion models with players of identical weight, which have incomplete information about each other's preferences. Beier *et al.* [2] focus on a service provider congestion game with incomplete information.

The Fully Mixed Nash Equilibrium Conjecture, states that for complete information routing games the fully mixed Nash equilibrium has worst Social Cost among all Nash equilibria; it was motivated by some results in [26], explicitly formulated in [15] and further studied in [24]. For Social Cost defined as the Sum of Individual Costs, the conjecture holds [14, 23]; it was recently disproved for Social Cost defined as the Expected Maximum Latency [11].

The *Coordination Ratio,* also known as *Price of Anarchy* [30], was first introduced and studied by Koutsoupias and Papadimitriou [22]. For complete information routing games and Social Cost defined as the Expected Maximum Latency, there exist *tight* bounds of $\Theta\left(\frac{\log m}{\log\log m}\right)$ for identical links [6, 21] and $\Theta\left(\frac{\log m}{\log\log\log m}\right)$ [6] for related links. For complete information routing games and Social Cost defined as the Sum of Individual Costs, Berenbrink *et al.* [3] give a lower bound of $\frac{n}{5}$ on the Coordination Ratio. Restricting to pure Nash equilibria, they also show an upper bound that solely depends on the user traffics.

Subsequently to our work Georgiou et al. [17] introduced a routing game with incomplete information where the players have complete information about each other's traffic but only incomplete information about the latency functions in the network.

1.4 Road Map

Section 2 introduces Bayesian routing games. Pure Bayesian Nash equilibria are studied in Section 3. Fully mixed Bayesian Nash equilibria are treated in Section 4. Section 5 studies the Coordination Ratio. We conclude, in Section 6, with a summary of our results and some open problems.

2 Model

Throughout, denote for each positive integer k, $[k] = \{1, \ldots, k\}$; take that $[0] = \emptyset$.

2.1 Bayesian Routing Games

A *Bayesian routing game* is a tuple $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$. Each of n users $1, 2, \ldots, n$ wishes to assign a particular amount of traffic to one of m *links* $1, 2, \ldots, m$. Throughout, we assume that $n \geq 2$ and $m \geq 2$. Denote $\mathbf{c} = (c_1, \ldots, c_m)$, where $c_j > 0$ is the *capacity* of link $j \in [m]$. In the case of *identical links*, all capacities equal 1. In this case, we write $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$. Link capacities vary arbitrarily in the case of *related links*. For each user $i \in [n]$, there is a finite set of possible types T_i; for each type $t \in T_i$, denote by $w(t)$ the *traffic* of type t, $w(t) \geq 0$. Denote $T = T_1 \times \ldots \times T_n$, the set of all possible *type profiles*. For each user $i \in [n]$ define $\tau_i = |T_i|$ as the number of types of user i. Define $\tau = \sum_{i \in [n]} \tau_i$ as the total number of types of the users. For simplicity we assume that the traffics $(w(t_i))_{t_i \in T_i, i \in [n]}$ are encoded in T, so we do not include them in the game tuple. We use the term *type agent* (i, t) to refer to the type $t \in T_i$ of user $i \in [n]$.

There is a joint probability distribution $\mathbf{p} = (p(t_1, \ldots, t_n))_{(t_1, \ldots, t_n) \in T}$, called *type distribution*, over the set of type profiles T; thus, $\mathbf{p}$ is a function $\mathbf{p} : T \rightarrow [0, 1]$ and $\sum_{(t_1, \ldots, t_n) \in T} p(t_1, \ldots, t_n) = 1$. Denote by $p(i, t)$ the probability that user i is of type t; so,

$$p(i, t) \quad = \sum_{(t_1, \ldots, t_n) \in T : t_i = t} p(t_1, \ldots, t_n).$$

We say that $\mathbf{p}$ is *independent* if

$$p(t_1, \ldots, t_n) = \prod_{i \in [n]} p(i, t_i) \quad \text{for all} \quad (t_1, \ldots, t_n) \in T,$$

otherwise, $\mathbf{p}$ is *correlated*. By the definition of conditional probability,

$$p(t_1, \ldots, t_{k-1}, t_{k+1}, \ldots, t_n | t_k = t) \quad = \quad \frac{p(t_1, \ldots, t_{k-1}, t, t_{k+1}, \ldots, t_n)}{p(k, t)};$$

that is, the probability of a type profile $(t_1, \ldots, t_n)$ *given that* $t_k = t$ is the probability of type profile $(t_1, \ldots, t_n)$ divided by the probability that user k is of type t. Throughout this paper we only consider instances where $p(k,t) > 0$ for all users $k \in [n]$ and all types $t \in T_k$. Denote by $W(i)$ the expected traffic of user $i \in [n]$; clearly,

$$\begin{aligned} W(i) &= \sum_{(t_1,\ldots,t_n)\in T} p(t_1,\ldots,t_n) \cdot w(t_i) \\ &= \sum_{t \in T_i} p(i,t) \cdot w(t). \end{aligned}$$

Furthermore, define the *expected total traffic* as

$$W = \sum_{i\in[n]} W(i).$$

For any pair of users $i, s \in [n]$ and for any type $t \in T_i$, define $W(s|t_i = t)$ as the *conditional expected traffic* of user s, given that user i has type t; so,

$$W(s|t_i = t) = \sum_{\substack{(t_1,\ldots,t_n)\in T: \\ t_i = t}} p(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i = t) w(t_s).$$

For the case of independent type distribution we have $W(s|t_i = t) = W(s)$ for all types $t \in T_i$ of user i.

A special instance of our Bayesian routing game in which each user has only a single type is a *complete information routing game*. For such a game, we write $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{c}, T, 1)$. Here, the set T contains only one type vector t that is used with probability 1.

2.2 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 $[m]$; so σ_i is a function $\sigma_i : T_i \to [m]$. Denote as Σ_i the set of all possible pure strategies for user $i \in [n]$; denote $\Sigma = \Sigma_1 \times \ldots \times \Sigma_n$. 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; here, $q(i,\sigma_i)$ denotes the probability that user i chooses the pure strategy σ_i.

The *support* of a mixed strategy Q_i for user $i \in [n]$, denoted $\mathsf{support}_{Q_i}(i)$, is the set of links to which user i assigns at least one type $t \in T_i$ with positive probability, that is,

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

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

$$\mathsf{support}_{Q_i}(t) \;=\; \{j \in [m] \,|\, \exists \sigma_i \in \Sigma_i \text{ with } q(i, \sigma_i) > 0 \text{ and } \sigma_i(t) = j\}.$$

Note that

$$\mathsf{support}_{Q_i}(i) \;=\; \bigcup_{t \in T_i} \mathsf{support}_{Q_i}(t).$$

A *pure strategy profile* σ is an n-tuple $(\sigma_1, \ldots, \sigma_n) \in \Sigma$. Call σ *normal* if $\sigma_i(t) = \sigma_i(t')$ for all types $t, t' \in T_i$ and for all users $i \in [n]$. So, each user $i \in [n]$ does not distinguish among its types in a normal pure strategy profile.

A *mixed strategy profile* $\mathbf{Q} = (Q_1, \ldots, Q_n)$ is an n-tuple of mixed strategies. Call a mixed strategy profile $\mathbf{F} = (F_1, \ldots, F_n)$ *fully mixed* if each user assigns strictly positive probability to each of its pure strategies; that is, $q(i, \sigma_i) > 0$ for all users $i \in [n]$ and all strategies $\sigma_i \in \Sigma_i$. Notice that $\mathsf{support}_{F_i}(i) = [m]$ for all users $i \in [n]$ and $\mathsf{support}_{F_i}(t) = [m]$ for all users $i \in [n]$ and types $t \in T_i$.

2.3 Individual Costs

2.3.1 Pure Strategy Profiles

Fix any type distribution $\mathbf{p}$ and a pure strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$. The *expected load* on link $j \in [m]$, denoted $\delta_j(\sigma, \mathbf{p})$, is defined by

$$\delta_j(\sigma, \mathbf{p}) \;=\; \sum_{(t_1, \ldots, t_n) \in T} p(t_1, \ldots, t_n) \sum_{\substack{i \in [n]: \\ \sigma_i(t_i) = j}} w(t_i).$$

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

$$\delta_j^{-k}(\sigma, (\mathbf{p}|t_k = t)) \;=\; \sum_{\substack{(t_1, \ldots, t_n) \in T: \\ t_k = t}} p(t_1, \ldots, t_{k-1}, t_{k+1}, \ldots, t_n | t_k = t) \sum_{\substack{i \in [n] \setminus \{k\}: \\ \sigma_i(t_i) = j}} w(t_i).$$

Denote as $\lambda_{(i,t)}^j(\sigma, \mathbf{p})$ the Individual Cost of type agent (i, t) when its traffic is assigned to link $j \in [m]$; so,

$$\lambda_{(i,t)}^j(\sigma, \mathbf{p}) = \frac{\delta_j^{-i}(\sigma, (\mathbf{p}|t_i = t)) + w(t)}{c_j}.$$

Denote as $v_{(i,t)}(\sigma, \mathbf{p})$ the *Conditional Individual Cost* of user $i \in [n]$, given that user i is of type t; this is also the Individual Cost of type agent (i,t); so

$$v_{(i,t)}(\sigma, \mathbf{p}) = \lambda_{(i,t)}^{\sigma_i(t)}(\sigma, \mathbf{p}).$$

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

Finally, denote as $u_i(\sigma, \mathbf{p})$ the *Individual Cost* of user i; clearly,

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

2.3.2 Mixed Strategy Profiles

Fix any type distribution $\mathbf{p}$ and a mixed strategy profile $\mathbf{Q}$. The *expected load* on link $j \in [m]$, denoted $\delta_j(\mathbf{Q}, \mathbf{p})$, is defined by

$$\delta_j(\mathbf{Q}, \mathbf{p}) \quad = \quad \sum_{\sigma \in \Sigma} \prod_{i \in [n]} q(i, \sigma_i) \cdot \delta_j(\sigma, \mathbf{p})$$

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

$$\delta_j^{-k}(\mathbf{Q}, (\mathbf{p}|t_k = t)) \quad = \quad \sum_{\sigma \in \Sigma} \prod_{i \in [n]} q(i, \sigma_i) \cdot \delta_j^{-k}(\sigma, (\mathbf{p}|t_k = t)).$$

For the case of an independent type distribution $\mathbf{p}$, we get that for all types $t, t' \in T_k$, $\delta_j^{-k}(\mathbf{Q}, (\mathbf{p}|t_k = t)) = \delta_j^{-k}(\mathbf{Q}, (\mathbf{p}|t_k = t'))$. Therefore, to simplify notation, we write in this case $\delta_j^{-k}(\mathbf{Q}, \mathbf{p})$.

Denote as $\lambda_{(i,t)}^{j}(\mathbf{Q}, \mathbf{p})$ the Individual Cost of type agent (i,t) when its traffic is assigned to link $j \in [m]$; so,

$$\lambda_{(i,t)}^{j}(\mathbf{Q}, \mathbf{p}) = \frac{\delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) + w(t)}{c_j}.$$

Denote as $v_{(i,t)}(\mathbf{Q}, \mathbf{p})$ the *Conditional Individual Cost* of user $i \in [n]$, given that user i is of type t; this is also the Individual Cost of type agent (i,t); so,

$$v_{(i,t)}(\mathbf{Q}, \mathbf{p}) = \sum_{\sigma_i \in \Sigma_i} q(i, \sigma_i) \cdot \lambda_{(i,t)}^{\sigma_i(t)}(\mathbf{Q}, \mathbf{p}).$$

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.

Finally, denote as $u_i(\mathbf{Q}, \mathbf{p})$ the *Individual Cost* of user i; clearly,

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

2.4 Bayesian Nash Equilibria

A strategy profile $\mathbf{Q}$ is a Bayesian Nash equilibrium, if in $\mathbf{Q}$ no user has an incentive to deviate from its (mixed) strategy; that is, no user can possibly decrease its Individual Cost when other users are sticking to their strategies. Formally, the mixed strategy profile $\mathbf{Q} = (Q_1, \ldots, Q_n)$ is a *Bayesian Nash equilibrium* if

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

for all mixed strategy profiles $\mathbf{Q}' = (Q_1, \ldots, Q'_i, \ldots, Q_n)$ and for all users $i \in [n]$. 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 mixed strategy profiles $\mathbf{Q}' = (Q_1, \ldots, Q'_i, \ldots, Q_n)$ and for all users $i \in [n]$ and types $t \in T_i$. Note that $\mathbf{Q}$ is a Bayesian Nash equilibrium if and only if for all users $i \in [n]$ and types $t \in T_i$,

$$\begin{aligned} v_{(i,t)}(\mathbf{Q}, \mathbf{p}) &= \lambda^j_{(i,t)}(\mathbf{Q}, \mathbf{p}), \quad \text{for } j \in \mathsf{support}_{Q_i}(t), \text{ and} \\ v_{(i,t)}(\mathbf{Q}, \mathbf{p}) &\leq \lambda^j_{(i,t)}(\mathbf{Q}, \mathbf{p}), \quad \text{for } j \notin \mathsf{support}_{Q_i}(t). \end{aligned}$$

We refer to these conditions as the *Bayesian Nash equilibrium conditions*.

2.5 Social Cost and Coordination Ratio

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

- the *Expected Maximum Latency*, which is the expectation over all user choices and type profiles, of the maximum latency on a link; so

$$\begin{aligned} \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q}, \Gamma) &= \sum_{(\sigma_1,\ldots,\sigma_n)\in\Sigma} \prod_{i\in[n]} q(i,\sigma_i) \sum_{(t_1,\ldots,t_n)\in T} p(t_1,\ldots,t_n) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma_i(t_i)=j}} w(t_i) \right\} \\ &= \sum_{(t_1,\ldots,t_n)\in T} p(t_1,\ldots,t_n) \sum_{(\sigma_1,\ldots,\sigma_n)\in\Sigma} \prod_{i\in[n]} q(i,\sigma_i) \cdot \max_{j\in[m]} \left\{ \frac{1}{c_j} \sum_{\substack{i\in[n],\\ \sigma_i(t_i)=j}} w(t_i) \right\}; \end{aligned}$$

- the *Sum of Individual Costs*,

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

- the *Maximum Individual Cost*,

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

Let $* \in \{\mathsf{MSP}, \mathsf{SUM}, \mathsf{MAX}\}$. Denote the corresponding *Optimum Social Cost* by $\mathsf{OPT}_*(\Gamma) = \min_{\mathbf{Q}} \mathsf{SC}_*(\mathbf{Q}, \Gamma)$. The *Coordination Ratio* CR_* is the supremum, over all instances Γ and Bayesian Nash equilibria $\mathbf{Q}$, of the ratio $\frac{\mathsf{SC}_*(\mathbf{Q},\Gamma)}{\mathsf{OPT}_*(\Gamma)}$; that is,

$$\mathsf{CR}_* \quad = \quad \sup_{\Gamma, \mathbf{Q}} \frac{\mathsf{SC}_*(\mathbf{Q}, \Gamma)}{\mathsf{OPT}_*(\Gamma)}.$$

2.6 Weighted Bayesian Congestion Games

A generalization of the Bayesian routing game considered in this paper is the *weighted Bayesian congestion game with linear cost functions.* In a congestion game [31], 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 $g_e(x) = a_e x + b_e$. For a Bayesian congestion game, a pure strategy profile σ is defined by $\sigma = (\sigma_1, \ldots, \sigma_n)$ with $\sigma_i : T_i \rightarrow S_i$ for all $i \in [n]$. Thus a pure strategy of a user can consist of *many* resources whereas in a Bayesian routing game a pure strategy is *one* link.

For a pure strategy profile σ, the *conditional expected load* of all users $i \in [n]$ other than k, on resource $e \in [m]$ given that $t_k = t$ is then

$$\delta_e^{-k}(\sigma, (\mathbf{p}|t_k = t)) \quad = \quad \sum_{\substack{(t_1,\ldots,t_n) \in T: \\ t_k = t}} p(t_1, \ldots, t_{k-1}, t_{k+1}, \ldots, t_n | t_k = t) \sum_{\substack{i \in [n] \setminus \{k\}: \\ e \in \sigma_i(t_i)}} w(t_i);$$

whereas the Conditional Individual Cost of user i, given that user i is of type $t \in T_i$ is then defined by

$$v_{(i,t)}(\sigma, \mathbf{p}) = \sum_{e \in \sigma_i(t)} g_e(\delta_e^{-i}(\sigma, (\mathbf{p}|t_i = t)) + w(t)).$$

3 Pure Bayesian Nash Equilibria

In this section, we study pure Bayesian Nash equilibria.

3.1 Existence

We prove:

Theorem 3.1 *Every weighted Bayesian congestion game* Γ *with linear cost functions has a pure Bayesian Nash equilibrium.*

Proof: Given a pure strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$, define the function

$$\Phi(\sigma) \quad = \quad \sum_{i \in [n]} \sum_{t \in T_i} \sum_{e \in \sigma_i(t)} p(i,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-i}(\sigma, (\mathbf{p}|t_i = t)) + w(t)) + g_e(w(t)) \right].$$

We will prove that any unilateral strategy change of a type agent that decreases its Individual Cost also decreases the value of the function Φ.

Given a pure strategy profile σ, define for every user $r \in [n]$ and type $t \in T_r$,

$$\Phi_{(r,t)}(\sigma) \quad = \quad \sum_{e \in \sigma_r(t)} p(r,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = t)) + w(t)) + g_e(w(t)) \right];$$

and for every resource $e \in [m]$ and user $r \in [n]$,

$$\Phi_e^{-r}(\sigma) = \sum_{i \in [n] \setminus \{r\}} \sum_{\substack{t \in T_i: \\ e \in \sigma_i(t)}} p(i,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-i}(\sigma, (\mathbf{p}|t_i = t)) + w(t)) + g_e(w(t)) \right].$$

Observe first that

$$\begin{aligned}
&\sum_{t \in T_r} \Phi_{(r,t)}(\sigma) + \sum_{e \in [m]} \Phi_e^{-r}(\sigma) \\
&= \sum_{t \in T_r} \sum_{e \in \sigma_r(t)} p(r,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = t)) + w(t)) + g_e(w(t)) \right] \\
&\quad + \sum_{e \in [m]} \sum_{i \in [n] \setminus \{r\}} \sum_{\substack{t \in T_i: \\ e \in \sigma_i(t)}} p(i,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-i}(\sigma, (\mathbf{p}|t_i = t)) + w(t)) + g_e(w(t)) \right] \\
&= \sum_{t \in T_r} \sum_{e \in \sigma_r(t)} p(r,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = t)) + w(t)) + g_e(w(t)) \right] \\
&\quad + \sum_{i \in [n] \setminus \{r\}} \sum_{t \in T_i} \sum_{e \in \sigma_i(t)} p(i,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-i}(\sigma, (\mathbf{p}|t_i = t)) + w(t)) + g_e(w(t)) \right] \\
&= \sum_{i \in [n]} \sum_{t \in T_i} \sum_{e \in \sigma_i(t)} p(i,t) \cdot w(t) \cdot \left[g_e(\delta_e^{-i}(\sigma, (\mathbf{p}|t_i = t)) + w(t)) + g_e(w(t)) \right] \\
&= \Phi(\sigma).
\end{aligned}$$

Consider a unilateral strategy change of type agent $(r, \hat{t})$ from the set of resources $\sigma_r(\hat{t}) \in S_r$ to the set of resources $\sigma_r'(\hat{t}) \in S_r$. Set $\sigma_r'(t) = \sigma_r(t)$ for all $t \in T_r \setminus \{\hat{t}\}$ and define $\sigma' = (\sigma_1, \ldots, \sigma_{r-1}, \sigma_r', \sigma_{r+1}, \ldots \sigma_n)$ as the pure strategy profile resulting from σ after this strategy change. Assume that $v_{(r,\hat{t})}(\sigma', \mathbf{p}) < v_{(r,\hat{t})}(\sigma, \mathbf{p})$, that is, the Individual Cost of type agent $(r, \hat{t})$ decreases. Thus,

$$\begin{aligned} & v_{(r,\hat{t})}(\sigma', \mathbf{p}) - v_{(r,\hat{t})}(\sigma, \mathbf{p}) \\ = & \sum_{e \in \sigma'_r(\hat{t})} g_e\left(\delta_e^{-r}(\sigma', (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right) - \sum_{e \in \sigma_r(\hat{t})} g_e\left(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right) \\ < & \ 0. \end{aligned}$$

Moreover,

- $\Phi_{(r,t)}(\sigma) = \Phi_{(r,t)}(\sigma')$ for all type agents (r, t) where $t \in T_r \setminus \{\hat{t}\}$, and
- $\Phi_e^{-r}(\sigma) = \Phi_e^{-r}(\sigma')$ for all resources e that are neither in $\sigma_r(\hat{t})$ nor in $\sigma'_r(\hat{t})$ or that are in both $\sigma_r(\hat{t})$ as well as in $\sigma'_r(\hat{t})$, that is, $e \in \left([m] \setminus (\sigma_r(\hat{t}) \cup \sigma'_r(\hat{t}))\right) \cup \left(\sigma_r(\hat{t}) \cap \sigma'_r(\hat{t})\right)$. Observe that these are the resources where the load does not change.

Now, consider the change $\Delta(\Phi)$ to the function Φ due to this strategy change of type agent $(r, \hat{t})$. Clearly,

$$\begin{aligned} & \Delta(\Phi) \\ = & \ \Phi(\sigma') - \Phi(\sigma) \\ = & \sum_{t \in T_r} \left(\Phi_{(r,t)}(\sigma') - \Phi_{(r,t)}(\sigma)\right) + \sum_{e \in [m]} \left(\Phi_e^{-r}(\sigma') - \Phi_e^{-r}(\sigma)\right) \\ = & \left(\Phi_{(r,\hat{t})}(\sigma') - \Phi_{(r,\hat{t})}(\sigma)\right) \\ & + \sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} \left(\Phi_e^{-r}(\sigma') - \Phi_e^{-r}(\sigma)\right) + \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} \left(\Phi_e^{-r}(\sigma') - \Phi_e^{-r}(\sigma)\right) \\ = & \ \Delta_1(\Phi) + \Delta_2(\Phi) + \Delta_3(\Phi), \end{aligned}$$

where:

$$\begin{aligned} \Delta_1(\Phi) &= \Phi_{(r,\hat{t})}(\sigma') - \Phi_{(r,\hat{t})}(\sigma); \\ \Delta_2(\Phi) &= \sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} \left(\Phi_e^{-r}(\sigma') - \Phi_e^{-r}(\sigma)\right); \\ \Delta_3(\Phi) &= \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} \left(\Phi_e^{-r}(\sigma') - \Phi_e^{-r}(\sigma)\right). \end{aligned}$$

Clearly,

$$\begin{aligned}
\Delta_1(\Phi) &= p(r,\hat{t})\cdot w(\hat{t})\cdot\Bigg(\sum_{e\in\sigma'_r(\hat{t})}\left[g_e\left(\delta_e^{-r}(\sigma',(\mathbf{p}|t_r=\hat{t}))+w(\hat{t})\right)+g_e(w(\hat{t}))\right] \\
&\qquad - \sum_{e\in\sigma_r(\hat{t})}\left[g_e\left(\delta_e^{-r}(\sigma,(\mathbf{p}|t_r=\hat{t}))+w(\hat{t})\right)+g_e(w(\hat{t}))\right]\Bigg) \\
&= p(r,\hat{t})\cdot w(\hat{t})\cdot\Bigg(\sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})}\left[g_e\left(\delta_e^{-r}(\sigma',(\mathbf{p}|t_r=\hat{t}))+w(\hat{t})\right)+g_e(w(\hat{t}))\right] \\
&\qquad - \sum_{e\in\sigma_r(\hat{t})\setminus\sigma'_r(\hat{t})}\left[g_e\left(\delta_e^{-r}(\sigma,(\mathbf{p}|t_r=\hat{t}))+w(\hat{t})\right)+g_e(w(\hat{t}))\right]\Bigg). \\
&= p(r,\hat{t})\cdot w(\hat{t})\cdot\Bigg(\sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})}\left[g_e\left(\delta_e^{-r}(\sigma,(\mathbf{p}|t_r=\hat{t}))+w(\hat{t})\right)+g_e(w(\hat{t}))\right] \\
&\qquad - \sum_{e\in\sigma_r(\hat{t})\setminus\sigma'_r(\hat{t})}\left[g_e\left(\delta_e^{-r}(\sigma,(\mathbf{p}|t_r=\hat{t}))+w(\hat{t})\right)+g_e(w(\hat{t}))\right]\Bigg).
\end{aligned}$$

Furthermore, due to the arrival of type agent $(r,\hat{t})$ on the resources $e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})$,

$$\begin{aligned}
&\Delta_2(\Phi) \\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})}\sum_{i\in[n]\setminus\{r\}}\sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} p(i,t)\cdot w(t)\cdot\Big(g_e(\delta_e^{-i}(\sigma',(\mathbf{p}|t_i=t))+w(t))+g_e(w(t)) \\
&\qquad - g_e(\delta_e^{-i}(\sigma,(\mathbf{p}|t_i=t))+w(t))-g_e(w(t))\Big) \\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})}\sum_{i\in[n]\setminus\{r\}}\sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} p(i,t)\cdot w(t)\cdot\Big(g_e(\delta_e^{-i}(\sigma',(\mathbf{p}|t_i=t)))-g_e(\delta_e^{-i}(\sigma,(\mathbf{p}|t_i=t)))\Big) \\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})}\sum_{i\in[n]\setminus\{r\}}\sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} p(i,t)\cdot w(t)\cdot a_e\Big(\delta_e^{-i}(\sigma',(\mathbf{p}|t_i=t))-\delta_e^{-i}(\sigma,(\mathbf{p}|t_i=t))\Big) \\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})}\sum_{i\in[n]\setminus\{r\}}\sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} p(i,t)\cdot w(t)\cdot a_e \\
&\qquad \cdot\sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} p(t_1,\ldots,t_{i-1},t_{i+1},\ldots,t_n|t_i=t)\cdot\left[\sum_{\substack{s\in[n]\setminus\{i\}:\\ e\in\sigma'_s(t_s)}} w(t_s)-\sum_{\substack{s\in[n]\setminus\{i\}:\\ e\in\sigma_s(t_s)}} w(t_s)\right].
\end{aligned}$$

Note that

$$\sum_{\substack{s\in[n]\setminus\{i\}:\\ e\in\sigma'_s(t_s)}} w(t_s) - \sum_{\substack{s\in[n]\setminus\{i\}:\\ e\in\sigma_s(t_s)}} w(t_s) \quad = \quad \begin{cases} w(\hat{t}), & \text{for all } (t_1,\dots,t_n)\in T \text{ where } t_r=\hat{t},\\ 0, & \text{else.}\end{cases}$$

Hence,

$$\begin{aligned}
&\Delta_2(\Phi)\\
&= \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} p(i,t)\cdot w(t) \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_i=t, t_r=\hat{t}}} p(t_1,\dots,t_{i-1},t_{i+1},\dots,t_n|t_i=t)\cdot w(\hat{t})\\
&= w(\hat{t}) \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} w(t) \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_i=t, t_r=\hat{t}}} p(t_1,\dots,t_n)\\
&= w(\hat{t}) \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_i=t, t_r=\hat{t}}} p(t_1,\dots,t_n)\cdot w(t)\\
&= p(r,\hat{t})\cdot w(\hat{t}) \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{t\in T_i:\\ e\in\sigma_i(t)}} \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_i=t, t_r=\hat{t}}} p(t_1,\dots,t_{r-1},t_{r+1},\dots,t_n|t_r=\hat{t})\cdot w(t)\\
&= p(r,\hat{t})\cdot w(\hat{t}) \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{i\in[n]\setminus\{r\}} \sum_{\substack{(t_1,\dots,t_n)\in T:\\ e\in\sigma_i(t_i), t_r=\hat{t}}} p(t_1,\dots,t_{r-1},t_{r+1},\dots,t_n|t_r=\hat{t})\cdot w(t_i)\\
&= p(r,\hat{t})\cdot w(\hat{t}) \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e \sum_{\substack{(t_1,\dots,t_n)\in T:\\ t_r=\hat{t}}} p(t_1,\dots,t_{r-1},t_{r+1},\dots,t_n|t_r=\hat{t}) \sum_{\substack{i\in[n]\setminus\{r\}:\\ e\in\sigma_i(t_i)}} w(t_i)\\
&= p(r,\hat{t})\cdot w(\hat{t})\cdot \sum_{e\in\sigma'_r(\hat{t})\setminus\sigma_r(\hat{t})} a_e\cdot \delta_e^{-r}(\sigma,(\mathbf{p}|t_r=\hat{t})).
\end{aligned}$$

Similarly, since type agent $(r,\hat{t})$ left the resources $e\in\sigma_r(\hat{t})\setminus\sigma'_r(\hat{t})$, we obtain that

$$\Delta_3(\Phi) \quad = \quad -p(r,\hat{t})\cdot w(\hat{t})\cdot \sum_{e\in\sigma_r(\hat{t})\setminus\sigma'_r(\hat{t})} a_e\cdot \delta_e^{-r}(\sigma,(\mathbf{p}|t_r=\hat{t})).$$

Hence,

$$
\begin{aligned}
&\Delta(\Phi)\\
&= \Delta_1(\Phi) + \Delta_2(\Phi) + \Delta_3(\Phi)\\
&= p(r,\hat{t}) \cdot w(\hat{t}) \cdot \Bigg(\sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t})) + a_e \cdot \delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t}))\right]\\
&\qquad - \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} \left[g_e\left(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right) + g_e(w(\hat{t})) + a_e \cdot \delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t}))\right] \Bigg)\\
&= 2 \cdot p(r,\hat{t}) \cdot w(\hat{t}) \cdot \Bigg(\sum_{e \in \sigma'_r(\hat{t}) \setminus \sigma_r(\hat{t})} g_e\left(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right)\\
&\qquad - \sum_{e \in \sigma_r(\hat{t}) \setminus \sigma'_r(\hat{t})} g_e\left(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right) \Bigg)\\
&= 2 \cdot p(r,\hat{t}) \cdot w(\hat{t}) \cdot \Bigg(\sum_{e \in \sigma'_r(\hat{t})} g_e\left(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right)\\
&\qquad - \sum_{e \in \sigma_r(\hat{t})} g_e\left(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right) \Bigg)\\
&= 2 \cdot p(r,\hat{t}) \cdot w(\hat{t}) \cdot \Bigg(\sum_{e \in \sigma'_r(\hat{t})} g_e\left(\delta_e^{-r}(\sigma', (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right)\\
&\qquad - \sum_{e \in \sigma_r(\hat{t})} g_e\left(\delta_e^{-r}(\sigma, (\mathbf{p}|t_r = \hat{t})) + w(\hat{t})\right) \Bigg)\\
&= 2 \cdot p(r,\hat{t}) \cdot w(\hat{t}) \cdot \left(v_{(r,\hat{t})}(\sigma', \mathbf{p}) - v_{(r,\hat{t})}(\sigma, \mathbf{p}) \right)\\
&< 0\,.
\end{aligned}
$$

Thus, any unilateral strategy change of a type agent that decreases its Individual Cost also decreases the value of the function Φ. Since the number of possible strategy profiles in Γ is finite, it follows that that there is a pure strategy profile that minimizes Φ. In this strategy profile, no type agent can decrease its Individual Cost by unilaterally changing its strategy. Hence, Γ has a pure Bayesian Nash equilibrium, as needed. ■

This generalizes a result of Fotakis *et al.* [13, Theorem 1] to the Bayesian setting. In particular our function Φ reduces to their potential function if each user has only a single type.

3.2 Computation

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 3.2 *Let $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ be a Bayesian routing game on identical links with independent type distribution. Then, a (normal) pure Bayesian Nash equilibrium for Γ can be computed in time polynomial in the size of Γ even if $\mathbf{p}$ is represented in a compact form by a set of probabilities $p(i, t)$ for $i \in [n]$ and $t \in T_i$.*

Proof: Given Γ compute a pure strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ as follows:

- Calculate for each user $i \in [n]$ its expected traffic $W(i)$.
- Construct a complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, \{(t'_1, ..., t'_n)\}, 1)$ where $w(t'_i) = W(i)$ for all $i \in [n]$.
- Compute a pure Nash equilibrium $\alpha : [n] \to [m]$ for Γ_{CI} in polynomial time with the LPT scheduling algorithm which assigns the users in order of non-increasing user traffics to minimum load links (see [12, 18]).
- Set $\sigma_i(t) = \alpha(i)$ for all users $i \in [n]$ and types $t \in T_i$.

We now show that $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a pure Bayesian Nash equilibrium for Γ. Recall that for an independent type distribution $\mathbf{p}$ we have that $W(k) = W(k|t_i = \hat{t})$ for all users $k, i \in [n]$ and all types $\hat{t} \in T_i$ of user $i \in [n]$. (This is the point in this proof where the independence is crucial. The rest of the used arguments also works for correlated type distributions.) With $W(k) = W(k|t_i = \hat{t})$ we get for any link $j \in [m]$,

$$
\begin{aligned}
\delta_j^{-i}(\alpha, 1) &= \sum_{\substack{k \in [n] \setminus \{i\}, \\ \alpha(k) = j}} W(k) \\
&= \sum_{\substack{k \in [n] \setminus \{i\}, \\ \alpha(k) = j}} W(k|t_i = \hat{t}) \\
&= \sum_{\substack{k \in [n] \setminus \{i\}, \\ \alpha(k) = j}} \sum_{\substack{(t_1, \ldots t_n) \in T: \\ t_i = \hat{t}}} p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i = \hat{t}) w(t_k) \\
&= \sum_{\substack{(t_1, \ldots t_n) \in T: \\ t_i = \hat{t}}} p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i = \hat{t}) \sum_{\substack{k \in [n] \setminus \{i\}, \\ \sigma_k(t_k) = j}} w(t_k) \\
&= \delta_j^{-i}(\sigma, (\mathbf{p}|t_i = \hat{t})).
\end{aligned}
$$

Assume, by way of contradiction, that σ is not a pure Bayesian Nash equilibrium for Γ. Then there exists a type agent $(i, \hat{t})$ for some user $i \in [n]$ and type $\hat{t} \in T_i$ that can improve by moving from $\sigma_i(\hat{t}) = \alpha(i)$ to another link $l \neq \sigma_i(\hat{t})$; thus

$$\begin{aligned}
\lambda_{(i,\hat{t})}^{\sigma_i(\hat{t})}(\sigma, \mathbf{p}) &= \delta_{\sigma_i(\hat{t})}^{-i}(\sigma, (\mathbf{p}|t_i = \hat{t})) + w(\hat{t}) \\
&> \delta_l^{-i}(\sigma, (\mathbf{p}|t_i = \hat{t})) + w(\hat{t}) \\
&= \lambda_{(i,\hat{t})}^{l}(\sigma, \mathbf{p}).
\end{aligned}$$

It follows that

$$\delta_{\sigma_i(\hat{t})}^{-i}(\sigma, (\mathbf{p}|t_i = \hat{t})) > \delta_l^{-i}(\sigma, (\mathbf{p}|t_i = \hat{t})).$$

This implies,

$$\begin{aligned}
v_{(i,t'_i)}(\alpha, \mathbf{p}) &= \delta_{\alpha(i)}^{-i}(\alpha, 1) + w(t'_i) \\
&= \delta_{\sigma_i(\hat{t})}^{-i}(\sigma, (\mathbf{p}|t_i = \hat{t})) + w(t'_i) \\
&> \delta_l^{-i}(\sigma, (\mathbf{p}|t_i = \hat{t})) + w(t'_i) \\
&= \delta_l^{-i}(\alpha, 1) + w(t'_i).
\end{aligned}$$

Therefore, in Γ_{CI}, user i can decrease its Individual Cost by switching from link $\alpha(i)$ to link l. Hence, α is not a Nash equilibrium for Γ_{CI}. A contradiction. ■

The algorithm used in the proof of Theorem 3.2 cannot be used to compute pure Bayesian Nash equilibria for the more general classes of Bayesian routing games either on related links or with correlated type distribution. The reason is that it always computes a *normal* Bayesian Nash equilibrium, while the following counter-examples show that a normal Bayesian Nash equilibrium does not exist in general.

Proposition 3.3 *There is a Bayesian routing game Γ on related links with independent type distribution that does not have a normal Bayesian Nash equilibrium.*

Proof: Consider the Bayesian routing game $\Gamma = (2, 2, \mathbf{c}, T_1 \times T_2, \mathbf{p})$ with two links of capacity $c_1 = 1$ and $c_2 = 5$. The two users have type sets $T_1 = \{t_1, t'_1\}$ and $T_2 = \{t_2\}$, where $w(t_1) = 1$, $w(t'_1) = 5$, $w(t_2) = 10$, and $p(1, t_1) = p(1, t'_1) = \frac{1}{2}$. We will now study the structure of pure Bayesian Nash equilibria for Γ and finally recognize that it has no *normal* pure Bayesian Nash equilibrium.

Let σ be an arbitrary pure Bayesian Nash equilibrium. Then,

$$\lambda^1_{(2,t_2)}(\sigma, \mathbf{p}) = \frac{\delta_1^{-2}(\sigma, \mathbf{p}) + w(t_2)}{c_1} \geq \frac{w(t_2)}{c_1} = 10$$

while

$$\lambda^2_{(2,t_2)}(\sigma, \mathbf{p}) = \frac{\delta_2^{-2}(\sigma, \mathbf{p}) + w(t_2)}{c_2} \leq \frac{\frac{1}{2} \cdot w(t_1) + \frac{1}{2} \cdot w(t'_1) + w(t_2)}{c_2} = \frac{13}{5} < 10.$$

Thus, σ assigns t_2 to link 2, so $\sigma_2(t_2) = 2$. Consider now the types of user 1. We have

$$\lambda^1_{(1,t_1)}(\sigma, \mathbf{p}) = \frac{w(t_1)}{c_1} = 1 \quad \text{and} \quad \lambda^2_{(1,t_1)}(\sigma, \mathbf{p}) = \frac{w(t_2) + w(t_1)}{c_2} = \frac{11}{5},$$
$$\lambda^1_{(1,t'_1)}(\sigma, \mathbf{p}) = \frac{w(t'_1)}{c_1} = 5 \quad \text{and} \quad \lambda^2_{(1,t'_1)}(\sigma, \mathbf{p}) = \frac{w(t_2) + w(t'_1)}{c_2} = 3.$$

So σ assigns t_1 to link 1 and t'_1 to link 2. It follows that σ is the unique pure Bayesian Nash equilibrium. However, σ is not a *normal* pure Bayesian Nash equilibrium. The claim follows. ■

Proposition 3.4 *There is a Bayesian routing game Γ on identical links with correlated type distribution that does not have a normal Bayesian Nash equilibrium.*

Proof: Consider the Bayesian routing game $\Gamma = (3, 2, \mathbf{1}, T_1 \times T_2 \times T_3, \mathbf{p})$ with 2 identical links and 3 users 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 $w(t_1) = w(t'_2) = w(t_3) = w(t'_3) = 1$ and $w(t'_1) = w(t_2) = 2$. The correlated distribution $\mathbf{p}$ is given by $p(t_1, t_2, t_3) = p(t'_1, t'_2, t'_3) = \frac{1}{2}$.

Assume, by way of contradiction, that a *normal* pure Bayesian Nash equilibrium σ exists; so, $\sigma_1(t_1) = \sigma_1(t'_1)$, $\sigma_2(t_2) = \sigma_2(t'_2)$, and $\sigma_3(t_3) = \sigma_3(t'_3)$. Let j and k be the two links. Without loss of generality, set $\sigma_1(t_1) = \sigma_1(t'_1) = j$. Then, clearly

$$\lambda^j_{(2,t'_2)}(\sigma, \mathbf{p}) \geq w(t'_1) + w(t'_2) = 3 \quad \text{while} \quad \lambda^k_{(2,t'_2)}(\sigma, \mathbf{p}) \leq w(t'_3) + w(t'_2) = 2.$$

Thus, $\sigma_2(t'_2) = k$; hence, $\sigma_2(t_2) = \sigma_2(t'_2) = k$ for all normal pure Bayesian Nash equilibria σ. For the types of user 3, note that

$$\lambda^j_{(3,t_3)}(\sigma, \mathbf{p}) = w(t_1) + w(t_3) = 2 \quad \text{while} \quad \lambda^k_{(3,t_3)}(\sigma, \mathbf{p}) = w(t_2) + w(t_3) = 3, \text{ and}$$
$$\lambda^j_{(3,t'_3)}(\sigma, \mathbf{p}) = w(t'_1) + w(t'_3) = 3 \quad \text{while} \quad \lambda^k_{(3,t'_3)}(\sigma, \mathbf{p}) = w(t'_2) + w(t'_3) = 2.$$

Since σ is a Bayesian Nash equilibrium, $\sigma_3(t_3) = j$ and $\sigma_3(t'_3) = k$. Hence, σ is not normal. A contradiction. ■

4 Fully Mixed Bayesian Nash Equilibria

In this section, we study fully mixed Bayesian Nash equilibria for the case of identical links. We start by proving a technical lemma that will be handy later on.

Lemma 4.1 *Consider a Bayesian routing game* $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ *on identical links and an associated mixed strategy profile* $\mathbf{Q}$*. Then, for each user* $i \in [n]$*,*

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

Proof: Clearly,

$$\begin{aligned}
&\sum_{j\in[m]} \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \\
&= \sum_{j\in[m]} \sum_{\sigma\in\Sigma} \prod_{s\in[n]} q(s, \sigma_s) \cdot \delta_j^{-i}(\sigma, (\mathbf{p}|t_i = t)) \\
&= \sum_{j\in[m]} \sum_{\sigma\in\Sigma} \prod_{s\in[n]} q(s, \sigma_s) \cdot \sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i = t) \sum_{\substack{s\in[n]\setminus\{i\}:\\ \sigma_s(t_s)=j}} w(t_s) \\
&= \sum_{\sigma\in\Sigma} \prod_{s\in[n]} q(s, \sigma_s) \cdot \sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i = t) \sum_{s\in[n]\setminus\{i\}} w(t_s) \\
&= \sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i = t) \sum_{s\in[n]\setminus\{i\}} w(t_s) \\
&= \sum_{s\in[n]\setminus\{i\}} \sum_{\substack{(t_1,\ldots,t_n)\in T:\\ t_i=t}} p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i = t) w(t_s) \\
&= \sum_{s\in[n]\setminus\{i\}} W(s|t_i = t),
\end{aligned}$$

as needed. ■

We continue to prove a simple expression for the Individual Cost for each user in a fully mixed Bayesian Nash equilibrium.

Proposition 4.2 *Consider a Bayesian routing game* $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ *on identical links and an associated fully mixed Bayesian Nash equilibrium* $\mathbf{F}$*. Then for each user* $i \in [n]$*,*

$$u_i(\mathbf{F}, \mathbf{p}) \quad = \quad \frac{W}{m} + \frac{m-1}{m} W(i).$$

Proof: Fix any user $i \in [n]$. Clearly, for any link $k \in \mathsf{support}_{F_i}(i) = [m]$, and by Lemma 4.1,

$$\begin{aligned}
u_i(\mathbf{F}, \mathbf{p}) &= \sum_{t \in T_i} p(i,t) \cdot v_{(i,t)}(\mathbf{F}, \mathbf{p}) \\
&= \sum_{t \in T_i} p(i,t) \cdot \big(w(t) + \delta_k^{-i}(\mathbf{F}, (\mathbf{p}|t_i = t))\big) \\
&= \sum_{t \in T_i} p(i,t) \cdot w(t) + \sum_{t \in T_i} p(i,t) \cdot \delta_k^{-i}(\mathbf{F}, (\mathbf{p}|t_i = t)) \\
&= W(i) + \sum_{t \in T_i} p(i,t) \cdot \frac{1}{m} \cdot \sum_{j \in [m]} \delta_j^{-i}(\mathbf{F}, (\mathbf{p}|t_i = t)) \\
&= W(i) + \frac{1}{m} \sum_{t \in T_i} p(i,t) \sum_{s \in [n] \setminus \{i\}} W(s|t_i = t) \qquad (1) \\
&= W(i) + \frac{1}{m} \sum_{s \in [n] \setminus \{i\}} \sum_{t \in T_i} p(i,t) \cdot W(s|t_i = t) \\
&= W(i) + \frac{1}{m} \sum_{s \in [n] \setminus \{i\}} W(s) \\
&= \frac{W}{m} + \frac{m-1}{m} W(i),
\end{aligned}$$

as needed. ■

We now prove that the Individual Cost of each user is maximized in a fully mixed Bayesian Nash equilibrium. For the special case of complete information routing games this result is already known [15].

Proposition 4.3 *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links and an associated fully mixed Bayesian Nash equilibrium $\mathbf{F}$ and Bayesian Nash equilibrium $\mathbf{Q}$. Then for each user $i \in [n]$,*

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

Proof: Fix any user $i \in [n]$. Then, for any link $j \in [m]$,

$$\begin{aligned}
u_i(\mathbf{Q}, \mathbf{p}) &= \sum_{t \in T_i} p(i,t) \cdot v_{(i,t)}(\mathbf{Q}, \mathbf{p}) \\
&\leq \sum_{t \in T_i} p(i,t) \cdot \Big(w(t) + \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t))\Big),
\end{aligned}$$

since $\mathbf{Q}$ is a Bayesian Nash equilibrium. In particular,

$$\begin{aligned}
u_i(\mathbf{Q}, \mathbf{p}) &\leq \sum_{t \in T_i} p(i,t) \left(w(t) + \min_{j \in [m]} \left\{ \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \right\} \right) \\
&\leq \sum_{t \in T_i} p(i,t) \left(w(t) + \frac{1}{m} \sum_{j \in [m]} \delta_j^{-i}(\mathbf{Q}, (\mathbf{p}|t_i = t)) \right) \\
&= \sum_{t \in T_i} p(i,t) \left(w(t) + \frac{1}{m} \sum_{s \in [n] \setminus \{i\}} W(s|t_i = t) \right) \\
&= W(i) + \frac{1}{m} \sum_{t \in T_i} p(i,t) \sum_{s \in [n] \setminus \{i\}} W(s|t_i = t) \\
&= u_i(\mathbf{F}, \mathbf{p}),
\end{aligned}$$

by Equation (1), as needed. ■

We proceed to define a particular fully mixed strategy profile $\overline{\mathbf{F}}$.

Definition 4.1 *The* standard *fully mixed strategy profile* $\overline{\mathbf{F}}$ *is 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 on identical links, the standard fully mixed strategy profile is a Bayesian Nash equilibrium. For the special case of complete information routing games, this fact was first stated in [26].

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 on identical links with independent type distribution. We start with an exact characterization of fully mixed Bayesian Nash equilibria in type agent representation.

Proposition 4.4 *Consider a Bayesian routing game* $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ *on identical links with independent type distribution and an associated fully mixed strategy profile* $\mathbf{F}$. *Then,* $\mathbf{F}$ *is a fully mixed Bayesian Nash equilibrium if and only if*

$$\frac{1}{m} \cdot W(i) = \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \sum_{\substack{t \in T_i: \\ \sigma_i(t) = j}} p(i,t) \cdot w(t)$$

for all users $i \in [n]$ *and links* $j \in [m]$.

Proof: For any user $i \in [n]$ and link $j \in [m]$, set

$$\mu(\mathbf{F}, i, j) = \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \sum_{\substack{t \in T_i: \\ \sigma_i(t) = j}} p(i, t) \cdot w(t).$$

Observe that for any user $i \in [n]$ and link $j \in [m]$,

$$\begin{aligned}
\delta_j^{-i}(\mathbf{F}, \mathbf{p}) &= \sum_{\sigma \in \Sigma} \prod_{s \in [n]} f(s, \sigma_s) \sum_{(t_1, \ldots, t_n) \in T} p(t_1, \ldots, t_n) \sum_{\substack{k \in [n] \setminus \{i\}: \\ \sigma_k(t_k) = j}} w(t_k) \\
&= \sum_{\sigma \in \Sigma} \prod_{s \in [n]} f(s, \sigma_s) \sum_{k \in [n] \setminus \{i\}} \sum_{\substack{t_k \in T_k: \\ \sigma_k(t_k) = j}} p(k, t_k) \cdot w(t_k) \\
&= \sum_{k \in [n] \setminus \{i\}} \sum_{\sigma \in \Sigma} \prod_{s \in [n]} f(s, \sigma_s) \sum_{\substack{t_k \in T_k: \\ \sigma_k(t_k) = j}} p(k, t_k) \cdot w(t_k) \\
&= \sum_{k \in [n] \setminus \{i\}} \sum_{\sigma'_k \in \Sigma_k} f(k, \sigma'_k) \sum_{\substack{\sigma \in \Sigma: \\ \sigma_k = \sigma'_k}} \prod_{s \in [n] \setminus \{k\}} f(s, \sigma_s) \sum_{\substack{t_k \in T_k: \\ \sigma_k(t_k) = j}} p(k, t_k) \cdot w(t_k) \\
&= \sum_{k \in [n] \setminus \{i\}} \sum_{\sigma'_k \in \Sigma_k} f(k, \sigma'_k) \sum_{\substack{t_k \in T_k: \\ \sigma_k(t_k) = j}} p(k, t_k) \cdot w(t_k) \\
&= \sum_{k \in [n] \setminus \{i\}} \mu(\mathbf{F}, k, j) \; .
\end{aligned}$$

Consider first an arbitrary fully mixed strategy profile $\mathbf{F}$ that satisfies $\mu(\mathbf{F}, i, j) = \frac{1}{m} \cdot W(i)$ for all users $i \in [n]$ and links $j \in [m]$. Then, for all users $i \in [n]$, types $t \in T_i$, and links $j \in [m]$,

$$\begin{aligned}
\lambda_{(i,t)}^{j}(\mathbf{F}, \mathbf{p}) &= \delta_j^{-i}(\mathbf{F}, (\mathbf{p} | t_i = t)) + w(t) \\
&= \delta_j^{-i}(\mathbf{F}, \mathbf{p}) + w(t) \\
&= \sum_{k \in [n] \setminus \{i\}} \mu(\mathbf{F}, k, j) + w(t) \\
&= \sum_{k \in [n] \setminus \{i\}} \frac{1}{m} \cdot W(k) + w(t).
\end{aligned}$$

Hence we get for the Individual Cost of type agent (i, t),

$$\begin{aligned}
v_{(i,t)}(\mathbf{F}, \mathbf{p}) &= \sum_{\sigma_i \in \Sigma_i} f(i, \sigma_i) \cdot \lambda_{(i,t)}^{\sigma_i(t)}(\mathbf{F}, \mathbf{p}) \\
&= \sum_{k \in [n] \setminus \{i\}} \frac{1}{m} \cdot W(k) + w(t).
\end{aligned}$$

So, $v_{(i,t)}(\mathbf{F}, \mathbf{p}) = \lambda^j_{(i,t)}(\mathbf{F}, \mathbf{p})$ for all users $i \in [n]$, types $t \in T_i$, and links $j \in [m]$. Thus, $\mathbf{F}$ is a fully mixed Bayesian Nash equilibrium.

Assume now that $\mathbf{F}$ is a fully mixed Bayesian Nash Equilibrium. So, $\mathsf{support}_{Q_i}(t) = [m]$ for all users $i \in [n]$ and types $t \in T_i$. Since $\mathbf{F}$ is a fully mixed Bayesian Nash Equilibrium and $\mathbf{p}$ is independent, it follows that for all links $j \in \mathsf{support}_{Q_i}(t) = [m]$,

$$
\begin{aligned}
v_{(i,t)}(\mathbf{F}, \mathbf{p}) &= \lambda^j_{(i,t)}(\mathbf{F}, \mathbf{p}) \\
&= \delta^{-i}_j(\mathbf{F}, (\mathbf{p}|t_i = t)) + w(t) \\
&= \delta^{-i}_j(\mathbf{F}, \mathbf{p}) + w(t).
\end{aligned}
$$

So, for all users $i \in [n]$ and pair of links $j, l \in [m]$,

$$\delta^{-i}_j(\mathbf{F}, \mathbf{p}) = \delta^{-i}_l(\mathbf{F}, \mathbf{p}).$$

Since $\delta^{-i}_j(\mathbf{F}, \mathbf{p}) = \sum_{k \in [n] \setminus \{i\}} \mu(\mathbf{F}, k, j)$ for any user i and link j, it follows that for an arbitrary pair of users $i_1, i_2 \in [n]$ with $i_1 \neq i_2$ and an arbitrary pair of links $j_1, j_2 \in [m]$ with $j_1 \neq j_2$,

$$\sum_{k \in [n] \setminus \{i_1\}} \mu(\mathbf{F}, k, j_1) = \sum_{k \in [n] \setminus \{i_1\}} \mu(\mathbf{F}, k, j_2) \tag{2}$$

and

$$\sum_{k \in [n] \setminus \{i_2\}} \mu(\mathbf{F}, k, j_1) = \sum_{k \in [n] \setminus \{i_2\}} \mu(\mathbf{F}, k, j_2). \tag{3}$$

Subtracting (3) from (2) yields that

$$\mu(\mathbf{F}, i_2, j_1) - \mu(\mathbf{F}, i_1, j_1) = \mu(\mathbf{F}, i_2, j_2) - \mu(\mathbf{F}, i_1, j_2),$$

or equivalently

$$0 = \mu(\mathbf{F}, i_2, j_1) - \mu(\mathbf{F}, i_2, j_2) + \mu(\mathbf{F}, i_1, j_2) - \mu(\mathbf{F}, i_1, j_1).$$

Summing up over all users $i_2 \in [n] \setminus \{i_1\}$ yields that

$$
\begin{aligned}
0 &= \sum_{i_2 \in [n] \setminus \{i_1\}} \mu(\mathbf{F}, i_2, j_1) - \sum_{i_2 \in [n] \setminus \{i_1\}} \mu(\mathbf{F}, i_2, j_2) + \sum_{i_2 \in [n] \setminus \{i_1\}} \mu(\mathbf{F}, i_1, j_2) - \sum_{i_2 \in [n] \setminus \{i_1\}} \mu(\mathbf{F}, i_1, j_1) \\
&= \delta^{-i_1}_{j_1}(\mathbf{F}, \mathbf{p}) - \delta^{-i_1}_{j_2}(\mathbf{F}, \mathbf{p}) + (n-1) \cdot \mu(\mathbf{F}, i_1, j_2) - (n-1) \cdot \mu(\mathbf{F}, i_1, j_1) \\
&= (n-1) \cdot \left(\mu(\mathbf{F}, i_1, j_2) - \mu(\mathbf{F}, i_1, j_1) \right).
\end{aligned}
$$

It follows that for all users $i_1 \in [n]$ and pair of links $j_1, j_2 \in [m]$,

$$\mu(\mathbf{F}, i_1, j_1) = \mu(\mathbf{F}, i_1, j_2).$$

Clearly, for any user $i \in [n]$,

$$\begin{aligned}
W(i) &= \sum_{t \in T_i} p(i,t) \cdot w(t) \\
&= \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \sum_{t \in T_i} p(i,t) \cdot w(t) \\
&= \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \sum_{j \in [m]} \sum_{\substack{t \in T_i: \\ \sigma_i(t)=j}} p(i,t) \cdot w(t) \\
&= \sum_{j \in [m]} \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \sum_{\substack{t \in T_i: \\ \sigma_i(t)=j}} p(i,t) \cdot w(t) \\
&= \sum_{j \in [m]} \mu(\mathbf{F}, i, j) \\
&= m \cdot \mu(\mathbf{F}, i, j),
\end{aligned}$$

for any link $j \in [m]$. This implies that for all users $i \in [n]$ and links $j \in [m]$,

$$\mu(\mathbf{F}, i, j) = \frac{1}{m} \cdot W(i)$$

or

$$\frac{1}{m} \cdot W(i) = \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \sum_{\substack{t \in T_i: \\ \sigma_i(t)=j}} p(i,t) \cdot w(t),$$

as needed. ■

We finally determine a lower bound on the dimension of the space of fully mixed Bayesian Nash equilibria.

Theorem 4.5 *Consider a Bayesian routing game Γ on identical links with independent type distribution. Then, the dimension of the space of fully mixed Bayesian Nash equilibria for Γ is at least $\sum_{i \in [n]} m^{\tau_i} - nm$.*

Proof: Let $\mathbf{F}$ be a fully mixed Bayesian Nash equilibrium. By Proposition 4.4, this is equivalent to $\mathbf{F}$ being a fully mixed strategy profile and

$$\sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \sum_{t \in T_i : \sigma_i(t)=j} p(i,t) \cdot w(t) \quad = \quad \frac{1}{m} \cdot W(i)$$

for all users $i \in [n]$ and links $j \in [m]$. So, $\mathbf{F}$ is a solution to the system of linear equations and inequalities:

$$
\begin{aligned}
&(1)\ f(i,\sigma_i) > 0 && \forall i \in [n], \forall \sigma_i \in \Sigma_i \\
&(2)\ \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) = 1 && \forall i \in [n] \\
&(3)\ \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \sum_{t \in T_i : \sigma_i(t) = j} p(i,t) \cdot w(t) = \frac{1}{m} \cdot W(i) && \forall i \in [n], \forall j \in [m]
\end{aligned}
$$

The dimension of the solution space of this system is the number of variables minus the number of independent equations. For each user $i \in [n]$ we have m^{τ_i} variables. Thus, the total number of variables is $\sum_{i \in [n]} m^{\tau_i}$. We now show an upper bound on the number of independent equations. Fix any user $i \in [n]$. Summing up the equations (3) for all links $j \in [m]$ yields

$$
\begin{aligned}
& \sum_{j \in [m]} \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \sum_{t \in T_i : \sigma_i(t) = j} p(i,t) \cdot w(t) = \sum_{j \in [m]} \frac{1}{m} \cdot W(i) \\
\Leftrightarrow\quad & \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \sum_{t \in T_i} p(i,t) \cdot w(t) = W(i) \\
\Leftrightarrow\quad & \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) \cdot W(i) = W(i) \\
\Leftrightarrow\quad & \sum_{\sigma_i \in \Sigma_i} f(i,\sigma_i) = 1
\end{aligned}
$$

It follows that all equations (2) are implied by a linear combination of equations in (3). Therefore, nm is an upper bound on the number of independent equations. The claim follows. ■

5 Social Cost and Coordination Ratio

In this section, we present bounds on the Coordination Ratio for three different Social Cost measures. All these results are for the case of *identical* links.

5.1 Social Cost as Expected Maximum Latency

In this section, we study Social Cost as the Expected Maximum Latency. For the special case of complete information routing games this Social Cost measure was introduced in [22] and asymptotic tight bounds on Coordination Ratio were given in [6, 21]. 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 $\mathcal{O}(\frac{\log m}{\log \log m})$. We prove that previous techniques cannot be applied to prove an upper bound on Coordination Ratio which is better than $\mathcal{O}(m)$.

Proposition 5.1 *For any $\epsilon > 0$, there is a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links with independent type distribution and an associated pure Bayesian Nash equilibrium σ with $\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma)$, such that for each link $j \in [m]$,*

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\delta_j(\sigma,\mathbf{p})} \geq \frac{m}{1+\epsilon}.$$

Proof: Set $n = m$; for each user $i \in [n]$, set $T_i = \{t_i, t'_i\}$ with $w(t_i) = 0$ and $w(t'_i) = a$; set also for each user $i \in [n]$, $p(i, t_i) = 1 - \frac{1}{a}$ and $p(i, t'_i) = \frac{1}{a}$. Let σ be the pure Bayesian Nash equilibrium that maps both types of user i to link i, where $i \in [n]$. Since each user is assigned to a different link, we have $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)$. Clearly, $\delta_j(\sigma,\mathbf{p}) = 1$ for all links $j \in [m]$. On the other hand,

$$\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma) \;=\; \left(1 - \left(1 - \frac{1}{a}\right)^m\right) a.$$

Note that

$$\begin{aligned}
\lim_{a\to\infty}\left(\left(1 - \left(1 - \frac{1}{a}\right)^m\right) a\right) &= \lim_{a\to\infty}\left(\left(1 - \sum_{i=0}^{m}\binom{m}{i}\left(-\frac{1}{a}\right)^i\right) a\right)\\
&= \lim_{a\to\infty}\left(\left(1 - 1 - \sum_{i=1}^{m}\binom{m}{i}(-1)^i\left(\frac{1}{a}\right)^i\right) a\right)\\
&= \lim_{a\to\infty}\left(\sum_{i=1}^{m}\binom{m}{i}(-1)^{i-1}\left(\frac{1}{a}\right)^{i-1}\right)\\
&= \lim_{a\to\infty}\left(m + \sum_{i=2}^{m}\binom{m}{i}(-1)^{i-1}\left(\frac{1}{a}\right)^{i-1}\right)\\
&= m.
\end{aligned}$$

The claim follows. ■

We now turn our attention to the standard fully mixed Bayesian Nash equilibrium on identical links. We prove:

Theorem 5.2 *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links and an associated standard fully mixed Bayesian Nash equilibrium $\overline{\mathbf{F}}$. Then,*

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}},\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} \;=\; \mathcal{O}\left(\frac{\log m}{\log\log m}\right).$$

Proof: Consider an arbitrary type profile $t = (t_1, ..., t_n) \in T$. Given t, we define the game $\Gamma_{\mathsf{CI}}(t) = (n, m, \mathbf{1}, \{(t_1, ..., t_n)\}, 1)$. Recall that for this complete information routing game $\Gamma_{\mathsf{CI}}(t)$, the unique fully mixed Nash equilibrium $\overline{\mathbf{Q}}(t)$ assigns each user to each link with probability $1/m$ (see [26, Lemma 15]). By [21, Theorem 4.4] or [6, Theorem 1.1], it holds that

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{Q}}, \Gamma_{\mathsf{CI}}(t))}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}(t))} = \mathcal{O}\left(\frac{\log m}{\log \log m}\right).$$

Recall that $\overline{\mathbf{F}}$ assigns every type agent to every link with probability $\frac{1}{m}$. Thus,

$$\begin{aligned}
\mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma) &= \sum_{t \in T} p(t) \cdot \sum_{(\sigma_1(t_1), ..., \sigma_n(t_n)) \in [m]^n} \left(\frac{1}{m}\right)^n \cdot \max_{j \in [m]} \left\{ \sum_{\substack{i \in [n]: \\ \sigma_i(t_i) = j}} w(t_i) \right\} \\
&= \sum_{t \in T} p(t) \cdot \mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{Q}}(t), \Gamma_{\mathsf{CI}}(t)) \\
&= \sum_{t \in T} p(t) \cdot \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}(t)) \cdot \mathcal{O}\left(\frac{\log m}{\log \log m}\right) \\
&= \mathsf{OPT}_{\mathsf{MSP}}(\Gamma) \cdot \mathcal{O}\left(\frac{\log m}{\log \log m}\right),
\end{aligned}$$

as needed. ■

Theorem 5.2 implies that for the standard fully mixed Nash equilibrium, incomplete information has no impact on the Coordination Ratio when Social Cost is taken as Expected Maximum Latency.

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

Proposition 5.3 *There exists a Bayesian routing game Γ on identical links and an associated fully mixed Bayesian Nash equilibrium $\mathbf{F}$ such that*

$$\mathsf{SC}_{\mathsf{MSP}}(\mathbf{F}, \Gamma) \quad > \quad \mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma).$$

Proof: Consider the Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ with $n = 2, m = 3$ and $T_i = \{t_i, t'_i\}$ with $w(t_i) = 2, w(t'_i) = 1$ for all users $i \in \{1, 2\}$; set $p(i, t_i) = p(i, t'_i) = \frac{1}{2}$ for all users $i \in \{1, 2\}$. Consider the standard fully mixed Bayesian Nash equilibrium $\overline{\mathbf{F}}$ and some other fully mixed Bayesian Nash equilibrium $\mathbf{F}$ which we define below:

- $\overline{\mathbf{F}}$ assigns each type to each link with a probability of $\frac{1}{3}$. Thus, the two users are assigned to the same link with a probability of $\frac{1}{3}$. In this case, the maximum latency can be 2, 3, or 4. With a probability of $\frac{2}{3}$, the users are assigned to different links. In this case the maximum latency can be 1 or 2. Hence, the Social Cost of the standard fully mixed Bayesian Nash equilibrium $\overline{\mathbf{F}}$ is

$$\mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma) = \frac{1}{3} \cdot \left(\frac{1}{4} \cdot 2 + \frac{1}{2} \cdot 3 + \frac{1}{4} \cdot 4 \right) + \frac{2}{3} \cdot \left(\frac{1}{4} \cdot 1 + \frac{3}{4} \cdot 2 \right) = \frac{13}{6}.$$

- The fully mixed strategy profile $\mathbf{F}$ assigns each type of traffic 1 to link 1 with a probability of $\frac{1}{2}$, to link 2 with a probability of $\frac{1}{4}$ and to link 3 with a probability of $\frac{1}{4}$. Each type of traffic 2 is assigned to link 1 with a probability of $\frac{1}{4}$, to link 2 with a probability of $\frac{3}{8}$ and to link 3 with a probability of $\frac{3}{8}$. Observe that for all $i \in \{1, 2\}$ we get $\delta_1^{-i}(\mathbf{F}) = \frac{1}{2} \cdot \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{4} \cdot 2 = \frac{1}{2}$ and $\delta_2^{-i}(\mathbf{F}) = \delta_3^{-i}(\mathbf{F}) = \frac{1}{2} \cdot \frac{1}{4} \cdot 1 + \frac{1}{2} \cdot \frac{3}{8} \cdot 2 = \frac{1}{2}$. Thus, $\mathbf{F}$ is a Bayesian Nash equilibrium.

  With probability $\frac{1}{4}$ both users are of traffic 1. In this case they use the same link with probability $(\frac{1}{2})^2 + 2 \cdot (\frac{1}{4})^2 = \frac{3}{8}$. With probability $\frac{1}{2}$, exactly one of the two users is of traffic 1. In this case, the users use the same link with probability $\frac{1}{2} \cdot \frac{1}{4} + 2 \cdot \frac{1}{4} \cdot \frac{3}{8} = \frac{5}{16}$. With the remaining probability $\frac{1}{4}$, both users are of traffic 2. In this case, the users use the same link with probability $(\frac{1}{4})^2 + 2 \cdot (\frac{3}{8})^2 = \frac{11}{32}$. Hence we get that the Social Cost of $\mathbf{F}$ is

$$\mathsf{SC}_{\mathsf{MSP}}(\mathbf{F}, \Gamma) = \frac{1}{4} \cdot \left(\frac{3}{8} \cdot 2 + \frac{5}{8} \cdot 1 \right) + \frac{1}{2} \cdot \left(\frac{5}{16} \cdot 3 + \frac{11}{16} \cdot 2 \right) + \frac{1}{4} \cdot \left(\frac{11}{32} \cdot 4 + \frac{21}{32} \cdot 2 \right) = \frac{139}{64}.$$

Observe that $\mathsf{SC}_{\mathsf{MSP}}(\mathbf{F}, \Gamma) = \frac{139}{64} = \frac{417}{192} > \frac{416}{192} = \frac{13}{6} = \mathsf{SC}_{\mathsf{MSP}}(\overline{\mathbf{F}}, \Gamma)$. ■

It is known (see [25, Section 8.E]) that mixed Nash equilibria in games with complete information are related to pure Bayesian Nash equilibria in a Bayesian game, where for each user all its types are identical. The following definition and theorem applies this to Bayesian routing games.

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

We call these games CI-like games (where CI stands for complete information) since they are similar to complete information routing games in the sense that the traffic of a user does not depend on its type. For complete information routing games, there exist asymptotically tight upper bounds on the Coordination Ratio for the cases of identical links [6, 21] and related links [6]. We use these bounds to prove:

Theorem 5.4 *Let $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ be a CI-like game with an associated pure Bayesian Nash equilibrium σ. Then*

(a) $\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} = \mathcal{O}\left(\frac{\log m}{\log\log m}\right)$*, for the case of identical links,*

(b) $\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} = \mathcal{O}\left(\frac{\log m}{\log\log\log m}\right)$*, for the case of related links,*

and there are CI-like games for which both bounds are asymptotically tight.

Proof: The proof is structured as follows: We first define a construction that maps any CI-like game Γ with an associated pure strategy profile σ to a complete information routing Γ_{CI} with associated (mixed) strategy profile $\mathbf{Q}$. For this construction, we show that $\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q}, \Gamma_{\mathsf{CI}})$, $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$, and that $\mathbf{Q}$ is a Nash equilibrium if σ is a Bayesian Nash equilibrium. From these properties of our construction, we derive that the corresponding upper bounds on the Coordination Ratio [6, 21] for complete information routing games also hold for CI-like games. To prove tightness, we show that for every complete information routing game Γ_{CI} with associated (mixed) Nash equilibrium $\mathbf{Q}$, we can define a CI-like game Γ with associated pure Bayesian Nash equilibrium σ, such that our construction maps Γ and σ to Γ_{CI} and $\mathbf{Q}$, respectively. This implies that also the lower bounds on the Coordination Ratio can be carried over to the CI-like games.

We start by defining our construction.

Construction $\Gamma \mapsto \Gamma_{\mathsf{CI}}$: Let $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ be a CI-like game. For each $i \in [n]$, denote by $w_i = w(t)$ the traffic of all types $t \in T_i$. Define a complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{c}, T', 1)$ where $T' = \{(t'_1, \ldots, t'_n)\}$ and $w(t'_i) = w_i$ for all $i \in [n]$.

Let $\sigma = (\sigma_1, \ldots, \sigma_n)$ be a pure strategy profile for the CI-like game Γ. Denote by Σ' the set of all pure strategy profiles for Γ_{CI}; thus, $\Sigma' = \Sigma'_1 \times \ldots \times \Sigma'_n$, where for each user $i \in [n]$, the set Σ'_i consists of all possible pure strategies $\sigma'_i : \{t'_i\} \to [m]$ for user i.

Define a mixed strategy profile $\mathbf{Q}$ for Γ_{CI}, where for each user $i \in [n]$ and all pure strategies $\sigma'_i \in \Sigma'_i$ the probability $q(i, \sigma'_i)$ is given by $q(i, \sigma'_i) = \sum_{t \in T_i : \sigma_i(t) = \sigma'_i(t'_i)} p(i, t)$.

We proceed by showing properties of our construction.

- $\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q}, \Gamma_{\mathsf{CI}})$: To show that the strategy profiles σ for Γ and $\mathbf{Q}$ for Γ_{CI}

are of the same Social Cost observe that

$$\begin{aligned}
\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma) &= \sum_{(t_1,\ldots,t_n)\in T} p(t_1,\ldots,t_n)\cdot \max_{j\in[m]}\left\{\frac{1}{c_j}\sum_{\substack{i\in[n],\\ \sigma_i(t_i)=j}} w(t_i)\right\}\\
&= \sum_{(t_1,\ldots,t_n)\in T}\prod_{i\in[n]} p(i,t_i)\cdot \max_{j\in[m]}\left\{\frac{1}{c_j}\sum_{\substack{i\in[n],\\ \sigma_i(t_i)=j}} w_i\right\}\\
&= \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'}\left(\sum_{\substack{(t_1,\ldots,t_n)\in T:\\ \sigma_i(t_i)=\sigma'_i(t'_i)\forall i\in[n]}}\prod_{i\in[n]} p(i,t_i)\right)\cdot \max_{j\in[m]}\left\{\frac{1}{c_j}\sum_{\substack{i\in[n],\\ \sigma'_i(t'_i)=j}} w_i\right\}\\
&= \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'}\left(\prod_{i\in[n]}\sum_{\substack{t\in T_i:\\ \sigma_i(t)=\sigma'_i(t'_i)}} p(i,t)\right)\cdot \max_{j\in[m]}\left\{\frac{1}{c_j}\sum_{\substack{i\in[n],\\ \sigma'_i(t'_i)=j}} w_i\right\}\\
&= \sum_{(\sigma'_1,\ldots,\sigma'_n)\in\Sigma'}\prod_{i\in[n]} q(i,\sigma'_i)\cdot \max_{j\in[m]}\left\{\frac{1}{c_j}\sum_{\substack{i\in[n],\\ \sigma'_i(t'_i)=j}} w_i\right\}\\
&= \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}}).
\end{aligned}$$

- $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)=\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$: To show $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)\geq\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$ observe that our construction maps a pure strategy profile for Γ of optimum Social Cost to a strategy profile for Γ_{CI} that has the same Social Cost.

  For the other direction $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)\leq\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$, observe that there always exists a *pure* strategy profile $\hat{\sigma}'$ for Γ_{CI} of optimum Social Cost, i.e. $\mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma}',\Gamma_{\mathsf{CI}})=\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$. Consider the normal pure strategy profile $\hat{\sigma}$ for Γ that assigns for each $i\in[n]$ all types of user i to the link to that $\hat{\sigma}'$ assigns user i, so $\hat{\sigma}_i(t)=\hat{\sigma}_i{}'(t'_i)$ for all users $i\in[n]$ and all types $t\in T_i$. Notice that our construction transforms Γ and $\hat{\sigma}$ back to Γ_{CI} and $\hat{\sigma}'$. Thus $\mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma},\Gamma)=\mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma}',\Gamma_{\mathsf{CI}})$. We get that

$$\begin{aligned}
\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) &\leq \mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma},\Gamma)\\
&= \mathsf{SC}_{\mathsf{MSP}}(\hat{\sigma}',\Gamma_{\mathsf{CI}})\\
&= \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}}).
\end{aligned}$$

- <u>Mapping of Equilibria</u>: Clearly, for all users $i \in [n]$, types $t \in T_i$, and links $j \in [m]$,

$$\begin{aligned}
\lambda^j_{(i,t)}(\sigma, \mathbf{p}) &= \frac{1}{c_j} \cdot \left(w(t) + \delta_j^{-i}(\sigma, \mathbf{p}) \right) \\
&= \frac{1}{c_j} \cdot \left(w(t) + \sum_{(t_1,\ldots,t_n) \in T} p(t_1, \ldots, t_n) \cdot \sum_{\substack{s \in [n] \setminus \{i\}: \\ \sigma_s(t_s) = j}} w(t_s) \right) \\
&= \frac{1}{c_j} \cdot \left(w_i + \sum_{(t_1,\ldots,t_n) \in T} \prod_{s \in [n]} p(s, t_s) \cdot \sum_{\substack{s \in [n] \setminus \{i\}: \\ \sigma_s(t_s) = j}} w_s \right) \\
&= \frac{1}{c_j} \cdot \left(w_i + \sum_{(\sigma'_1,\ldots,\sigma'_n) \in \Sigma'} \left(\sum_{\substack{(t_1,\ldots,t_n) \in T: \\ \sigma_s(t_s) = \sigma'_s(t'_s) \forall s \in [n]}} \prod_{s \in [n]} p(s, t_s) \right) \cdot \sum_{\substack{s \in [n] \setminus \{i\}: \\ \sigma'_s(t'_s) = j}} w_s \right) \\
&= \frac{1}{c_j} \cdot \left(w_i + \sum_{(\sigma'_1,\ldots,\sigma'_n) \in \Sigma'} \left(\prod_{s \in [n]} \sum_{\substack{t_s \in T_s: \\ \sigma_s(t_s) = \sigma'_s(t'_s)}} p(s, t_s) \right) \cdot \sum_{\substack{s \in [n] \setminus \{i\}: \\ \sigma'_s(t'_s) = j}} w_s \right) \\
&= \frac{1}{c_j} \cdot \left(w(t'_i) + \sum_{(\sigma'_1,\ldots,\sigma'_n) \in \Sigma'} \prod_{s \in [n]} q(s, \sigma'_s) \cdot \sum_{\substack{s \in [n] \setminus \{i\}: \\ \sigma'_s(t'_s) = j}} w(t'_s) \right) \\
&= \frac{1}{c_j} \cdot \left(w(t'_i) + \delta_j^{-i}(\mathbf{Q}, 1) \right) \\
&= \lambda^j_{(i,t'_i)}(\mathbf{Q}, 1).
\end{aligned}$$

We now use this property to show that $\mathbf{Q}$ is a Nash equilibrium for Γ_{CI} if σ is a pure Bayesian Nash equilibrium for Γ. So, let σ be a pure Bayesian Nash equilibrium for Γ. Fix an arbitrary user $i \in [n]$. Remember that in Γ all types of user i have the same traffic. Thus,

$$v_{(i,t)}(\sigma, \mathbf{p}) = v_{(i,\hat{t})}(\sigma, \mathbf{p})$$

for all pairs of types $t, \hat{t} \in T_i$. Since σ is a pure Bayesian Nash equilibrium for Γ, this implies that for all types $t \in T_i$,

$$\begin{aligned}
v_{(i,t)}(\sigma, \mathbf{p}) = \lambda^j_{(i,t)}(\sigma, \mathbf{p}) &\quad \text{for all } j \in \mathsf{support}_{\sigma_i}(i) \text{ and} \\
v_{(i,t)}(\sigma, \mathbf{p}) \leq \lambda^j_{(i,t)}(\sigma, \mathbf{p}) &\quad \text{for all } j \notin \mathsf{support}_{\sigma_i}(i).
\end{aligned}$$

By definition of $\mathbf{Q}$,

$$\mathsf{support}_{\sigma_i}(i) \quad = \quad \mathsf{support}_{\mathbf{Q}_i}(t'_i).$$

It follows that

$$\begin{array}{ll} v_{(i,t'_i)}(\mathbf{Q},1) = \lambda^j_{(i,t'_i)}(\mathbf{Q},1) & \text{for all } j \in \mathsf{support}_{\mathbf{Q}_i}(t'_i) \text{ and} \\ v_{(i,t'_i)}(\mathbf{Q},1) \leq \lambda^j_{(i,t'_i)}(\mathbf{Q},1) & \text{for all } j \notin \mathsf{support}_{\mathbf{Q}_i}(t'_i), \end{array}$$

so that $\mathbf{Q}$ is a Nash equilibrium.

Upper bounds on Coordination Ratio: Recall that by our construction, we have that $\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})$ and $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$. Thus, resorting to the corresponding upper bounds on Coordination Ratio from [21] and [6], we get

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} = \frac{\mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})} \quad = \quad \begin{cases} \mathcal{O}\left(\frac{\log m}{\log\log m}\right), & \text{for the case of identical links,} \\ \mathcal{O}\left(\frac{\log m}{\log\log\log m}\right), & \text{for the case of related links.} \end{cases}$$

This completes the proof of the upper bounds.

Tightness of the upper bounds: From [21] and [6], there exist complete information routing games Γ_{CI} with an associated mixed Nash equilibrium $\mathbf{Q}$ such that

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})} \quad = \quad \begin{cases} \Omega\left(\frac{\log m}{\log\log m}\right) & \text{, for the case of identical links,} \\ \Omega\left(\frac{\log m}{\log\log\log m}\right) & \text{, for the case of related links.} \end{cases}$$

Let $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{c}, T', 1)$, $T' = \{(t'_1, \ldots, t'_n)\}$, be such a complete information routing game with an associated mixed Nash equilibrium $\mathbf{Q}$. With a slight abuse of notation, we denote $\mathbf{Q} = (q(i,j))_{i\in[n],j\in[m]}$ where $q(i,j)$ is the probability that type $t'_i \in T'_i$ is assigned to link $j \in [m]$.

We define a CI-like game $\Gamma = (n, m, \mathbf{c}, T, \mathbf{p})$ and an associated pure strategy profile σ as follows:

For each user $i \in [n]$, T_i consists of $|\mathsf{support}_{\mathbf{Q}_i}(i)|$ types, where we have a type t^j_i for every link $j \in \mathsf{support}_{\mathbf{Q}_i}(i)$. For all users $i \in [n]$ and links $j \in \mathsf{support}_{\mathbf{Q}_i}(i)$, define $p(i, t^j_i) = q(i,j)$ and $\sigma_i(t^j_i) = j$.

Notice that our construction $\Gamma \mapsto \Gamma_{\mathsf{CI}}$ transforms the CI-like game Γ with associated pure strategy profile σ back to the complete information routing game Γ_{CI} with associated (mixed) Nash equilibrium $\mathbf{Q}$. It follows that $\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma) = \mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})$, $\mathsf{OPT}_{\mathsf{MSP}}(\Gamma) = \mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})$

and $\lambda^l_{(i,t_i^j)}(\sigma,\mathbf{p}) = \lambda^l_{(i,t_i')}(\mathbf{Q},1)$ for all users $i \in [n]$, for all links $l \in [m]$, and for all $j \in \mathsf{support}_{\mathbf{Q}_i}(i)$. Since $\mathbf{Q}$ is a Nash equilibrium we have

$$v_{(i,t_i')}(\mathbf{Q},1) = \lambda^j_{(i,t_i')}(\mathbf{Q},1) \quad \text{for all } j \in \mathsf{support}_{\mathbf{Q}_i}(t_i') \text{ and}$$
$$v_{(i,t_i')}(\mathbf{Q},1) \leq \lambda^j_{(i,t_i')}(\mathbf{Q},1) \quad \text{for all } j \notin \mathsf{support}_{\mathbf{Q}_i}(t_i').$$

Furthermore $\mathsf{support}_{\sigma_i}(i) = \mathsf{support}_{\mathbf{Q}_i}(i)$ for all $i \in [n]$, and $\lambda^l_{(i,t_i^j)}(\sigma,\mathbf{p}) = \lambda^l_{(i,t_i')}(\mathbf{Q},1)$ for all users $i \in [n]$, for all links $l \in [m]$, and for all $j \in \mathsf{support}_{\mathbf{Q}_i}(i)$. It follows that σ is a pure Bayesian Nash equilibrium with

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma,\Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} = \frac{\mathsf{SC}_{\mathsf{MSP}}(\mathbf{Q},\Gamma_{\mathsf{CI}})}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma_{\mathsf{CI}})} = \begin{cases} \Omega\left(\frac{\log m}{\log\log m}\right) & \text{, for the case of identical links,} \\ \Omega\left(\frac{\log m}{\log\log\log m}\right) & \text{, for the case of related links.} \end{cases}$$

This completes the proof. ■

We conclude with a lower bound on Coordination Ratio as Expected Maximum Latency for normal pure Bayesian Nash equilibria.

Theorem 5.5 *There exists a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links and an associated normal pure Bayesian Nash equilibrium σ such that*

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

Proof: Let $m \in \mathbb{N}$ be a perfect square. Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links with independent type distribution $\mathbf{p}$. There are two classes of users, $\mathcal{U}_1$ and $\mathcal{U}_2$:

- The class $\mathcal{U}_1$ consists of m users with type set $T_i = \{t_i, t_i'\}$, where $w(t_i) = 1, w(t_i') = 0, p(i,t_i) = \frac{1}{\sqrt{m}}$ and $p(i,t_i') = 1 - \frac{1}{\sqrt{m}}$ for all users $i \in \mathcal{U}_1$.
- The class $\mathcal{U}_2$ consists of $(\sqrt{m}-1)m$ users with type set $T_i = \{t_i\}$, where $w(t_i) = \frac{1}{\sqrt{m}}$ and $p(i,t_i) = 1$ all users $i \in \mathcal{U}_2$.

Consider the pure strategy profile σ' 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 σ', we get

$$\mathsf{SC}_{\mathsf{MSP}}(\sigma',\Gamma) \leq 1 + (\sqrt{m}-1)\cdot\frac{1}{\sqrt{m}} < 2.$$

Now consider the normal pure strategy profile σ 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. Clearly, σ is a normal pure Bayesian Nash equilibrium.

To show a lower bound on $\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma)$ we consider any link $j \in [\sqrt{m}]$. The actual load, say X_j, on link $j \in [\sqrt{m}]$ is a random variable which is a sum of $\sqrt{m}$ independent random variables with $\mathbb{E}(X_j) = 1$. Let $1 \leq k \leq \sqrt{m}, k \in \mathbb{N}$; the precise choice of k will be made later. Clearly,

$$\begin{aligned}
\Pr(X_j \geq k) &\geq \Pr(X_j = k) \\
&= \binom{\sqrt{m}}{k} \cdot \left(\frac{1}{\sqrt{m}}\right)^k \cdot \left(1 - \frac{1}{\sqrt{m}}\right)^{\sqrt{m}-k} \\
&\geq \binom{\sqrt{m}}{k} \cdot \left(\frac{1}{\sqrt{m}}\right)^k \cdot \frac{1}{\mathsf{e}} \qquad \text{(since } k \geq 1\text{)} \\
&= \frac{\sqrt{m} \cdot \ldots \cdot (\sqrt{m} - k + 1)}{\sqrt{m}^k} \cdot \frac{1}{k!} \cdot \frac{1}{\mathsf{e}}.
\end{aligned}$$

Now, observe that $\frac{\sqrt{m} \cdot \ldots \cdot (\sqrt{m} - k + 1)}{\sqrt{m}^k}$ is monotonically increasing in $\sqrt{m}$ and $\sqrt{m} \geq k$. Thus,

$$\frac{\sqrt{m} \cdot \ldots \cdot (\sqrt{m} - k + 1)}{\sqrt{m}^k} \geq \frac{k!}{k^k}.$$

It follows that

$$\begin{aligned}
\Pr(X_j \geq k) &\geq \frac{k!}{k^k} \cdot \frac{1}{k!} \cdot \frac{1}{\mathsf{e}} \\
&= \frac{1}{\mathsf{e} \cdot k^k},
\end{aligned}$$

so that

$$\Pr(X_j < k) \leq 1 - \frac{1}{\mathsf{e} \cdot k^k}.$$

Now, since the actual loads $X_1, \ldots, X_{\sqrt{m}}$ are independent of each other, we have

$$\begin{aligned}
\Pr((X_1 < k) \wedge \ldots \wedge (X_{\sqrt{m}} < k)) &= \prod_{j \in [\sqrt{m}]} \Pr(X_j < k) \\
&\leq \left(1 - \frac{1}{\mathsf{e} \cdot k^k}\right)^{\sqrt{m}} \\
&\leq e^{-\frac{1}{\mathsf{e} \cdot k^k} \cdot \sqrt{m}}.
\end{aligned}$$

Define now $\alpha > 0$ so that $\left(\frac{\alpha}{\mathsf{e}}\right)^\alpha = m$. Then, clearly, $\alpha = \Theta\left(\frac{\log m}{\log \log m}\right)$. Choose $k = \frac{\alpha}{\mathsf{e}}$. Then

$k^k = m^{\frac{1}{\mathsf{e}}}$ which implies

$$\begin{aligned}\Pr((X_1 < k) \wedge ... \wedge (X_{\sqrt{m}} < k)) &\leq e^{-\frac{1}{\mathsf{e}\cdot k^k}\cdot\sqrt{m}} \\ &= e^{-\frac{1}{\mathsf{e}}\cdot m^{\frac{1}{2}-\frac{1}{\mathsf{e}}}} \\ &\leq \frac{1}{m},\end{aligned}$$

for suitably large m. This implies that

$$\begin{aligned}\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma) &\geq \Pr((X_1 \geq k) \vee ... \vee (X_{\sqrt{m}} \geq k)) \cdot k \\ &= \Big(1 - \Pr((X_1 < k) \wedge ... \wedge (X_{\sqrt{m}} < k))\Big) \cdot k \\ &\geq \left(1 - \frac{1}{m}\right) \cdot \frac{\alpha}{\mathsf{e}} \\ &= \Theta\left(\frac{\log m}{\log\log m}\right).\end{aligned}$$

Thus,

$$\frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma)}{\mathsf{OPT}_{\mathsf{MSP}}(\Gamma)} \geq \frac{\mathsf{SC}_{\mathsf{MSP}}(\sigma, \Gamma)}{\mathsf{SC}_{\mathsf{MSP}}(\sigma', \Gamma)} = \Omega\left(\frac{\log m}{\log\log m}\right),$$

as needed. ■

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 Costs.

Proposition 4.3 implies that fully mixed Bayesian Nash equilibria have worst Social Cost as Sum of Individual Costs. Hence, we obtain:

Theorem 5.6 *Consider a Bayesian routing game Γ on identical links and an associated fully mixed Bayesian Nash equilibrium $\mathbf{F}$ and a Bayesian Nash equilibrium $\mathbf{Q}$. Then,*

$$\mathsf{SC}_{\mathsf{SUM}}(\mathbf{Q}, \Gamma) \quad \leq \quad \mathsf{SC}_{\mathsf{SUM}}(\mathbf{F}, \Gamma).$$

We now use Theorem 5.6 to prove an asymptotically tight bound on Coordination Ratio for the case of identical links (and Social Cost as Sum of Individual Costs).

Theorem 5.7 *Consider a Bayesian routing game $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ on identical links and an associated Bayesian Nash equilibrium $\mathbf{Q}$. Then,*

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

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

Proof: By Theorem 5.6, it suffices to prove the upper bound for a fully mixed Bayesian Nash equilibrium $\mathbf{F}$. Clearly,

$$\begin{aligned}\mathsf{SC}_{\mathsf{SUM}}(\mathbf{F},\Gamma) &= \sum_{i\in[n]} u_i(\mathbf{F},\mathbf{p}) \\ &= \sum_{i\in[n]} \left(\frac{W}{m} + \frac{m-1}{m}W(i)\right) \quad \text{(by Proposition 4.2)} \\ &= \frac{nW}{m} + \frac{m-1}{m}W \\ &= \frac{m+n-1}{m}W.\end{aligned}$$

On the other hand, $u_i(\mathbf{Q},\mathbf{p}) \geq W(i)$ for any user $i \in [s]$ and any strategy profile $\mathbf{Q}$; hence,

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

The upper bound follows.

We now prove that this upper bound is tight even for complete information routing games. To do so, we will prove that for any $\varepsilon > 0$, there is a complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, T, 1)$ such that $\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) \leq (1+\varepsilon)\cdot W$. We proceed by case analysis on the relation between n and m.

- Assume first that $n \leq m$. Let Γ_{CI} be an arbitrary complete information routing game with $n \leq m$. Then we can assign each user to a separate link which yields $\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) = W$.
- Assume now that $n > m$. Define the complete information routing game Γ_{CI} as follows:

  There are two sets of users $\mathcal{U}_1$, and $\mathcal{U}_2$. 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 $i \in \mathcal{U}_2$ where $k \in \mathbb{N}$ is a constant to be determined later.

  For the (expected) total traffic, we get

  $$W = n - m + 1 + (m-1)k.$$

  Let σ 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 $[m-1]$. Thus,

$$\begin{aligned}
\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) &\leq \mathsf{SC}_{\mathsf{SUM}}(\sigma, \Gamma_{\mathsf{CI}}) \\
&= (n-m+1)^2 + (m-1)k \\
&= \frac{(n-m+1)^2 + (m-1)k}{n-m+1+(m-1)k} \cdot W \\
&= \frac{(n-m+1)\cdot(n-m) + (n-m+1) + (m-1)\cdot k}{n+(m-1)(k-1)} \cdot W \\
&= \left(1 + \frac{(n-m)(n-m+1)}{n+(m-1)(k-1)}\right) \cdot W.
\end{aligned}$$

Clearly, for any $\varepsilon > 0$, there is a $k \in \mathbb{N}$ such that $\frac{(n-m)(n-m+1)}{n+(m-1)(k-1)} \leq \varepsilon$. Hence, for any $\varepsilon > 0$, there is a complete information routing game Γ_{CI} such that $\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) \leq (1+\varepsilon) \cdot W$. This completes the proof for the case $n > m$.

In all cases, there is a complete information routing game Γ_{CI} such that $\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}}) \leq (1+\varepsilon) \cdot W$. Since $\mathsf{SC}_{\mathsf{SUM}}(\mathbf{F}, \Gamma_{\mathsf{CI}}) = \frac{m+n-1}{m} W$, it follows that

$$\frac{\mathsf{SC}_{\mathsf{SUM}}(\mathbf{F}, \Gamma_{\mathsf{CI}})}{\mathsf{OPT}_{\mathsf{SUM}}(\Gamma_{\mathsf{CI}})} \geq \frac{1}{1+\varepsilon} \cdot \frac{m+n-1}{m},$$

as needed. ■

Berenbrink *et al.* [3] showed that the Coordination Ratio for complete information routing games and Social Cost as Sum of Individual Costs grows at least linearly with the number of users. In particular they proved that $\frac{n}{5}$ is a lower bound on the Coordination Ratio. Theorem 5.7 implies that the Coordination Ratio increases at most linear with n and also shows the impact of the number of links.

Another interesting insight of Theorem 5.7 is that the Coordination Ratio does *not* increase if we allow incomplete information. This is not the case if Social Cost is defined as the Maximum Individual Cost, as we will see next.

5.3 Social Cost as Maximum Individual Cost

In this section, we study the Coordination Ratio for Social Cost as the Maximum Individual Cost.

Proposition 4.3 implies that fully mixed Bayesian Nash equilibria have worst Social Cost as Maximum Individual Cost. Hence, we obtain:

Theorem 5.8 *Consider a Bayesian routing game* Γ *on identical links and an associated fully mixed Bayesian Nash equilibrium* $\mathbf{F}$ *and a Bayesian Nash equilibrium* $\mathbf{Q}$*. Then,*

$$\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q},\Gamma) \quad\leq\quad \mathsf{SC}_{\mathsf{MAX}}(\mathbf{F},\Gamma).$$

We now use Theorem 5.8 to prove asymptotically tight bounds on Coordination Ratio for the case of identical links.

Theorem 5.9 *Consider a Bayesian routing game* $\Gamma = (n, m, \mathbf{1}, T, \mathbf{p})$ *on identical links and an associated Bayesian Nash equilibrium* $\mathbf{Q}$*. Then,*

$$\text{(a)}\quad \frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q},\Gamma)}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} \leq \frac{m+n-1}{m}, \text{ and}$$

$$\text{(b)}\quad \frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q},\Gamma)}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} \leq 2-\frac{1}{m}, \qquad \text{if } \Gamma \text{ is a complete information routing game.}$$

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

Proof: Let $\mathbf{F}$ be a fully mixed Bayesian Nash equilibrium for Γ. Propositions 4.2 and 4.3 imply together that

$$u_i(\mathbf{Q},\mathbf{p}) \leq u_i(\mathbf{F},\mathbf{p}) = \frac{W}{m} + \frac{m-1}{m}W(i), \tag{4}$$

for each user $i \in [n]$. We consider the two cases from the theorem.

Case (a):

Upper bound: Clearly, for any strategy profile $\mathbf{Q}'$ and for any user $i \in [n]$, $u_i(\mathbf{Q}',\mathbf{p}) \geq W(i)$; hence, $\sum_{i\in[n]} u_i(\mathbf{Q}',\mathbf{p}) \geq W$. This implies that

$$\mathsf{OPT}_{\mathsf{MAX}}(\Gamma) \geq \frac{W}{n}. \tag{5}$$

Clearly, $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma) \geq W(i)$ for all $i \in [n]$. Fix any user $i \in [n]$. By (4) and (5),

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

and the upper bound follows.

Lower bound: Fix any arbitrary $k, a, r \in \mathbb{N}$, which will be determined later. Consider the Bayesian routing game $\Gamma_{k,a,r} = (n, m, \mathbf{1}, T, \mathbf{p})$ with independent type distribution and $n =$

$k \cdot (m-1)$ users. Each user $i \in [n]$ has type set $T_i = \{t_i, t'_i\}$ with traffics $w(t_i) = 1$, $w(t'_i) = a \cdot r$ and probabilities $p(i, t_i) = 1 - \frac{1}{a}$, $p(i, t'_i) = \frac{1}{a}$. Clearly, for user $i \in [n]$, $W(i) = r + 1 - \frac{1}{a}$.

Define a pure strategy profile σ 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 $[m-1]$; so, σ assigns exactly k of these types to each link in $[m-1]$. Now for each user $i \in [n]$,

$$\begin{aligned} u_i(\sigma, \mathbf{p}) &= \left(1 - \frac{1}{a}\right) \cdot \left(1 + (k-1) \cdot \left(1 - \frac{1}{a}\right)\right) + \frac{1}{a} \cdot ((n-1)r + r \cdot a) \\ &= \left(1 - \frac{1}{a}\right) \cdot \left(\frac{1}{a} + k \cdot \left(1 - \frac{1}{a}\right)\right) + r \cdot \left(\frac{(n-1)}{a} + 1\right); \end{aligned}$$

so, for any $\epsilon' > 0$, there is a sufficiently large a such that for each user $i \in [n]$,

$$u_i(\sigma, \mathbf{p}) \leq (k + r)(1 + \epsilon').$$

Hence, $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{k,a,r}) \leq (k+r)(1+\epsilon')$. Fix now any fully mixed Bayesian Nash equilibrium $\mathbf{F}$. Proposition 4.2 implies that for each user $i \in [n]$,

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

Thus $\mathsf{SC}_{\mathsf{MAX}}(\mathbf{F}, \Gamma_{k,a,r}) = \frac{m+n-1}{m} \cdot \left(r + 1 - \frac{1}{a}\right)$ and we can conclude that

$$\frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{F}, \Gamma_{k,a,r})}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{k,a,r})} \geq \frac{(r + 1 - \frac{1}{a})}{(k+r)(1+\epsilon')} \cdot \frac{m+n-1}{m};$$

so, for any $\epsilon > \epsilon'$, there is a sufficiently large r such that

$$\frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{F}, \Gamma_{k,a,r})}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{k,a,r})} \geq \frac{m+n-1}{m} \cdot \frac{1}{1+\epsilon}.$$

This proves that the upper bound shown before is tight up to a factor of $(1+\epsilon)$.

Case (b):

Upper bound: Consider the complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, \{(t_1, \ldots, t_n)\}, 1)$. Here, $W(i) = w(t_i)$ for all $i \in [n]$. Clearly, $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}) \geq W(i)$ for all $i \in [n]$ and $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}) \geq \frac{W}{m}$. By Equation (4),

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

so that

$$\frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{Q},\Gamma)}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} \leq 2-\frac{1}{m}$$

as needed. The upper bound follows.

Lower bound: Consider the complete information routing game $\Gamma_{\mathsf{CI}} = (n, m, \mathbf{1}, \{(t_1,\dots,t_n)\}, 1)$ with $n = m$ and $w(t_1) = \dots = w(t_n) = 1$. Clearly, $W(i) = w(t_i) = 1$ for all $i \in [n]$, $W = m$ and $\mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}) = 1$. Now, for the fully mixed Nash equilibrium $\mathbf{F}$ and any user $i \in [n]$, by Equation (4),

$$\begin{aligned} u_i(\mathbf{F},\mathbf{p}) &= \frac{W}{m} + \frac{m-1}{m}W(i) \\ &= (2-\frac{1}{m})\cdot \mathsf{OPT}_{\mathsf{MAX}}(\Gamma_{\mathsf{CI}}), \end{aligned}$$

so that

$$\frac{\mathsf{SC}_{\mathsf{MAX}}(\mathbf{F},\Gamma)}{\mathsf{OPT}_{\mathsf{MAX}}(\Gamma)} = 2-\frac{1}{m},$$

as needed. ■

6 Conclusions and Discussion

We have introduced the dimension of incomplete information into the class of routing games on parallel links. In this setting, we have studied the existence and computational complexity of pure Bayesian Nash equilibria, structural properties of fully mixed Bayesian Nash equilibria and the Coordination Ratio for different Social Cost measures.

Our work leaves open several interesting problems. On the most concrete level, we would like to ask:

- Can pure Bayesian Nash equilibria be computed in polynomial time?
- What is the exact value of Coordination Ratio for identical links if Social Cost is defined as Expected Maximum Latency?
- What is the Coordination Ratio for all three considered Social Cost measures in the case of related links?

Acknowledgment. We would like to thank Thomas Lücking, Marios Mavronicolas, and Florian Schoppmann for many fruitful discussions and helpful comments. We would also like to thank the anonymous referee for many valuable comments on improving the manuscript.

References

[1] B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 57–66, 2005.

[2] 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*, pages 746–755, 2004.

[3] P. Berenbrink, L. A. Goldberg, P. W. Goldberg, and R. Martin. Utilitarian Resource Assignment. *Journal of Discrete Algorithms*, (http://dx.doi.org/10.1016/j.jda.2005.06.009), to appear.

[4] J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Second Edition. Springer Verlag, 1980.

[5] 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.

[6] 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. Also accepted to *Journal of Algorithms* as Special Issue of SODA'02.

[7] 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*, pages 604–612, 2004.

[8] G. Facchini, F. van Megan, P. Borm, and S. Tijs. Congestion Models and Weighted Bayesian Potential Games. *Theory and Decision*, 42:193–206, 1997.

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

[10] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish Routing in Non-Cooperative Networks: A Survey. In B. Rovan and P. Vojtás, editors, *Proceedings of the*

28th International Symposium on Mathematical Foundations of Computer Science, Lecture Notes in Computer Science, Vol. 2747, Springer Verlag, pages 21–45, 2003.

[11] S. Fischer and B. Vöcking. On the Structure and Complexity of Worst-Case Equilibria. In X. Deng and Y. Ye, editors, *Proceedings of the 1st International Workshop on Internet and Network Economics*, Lecture Notes in Computer Science, Vol. 3828, Springer Verlag, pages 151–160, 2005.

[12] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In P. Widmayer, F. T. Ruiz, R. M. Bueno, M. Hennessy, S. Eidenbenz, and R. Conejo, editors, *Proceedings of the 29th International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 2380, Springer Verlag, pages 123–134, 2002.

[13] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. In J. Diaz, J. Karhumäki, A. Lepistö, and D. Sannella, editors, *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 3142, Springer Verlag, pages 593–605, 2004.

[14] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. In J. Diaz, J. Karhumäki, A. Lepistö, and D. Sannella, editors, *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 3142, Springer Verlag, pages 645–657, 2004.

[15] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and P. Spirakis. Structure and Complexity of Extreme Nash Equilibria. *Theoretical Computer Science*, 343:133–157, 2005.

[16] M. Gairing, T. Lücking, B. Monien, and K. Tiemann. Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture. In *Proceedings of the 32nd International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science, Vol. 3580, Springer Verlag, pages 51–65, 2005.

[17] C. Georgiou, T. Pavlides, and A. Philippou. Network Uncertainty in Selfish Routing. In *Proc. of the 20th IEEE International Parallel & Distributed Processing Symposium*, page 105, 2006.

[18] R. L. Graham. Bounds on Multiprocessing Timing Anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.

[19] J. C. Harsanyi. Games with Incomplete Information Played by Bayesian Players, I, II, III. *Management Science*, 14:159–182, 320–332, 468–502, 1967.

[20] J. C. Harsanyi. Games with Randomly Disturbed Payoffs. *Internation Journal on Game Theory*, 21:1–23, 1973.

[21] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate Equilibria and Ball Fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.

[22] E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In C. Meinel and S. Tison, editors, *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Vol. 1563, Springer Verlag, pages 404–413, 1999.

[23] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. In V. Diekert and M. Habib, editors, *Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Vol. 2996, Springer Verlag, pages 547–558, 2004.

[24] T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the Worst-Case Nash Equilibrium? In B. Rovan and P. Vojtás, editors, *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science*, Lecture Notes in Computer Science, Vol. 2747, Springer Verlag, pages 551–561, 2003.

[25] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory.* Oxford University Press, 1995.

[26] M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 510–519, 2001.

[27] I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

[28] D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[29] R. B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, 1997.

[30] C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 749–753, 2001.

[31] R. W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[32] 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.