The Price of Selfish Routing

Marios Mavronicolas*
Dept. of Computer Science
University of Cyprus
Nicosia CY-1678, Cyprus
mavronic@ucy.ac.cy

Paul Spirakis†
Dept. of Computer Engineering and Informatics
University of Patras
& Computer Technology Institute
261 10 Patras, Greece
spirakis@cti.gr

ABSTRACT

We study the problem of *routing* traffic through a congested network. We focus on the simplest case of a network consisting of m parallel *links.* We assume a collection of n network *users,* each employing a *mixed strategy* which is a probability distribution over links, to control the shipping of its own assigned traffic. Given a *capacity* for each link specifying the rate at which the link processes traffic, the objective is to route traffic so that the maximum expected *latency* over all links is minimized. We consider both *uniform* and *non-uniform* link capacities.

How much decrease in global performace is necessary due to the absence of some central authority to regulate network traffic and implement an optimal assignment of traffic to links? We investigate this fundamental question in the context of *Nash equilibria* for such a system, where each network user selfishly routes its traffic only on those links available to it that minimize its expected latency cost, given the network congestion caused by the other users. We use the *coordination ratio,* defined by Koutsoupias and Papadimitriou [25] as the ratio of the maximum (over all links) expected latency in the *worst* possible Nash equlibrium, over the *least* possible maximum latency had global regulation been available, as a measure of the cost of lack of coordination among the network users.

Our point of departure is a set of combinatorial *minimum expected latency cost equations,* one per network user, that characterize the Nash equilibria of this system. These are linear equations in the minimum expected latencies, involving the users' traffics, the link capacities and the routing pattern determined by the mixed strategies. In turn, we solve these equations in the case of *fully mixed strategies,* where each user assigns its traffic with a non-zero probability to every link, to derive the first *existence* and *uniqueness* results for Nash equilibria in this setting. Most importantly, we use the derived characterizations of Nash equilibria to show, under the assumption of fully mixed strategies, *tight* upper bounds of no worse than $O(\ln n / \ln \ln n)$ on the coordination ratio for *(i)* the case of uniform link capacities and arbitrary traffics, and *(ii)* the case of non-uniform link capacities and identical traffics.

*Supported by funds from the Joint Program of Scientific and Technological Collaboration between Greece and Cyprus, and by funds for the promotion of research at University of Cyprus.

†Supported in part by European Union's ESPRIT/IST Long Term research projects ALCOM-FT (contract # 20244), FET-OPEN, IMPROVING RTN ARACNE, by the Greek General Secretariat for Research and Technology, by the Greek Ministry of Education, and by funds from the Joint Program of Scientific and Technological Collaboration between Greece and Cyprus.

STOC'01 July 6-8, 2001, Hersonissos, Crete, Greece.

1. INTRODUCTION

1.1 Motivation-Framework

We study a *routing* problem in communication networks; in this problem, paths from a *source* to a *destination* are to be established by a collection of *non-cooperating* entities, which we call *users.* Thus, users correspond to different traffic sources, each seeking to determine the shipping of its own traffic over a shared network. However, in doing so, different users may have to optimize completely different (and even conflicting) measures of performance and demand. Such networks are henceforth called *noncooperative* (cf. [12, 21, 24, 27, 32]).

Such noncooperative and antagonistic scenaria apply to various modern networking environments, where a single performance objective, which is regulated via some global-control mechanism, is no longer a valid assumption. For example, the Internet Protocol (both IPv4 and the current IPv6 specification [10]) provides the option of *source routing* that enables the user to determine the path its traffic follows from source to destination. Another example is the *flexible routing service* specified in the Q.1211 CCITT Recommendation for the standardized capability set of Intelligent Networks [8]; a goal of this service is to route calls over particular facilities based on the subscriber's preference list.

A natural framework in which to study such multiobjective optimization problems with local payoffs in a noncooperative network is (noncooperative) *game theory* [6, 33, 34]. An appropriate, game-theoretic concept for the solution is *Nash equilibrium* [31]. Roughly speaking, the operating points of a noncooperative network are the Nash equilibria of the underlying game; these are points where unilateral deviation does not help any user to improve its performance. Game-theoretic models, concepts and techniques have been employed recently in the context of various networking problems such as *flow control* [1, 16, 20, 37], *routing* [2, 12, 21, 22, 23, 25, 27, 32, 36], *bandwidth allocation* [26], *Web access* [35], *multicasting* [13], and *congestion control* [18]. Moreover, applications of game theory to the entire discipline of computer science have attracted recently a lot of interest and attention, and have become a currently major trend.

The mission of this work is to study, within the game-theoretic framework, the inherent costs due to the lack of a central authority to monitor and regulate network operation according to global objectives. More specifically, we adopt *coordination ratio* as the measure of performance loss due to the lack of coordination; this measure was introduced in a very recent paper by Koutsoupias and Papadimitriou [25]. Roughly speaking, the coordination ratio is the ratio between the *social cost* (specifically, the *maximum expected latency* in the setting we consider) in the *worst* possible Nash equlibrium, and the *social optimum,* which is the best "offline" global cost (specifically, the *minimum latency* in our setting) had all information been available to a central network authority regulating traffic.

We follow Koutsoupias and Papadimitriou [25] to continue the study of the simplest case of a network consisting of m parallel *links*. Systems of parallel links, albeit simple, represent an appropriate model for several, diverse networking problems. Consider, for example, broadband networks where bandwidth is preallocated to different virtual paths that do not interfere; thus, these paths result effectively in a system of parallel links between source/destination pairs. As a second example, consider a multimedia network with several servers that are shared by the network customers; each customer distributes its applications among the servers, while competing with the other customers on the common available resources. Modeling each server as a link, the parallel links model considered in our study fits well such a framework.

In the model of parallel links we consider, each of n users fixes a *mixed strategy,* which is a probability distribution over links; the distribution determines the (possibly zero) probability for the user to ship its traffic through each link. We model the *latency* over each link as the ratio of the total traffic assigned to the link over its *capacity.* Thus, in this setting, Nash equilibrium requires that for each specific user, the *expected latency* is constant across all links that are potential bearers of the user's traffic, given the congestion caused by other users.

1.2 Summary of Results

Our point of departure is a linear system of equations for the minimum expected latency costs (Proposition 8), which we call *minimum expected latency cost equations.* These equations are specific to Nash equilibria of the system and they are inspired by (and reminiscent of) classical equations of stochastic equilibrium such as the *Chapman-Kolmogorov equation* that describes the steady-state equilibrium of a Markov chain (cf. [14, Section 6.1]). The coefficients and constant terms of the minimum expected latency cost equations depend on the link and user parameters and on the routing pattern (which user uses which links); the equations hold for *any* routing pattern.

For the rest of our study, we focus on a special kind of routing pattern, that we call *fully mixed strategies*; here, each user assigns *non-zero* probability to each and every link. For fully mixed strategies, we are able to explicitly solve the minimum expected latency equations and derive the user probabilities in a Nash equilibrium, henceforth refered to as *Nash probabilities.* We discover that the Nash probabilities enjoy a particularly nice form, as functions of the link capacities and the user traffics (Proposition 11), in the case of fully mixed strategies.

The requirement that the computed values for the Nash probabilities do indeed represent probabilities yields the *first* inexistence result for Nash equilibrium (Corollary 12) in this setting. In addition, we prove that these probabilities are also sufficient to give rise to a Nash equilibrium (Proposition 13), culminating in the first existence and uniqueness result for Nash equilibria in the setting we consider (Theorem 14). Our study reveals a rich structure for Nash equilibria in the case of fully mixed strategies.

How useful is this understanding of the structure of Nash equilibria for fully mixed strategies? We credit its usefulness by using it to derive some new, improved bounds on coordination ratio in a number of significant cases.

- We first consider the case where capacities are uniform, while user traffics are allowed to vary arbitrarily. For this case, we observe that all links are *equiprobable* for each user. This fact allows us to use simple results from the classical theory of *random allocations,* where each of n *balls* is put into one of m *bins,* chosen uniformly at random; properties of such random allocation processes have been studied extensively in the probability and statistics literature (see, e.g., [17, 19]). Thus, in our study, we treat users and links as balls and bins, respectively.

- We also consider the case where all user traffics are identical, but capacities may now vary arbitrarily, subject, however, to the constraints that are necessary for a Nash equilibrium to exist (Corollary 12). Although, in this case, links are no longer equiprobable for the same user, the constraints imposed by Nash equilibrium still allow us to benefit from envisioning our problem as an occupancy problem.

To prove our bounds on coordination ratio, we develop a modular methodology that may be applicable to other instances of the problem we consider, and even in other settings and for different performance measures as well. This methodology consists of three major components.

- The first one establishes a probabilistic *tail lemma,* showing that if the number of users choosing any particular link enjoys (as a random variable) a sharp concentration around its expectation (this fact being properly formalized), then an upper bound on social cost holds that depends also on parameters specifying the sharpness of the concentration (see Lemmas 16 and 23 for the two cases we consider, respectively).

- The second component quantitatively determines this sharpness of concentration; it employs the form of Nash probabilities, the constraints on system parameters (capacities and traffics) for a Nash equilibrium to exist, and standard tail inequalities for occupancy problems (cf. [19, 30]).

Thus, these two components together establish a concrete upper bound on social cost for each specific case.

- The third component complements the first two by showing a *lower* bound on social optimum, thus implying a concrete upper bound on coordination ratio for each specific case. (See Lemmas 19 and 25 for the two cases we consider, respectively.)

The concrete upper bounds on coordination ratio obtained by following this methodology are as follows:

- For the case of uniform capacities and arbitrary traffics:
  - Assuming that $m = n$, we prove an upper bound of $(e \ln n)/(\ln \ln n) + 1$ (Theorem 20).
  - Assuming that $m \leq n/(24 \ln n)$, we prove an upper bound of $(3/2+\varepsilon)$ times the ratio of the maximum over the minimum traffics, for any arbitrary constant $\varepsilon > 0$ and for sufficiently large values of m and n (Theorem 21).
- For the case of non-uniform capacities and identical traffics, and assuming that $m \leq n$, we prove an upper bound of $(2e \ln n)/(\ln \ln n) + 2(2e + 1)$ (Theorem 26).

Our first and third bounds on coordination ratio match a corresponding asymptotic lower bound of $\Omega(\ln n/ \ln \ln n)$ shown by Koutsoupias and Papadimitriou [25, Theorem 6] and conjectured by them to be the precise bound. On the other hand, our second bound identifies the *first* conditions on (non-constant) m and n allowing for a bound independent of m and n. Our first and second bounds surpasses a *general* upper bound of $O(\sqrt{m \ln m}$ shown by Koutsoupias and Papadimitriou [25, Theorem 8] for the model of uniform capacities; that bound, however, holds for *all* possible Nash equilibria, while our bounds hold for the fully mixed case and under particular assumptions on n and m. In the same vein, our third bound surpasses a corresponding upper bound of Koutsoupias and Papadimitriou [25, Theorem 9] for the model of non-uniform capacities.

The contribution of our work is two-fold. First, it offers a substantially deep understanding of the structure of Nash equilibria, especially for the case of fully mixed strategies. Second, it integrates the benefits of this understanding into a modular methodology for proving upper bounds on coordination ratio, hereby yielding improved and tight (within small constants) upper bounds for several interesting cases. We believe that our structural findings and proposed methodology will be instrumental to settling other problem instances as well, and even applicable and extendible to other settings and performance measures.

1.3 Related Work

Our work continues and complements the recent work of Koutsoupias and Papadimitriou [25]; that work formulated selfish routing as a noncooperative game and initiated the study of performance degradation caused by a lack of regulation in a congested network. Koutsoupias and Papadimitriou [25] focused too on the network consisting of m parallel links and obtained tight bounds on coordination ratio for the case where $m = 2$ and less tight ones for the general case (under both uniform and non-uniform capacities). Some additional bounds for the cases $m = 2$ and 3 were proved by Mavronicolas *et al.* [28].

The very recent paper of Roughgarden and Tardos [36] is very similar in motivation to our paper, investigating the degradation in network performance due to unregulated traffic. However, the model considered in [36] is different than the one considered here (and introduced in [25]) in that it assumes *pure strategies* for the users (as opposed to mixed ones); moreover, the global objective adopted in [36] is to minimize the *total* latency (as opposed to the expected latency on the most congested link that we adopt, following [25]). Roughgarden and Tardos [36] show that in their model (and assuming linear latency functions), the total latency in a Nash equilibrium is no worse than 4/3 times the *optimal* total latency; they also prove that this upper bound is tight.

The questions of existence, uniqueness, efficiency and computation of Nash equilibria for selfish (noncooperative) routing have been investigated over various settings in the networking literature (see, e.g., [2, 11, 12, 20, 21, 22, 23, 27, 32]). For more general results on the computation of Nash equilibria, see the survey [29]. The general inefficiency of Nash equilibria is discussed in [11].

1.4 Organization

The rest of this paper is organized as follows. Section 2 introduces the network model we consider, summarizes some background material, and establishes some preliminary facts. Section 3 outlines some material on random allocations. Our results on the structure of Nash equilibria appear in Section 4. The case of fully mixed strategies is treated in Sections 5 and 6 under the models of uniform capacities and non-uniform capacities, respectively. We conclude, in Section 7, with a discussion of our results and suggestions for further research.

2. DEFINITIONS AND PRELIMINARIES

2.1 Notation and Facts

For any integer $m \geq 2$, denote $[m] = \{1, \ldots, m\}$. For a real interval (a, b) and a real $\delta > 0$, $(a, b) + \delta$ denotes the real interval $(a + \delta, b + \delta)$. For all integers $m \geq 2$ and $n \geq 2$, denote $\mathbf{J}_{m \times n}$ the matrix with all entries in its m rows and n columns equal to 1; denote $\mathbf{I}_{n \times n}$ the *identity matrix* with n rows and n columns, all of its entries vanish except for those on the main diagonal that are equal to 1. For a vector $\mathbf{w}$ with positive entries, denote $\mathsf{max/min}(\mathbf{w})$ the ratio of the maximum over the minimum entry. Denote e the base of the natural logarithm. We shall use a combinatorial inequality stating that for any sufficiently large integer n and for an integer $\vartheta \leq n$, $\binom{n}{\vartheta} \leq (ne/\vartheta)^{\vartheta}$. For an event E in a sample space, denote $\mathbf{Pr}(E)$ the probability of event E happening. For a random variable X, denote $\mathcal{E}(X)$ the *expectation* of X.

2.2 Model

Our model and presentation are patterned after those in [25, Sections 1 & 2].

We consider a *network* consisting of a set of m parallel *links* $1, 2, \ldots, m$ from a *source* node to a *destination* node. Each of n *network users* $1, 2, \ldots, n$, or *users* for short, wishes to route a particular amount of traffic along a (non-fixed) link from source to destination; denote w_i the *traffic* of user i, $i \in [n]$. Define the $n \times 1$ *traffic vector* $\mathbf{w}$ in the natural way.

A *pure strategy* for user i is some specific link; a *mixed strategy* for user i is a probability distribution on the set of links. Throughout, we will be using subscripts for users and superscripts for links. A *set of pure strategies,* one per user, is represented by an n-tuple $\langle \ell_1, \ell_2, \ldots, \ell_n \rangle \in [m]^n$; a *set of mixed strategies,* one per user, is represented by an $m \times n$ *probability matrix* $\mathbf{P}$ of mn probabilities p_i^ℓ, $i \in [n]$ and $\ell \in [m]$, where p_i^ℓ is the probability that user i selects link ℓ. For a probability matrix $\mathbf{P}$, define *indicator variables* $I_i^\ell \in \{0,1\}$, $i \in [n]$ and $\ell \in [m]$, such that $I_i^\ell = 1$ if and only if $p_i^\ell > 0$. In the *fully mixed case,* $I_i^\ell = 1$ for all users i and links ℓ; here, each user assigns its traffic on each link with non-zero probability.

A *solo link* is a link ℓ such that $\sum_{k=1}^n I_k^\ell = 1$. By definition of indicator variables, it follows that there is a single user $s(\ell)$ such that $I_k^\ell = 1$ if $k = s(\ell)$, while $I_k^\ell = 0$ otherwise. Roughly speaking, the solo link ℓ can be traversed only by user $s(\ell)$. Denote $\mathcal{S}$ the set of solo links; clearly, $\mathcal{S} \subseteq [m]$. A *non-solo link* is a link that is not solo. For each link ℓ, define the random variable θ_ℓ to be the number of users that choose link ℓ.

Denote c_ℓ the *capacity* of link ℓ, representing the rate at which the link processes traffic. So, the *latency* for traffic w through link ℓ equals w/c_ℓ. In the model of *uniform capacities,* all link capacities are equal to c, for some constant $c > 0$; link capacities may vary arbitrarily in the model of *non-uniform capacities.* Define the $m \times n$ *capacity matrix* $\mathbf{C}$ with all entries in row ℓ equal to c_ℓ.

For any set of pure strategies $\langle \ell_1, \ell_2, \ldots, \ell_n \rangle$, the *latency cost for user i,* denoted λ_i, is $(\sum_{k:\ell_k=\ell_i} w_k)/c_{\ell_i}$; that is, the latency cost for user i is the latency of the link it chooses.

For any set of mixed strategies $\mathbf{P}$, denote W^ℓ the *expected traffic* on link ℓ; clearly, $W^\ell = \sum_{i=1}^n p_i^\ell w_i$. Given $\mathbf{P}$, define the $m \times 1$ *expected traffic vector* $\mathbf{W}$ in the natural way. For any set of mixed strategies $\mathbf{P}$, the *expected latency cost* for user i on link ℓ, denoted λ_i^ℓ, is the expectation, over all random choices of the remaining users, of the latency cost for user i when its traffic is assigned to link ℓ; thus,

$$\begin{aligned}
\lambda_i^\ell &= \frac{w_i + \sum_{k=1, k\neq i}^n p_k^\ell w_k}{c_\ell} \\
&= \frac{w_i - p_i^\ell w_i + \sum_{k=1}^n p_k^\ell w_k}{c_\ell} \\
&= \frac{(1 - p_i^\ell) w_i + W^\ell}{c_\ell}.
\end{aligned}$$

For each user i, the *minimum expected latency cost,* denoted λ_i, is the minimum, over all links ℓ, of the expected latency cost for user i on link ℓ; thus, $\lambda_i = \min_{\ell \in [m]} \lambda_i^\ell$. For a probability matrix $\mathbf{P}$, define the $n \times 1$ *minimum expected latency cost vector* λ in the natural way.

We are interested in a special class of mixed strategies called Nash equilibria [31] that we describe below. Formally, the probability matrix $\mathbf{P}$ is a *Nash equilibrium* if for all users i and links ℓ, $\lambda_i^\ell = \lambda_i$ if and only if $I_i^\ell > 0$. Thus, each user i assigns its traffic with non-zero probability only on links (possibly more than one of them) on which its expected latency cost is minimized; this implies that there is no incentive for user i to unilaterally deviate from its mixed strategy in order to avoid links on which its expected latency cost is higher than necessary. Call *Nash probabilities* the probabilities in a Nash equilibrium.

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

$$\begin{aligned}
\mathsf{SC}(\mathbf{w}, \mathbf{P}) &= \mathcal{E}\left(\max_{\ell \in [m]} \frac{\sum_{k:\ell_k=\ell} w_k}{c_\ell}\right) \\
&= \sum_{\langle \ell_1, \ell_2, \ldots, \ell_n \rangle \in [m]^n} \left(\prod_{k=1}^n p_k^{\ell_k} \cdot \max_{\ell \in [m]} \frac{\sum_{k:\ell_k=\ell} w_k}{c_\ell}\right).
\end{aligned}$$

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

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

The *coordination ratio* [25], denoted CR, is the *maximum* value, over all traffic vectors $\mathbf{w}$ and Nash equilibria $\mathbf{P}$, of the ratio $\mathsf{SC}(\mathbf{w}, \mathbf{P})/\mathsf{SO}(\mathbf{w})$.

2.3 Properties of Nash Equilibria

Koutsoupias and Papadimitriou [25, Section 2] provide necessary conditions for Nash equilibria.

Proposition 1 ([25]). *In a Nash equilibrium, for any user $i \in [n]$ and link $\ell \in [m]$, $p_i^\ell = \frac{W^\ell + w_i - c_\ell \lambda_i}{w_i}$, subject to*

(1) for all links $\ell \in [m]$, $W^\ell = \sum_{k=1}^n I_k^\ell (W^\ell + w_k - c_\ell \lambda_k)$, and

(2) for all users $i \in [m]$, $w_i = \sum_{j=1}^m I_i^j (W^j + w_i - c_j \lambda_i)$.

We remark that the necessary conditions in Proposition 1 neither provide any apparent way of computing Nash equilibria nor say anything about their existence and uniqueness. It appears that existence and uniqueness are contingent upon the corresponding existence and uniqueness of solutions for $\mathbf{W}$ and λ to the conditions (1) and (2). However, we observe here that simple expressions for expected traffics on non-solo links may be easily derived, implying their existence and uniqueness in a Nash equilibrium. The expressions are derived by appealing to Proposition 1 (Condition (1)). We prove:

Lemma 2. *In a Nash equilibrium, for any non-solo link $\ell \in [m]$, $W^\ell = \frac{-\sum_{k=1}^n I_k^\ell w_k + c_\ell \sum_{k=1}^n I_k^\ell \lambda_k}{\sum_{k=1}^n I_k^\ell - 1}$.*

2.4 An Exact Lower Bound

We conclude this section by establishing a simple lower bound on coordination ratio for the fully mixed case and under the model of uniform capacities. Unlike the upper bounds that we will show in Sections 5 and 6, which are asymptotic, this bound is *exact*; it is shown from first principles.

PROPOSITION 3. *Consider the fully mixed case under the model of uniform capacities. Then,*

$$\mathsf{CR} \geq m - \frac{1}{m^m}\sum_{\vartheta=1}^{m-1}\left(\sum_{\theta_1,\theta_2,\ldots,\theta_m \leq \vartheta}\binom{m}{\theta_1,\theta_2,\ldots,\theta_m}\right).$$

PROOF. We give an informal outline of our proof. We will set $n = m$ and assume unit traffics. We will also set $c = 1$. We notice that for this particular choice of traffics, $\mathsf{SO} = 1$. Hence, the corresponding social cost is a lower bound on coordination ratio. We use the fact that social cost is the expectation of the random variable $\max_{\ell\in[m]}\theta_\ell$ and apply sum telescoping to derive the claim. $\square$

We emphasize that Proposition 3 yields an *exact* (as opposed to asymptotic) lower bound on the coordination ratio for *any* particular value of m. For example, for $m = 2$ and $m = 3$, it yields lower bounds of $3/2$ and $51/27 \approx 1.889$, respectively, on the coordination ratio. (The lower bound of $3/2$ for $m = 2$ was shown before by Koutsoupias and Papadimitriou [25, Theorem 1].)

3. RANDOM ALLOCATIONS

In this section, we outline some material on random allocations that will be used in our later analysis.

Say that a discrete random variable X follows the *binomial distribution with parameters n and p* if for each integer ϑ, $0 \leq \vartheta \leq n$, $\mathbf{Pr}(X=\vartheta) = \binom{n}{\vartheta}p^{\vartheta}(1-p)^{n-\vartheta}$; we will use the fact that $\mathcal{E}(X) = np$. We shall use concepts and tools from the classical theory of *random allocations* (see, e.g., [17, 19]), studying the size of the fullest *bin* when each of n *balls* is independently put into one of m bins, which is chosen according to some specific probability distribution. We shall exploit arising analogies between selfish routing and random allocation problems, and we shall interchangeably use the terms balls and users, and bins and boxes, respectively. For any bin $\ell \in [m]$, denote θ_ℓ the random variable representing the number of balls put into it. In the special case where all users choose link $\ell \in [m]$, $\ell \in [m]$, with the same probability α_ℓ, each random variable θ_ℓ, $\ell \in [m]$ may be cast as a sum of n independent but identical *Bernoulli trials,* each representing the choice made by each specific ball; thus, in this case, θ_ℓ follows the binomial distribution with parameters n and α_ℓ.

For the case where $m = n$ and each ball chooses a bin uniformly at random, so that $\alpha_\ell = 1/m$ for all bins $\ell \in [m]$, a classical tool from the theory of random allocations shows that, with high probability, the size of the fullest bin does not exceed $\Theta(\ln n/\ln\ln n)$.

LEMMA 4. *Assume each of n balls is put uniformly at random into one of n bins. Then,* $\mathbf{Pr}\left(\max_{\ell\in[n]}\theta_\ell \leq \frac{e\ln n}{\ln\ln n}\right) \geq 1 - \frac{1}{n}$.

For the general case, where arbitrary probabilities α_ℓ, $\ell \in [m]$, are allowed and $m \neq n$, classical Chernoff-type results [9] bound the *tail probability* of θ_ℓ, which is the probability that θ_ℓ exceeds a suitable fraction of its expectation. We will use a particular such result derived by Angluin and Valiant [3].

LEMMA 5. *For any link $\ell \in [m]$, for any parameter $\beta \in (0,1)$,* $\mathbf{Pr}(\theta_\ell > (1+\beta)\mathcal{E}(\theta_\ell)) \leq \exp\left(-\frac{\beta^2}{3}\cdot\mathcal{E}(\theta_\ell)\right)$.

We conclude this section with two technical lemmas that will be used in our later proofs. These lemmas derive bounds on tail probability under some particular assumptions on expectation. We start by proving:

LEMMA 6. *Consider any link $\ell \in [m]$ such that $\mathcal{E}(\theta_\ell) \leq 1$. Then,* $\mathbf{Pr}\left(\theta_\ell > \frac{e\ln n}{\ln\ln n}\right) < \frac{1}{n^2}$.

We also prove:

LEMMA 7. *Consider any link $\ell \in [m]$ such that $\mathcal{E}(\theta_\ell) > 1$. Then,* $\mathbf{Pr}\left(\theta_\ell > \frac{e\ln n}{\ln\ln n}\mathcal{E}(\theta_\ell)\right) < \frac{2e\,\alpha_\ell}{n}$.

4. STRUCTURE OF NASH EQUILIBRIA

This section is organized as follows. Section 4.1 derives the minimum expected latency cost equations. These equations are used in Section 4.2 for establishing existence and uniqueness results for Nash equilibria in the fully mixed case.

4.1 Minimum Expected Latency Cost Equations

We show:

PROPOSITION 8. *In a Nash equilibrium, for any user $i \in [n]$,*

$$\begin{aligned}
&\lambda_i\left(\sum_{j=1}^{m} I_i^j c_j - \sum_{j=1, j\notin\mathcal{S}}^{m}\frac{I_i^j}{\sum_{k=1}^{n} I_k^j - 1}c_j\right) \\
&- \sum_{k=1,k\neq i}^{n}\lambda_k\left(\sum_{j=1,j\notin\mathcal{S}}^{m}\frac{I_i^j I_k^j}{\sum_{k=1}^{n} I_k^j - 1}c_j\right) \\
= \;& w_i\left(\sum_{j=1}^{m} I_i^j - 1 - \sum_{j=1,j\notin\mathcal{S}}^{m}\frac{I_i^j}{\sum_{k=1}^{n} I_k^j - 1}\right) \\
&- \sum_{k=1,k\neq i}^{n} w_k\left(\sum_{j=1,j\notin\mathcal{S}}^{m}\frac{I_i^j I_k^j}{\sum_{k=1}^{n} I_k^j - 1}\right) + \sum_{j=1,j\in\mathcal{S}}^{m} I_i^j W^j.
\end{aligned}$$

PROOF. We give an informal outline of our proof. Fix any user $i \in [n]$. We derive two alternative expressions for $\sum_{j=1}^m I_i^j W^j$. The first $\sum_{j=1}^m I_i^j W^j = w_i\left(1 - \sum_{j=1}^m I_i^j\right) + \lambda_i\sum_{j=1}^m I_i^j c_j$ follows directly from Proposition 1 (Condition (2)). The second one will follow by using the expressions for the expected traffics that were derived in Lemma 2. Equating the two derived expressions eliminates the expected traffics W^j, $1 \leq j \leq m$ and $j \notin \mathcal{S}$, on non-solo links and yields through algebraic manipulation the minimum expected latency equations. $\square$

4.2 Fully Mixed Strategies

We now focus on the fully mixed case. (Recall that there are no solo links in the fully mixed case.) We set $I_i^\ell = 1$ for all users $i \in [n]$ and links $\ell \in [m]$ in the minimum expected latency cost equations (Proposition 8) and solve the resulting linear system to obtain:

LEMMA 9. *In the fully mixed case, in a Nash equilibrium, λ is a linear transformation of* $\mathbf{w}$*:*

$$\lambda = \frac{1}{\sum_{j=1}^{m} c_j} \begin{pmatrix} m & 1 & \dots & 1 \\ 1 & m & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & m \end{pmatrix} \cdot \mathbf{w} .$$

We plug the expressions for the minimum expected latency costs derived in Lemma 9 into the expressions for the expected traffics on links provided by Lemma 2 to obtain:

LEMMA 10. *In the fully mixed case, in a Nash equilibrium,* $\mathbf{W}$ *is a linear transformation of* $\mathbf{w}$*:*

$$\mathbf{W} = \frac{1}{n-1} \left(-\mathbf{J}_{m\times n} + + \frac{m+n-1}{\sum_{j=1}^{m} c_j} \mathbf{C} \right) \cdot \mathbf{w} .$$

We are now ready to derive expressions for the Nash probabilities in the fully mixed case. We prove:

PROPOSITION 11. *Consider the fully mixed case in Nash equilibrium. Then, for all users $i \in [n]$ and link $\ell \in [m]$,*

$$p_i^\ell = \left(1 - \frac{m c_\ell}{\sum_{j=1}^{m} c_j}\right) \left(1 - \frac{\sum_{k=1}^{n} w_k}{(n-1) w_i}\right) + \frac{c_\ell}{\sum_{j=1}^{m} c_j} .$$

PROOF. The proof amounts to plugging the expressions for λ_i, $i \in [n]$, and W^ℓ, $\ell \in [m]$, derived in Lemmas 9 and 10, respectively, into the expressions for the Nash probabilities p_i^ℓ in Proposition 1. □

When do the quantities p_i^ℓ determined in Proposition 11 represent probabilities? For this, it must be that for each user $i \in [n]$, *(1)* $\sum_{j=1}^{m} p_i^j = 1$, and *(2)* for each link $\ell \in [m]$, $0 \leq p_i^\ell \leq 1$. Also, since these quantities were derived for the fully mixed case, condition (2) should actually be stated as *(2′)* for each link e, $1 \leq e \leq m$, $0 < p_i^e < 1$. A straightforward calculation verifies that condition (1) and condition (2′) may or may not hold, depending on the particular values of the traffics and link capacities. Hence, we obtain an impossibility result for Nash equilibria in the fully mixed case.

COROLLARY 12. *Assume that there exist a user $i \in [n]$ and a link $\ell \in [m]$ such that*

$$\left(1 - \frac{m c_\ell}{\sum_{j=1}^{m} c_j}\right) \left(1 - \frac{\sum_{k=1}^{n} w_k}{(n-1) w_i}\right) + \frac{c_\ell}{\sum_{j=1}^{m} c_j} \not\in (0,1) .$$

Then, in the fully mixed case, there is no Nash equilibrium.

Moreover, we show that the necessary condition for a Nash equilibrium (in the fully mixed case) determined in Corollary 12 is also sufficient.

PROPOSITION 13. *Assume that for each user $i \in [m]$ and link $\ell \in [m]$,*

$$\left(1 - \frac{m c_\ell}{\sum_{j=1}^{m} c_j}\right) \cdot \left(1 - \frac{\sum_{k=1}^{n} w_k}{(n-1) w_i}\right) + \frac{c_\ell}{\sum_{j=1}^{m} c_j} \in (0,1) .$$

Then, in the fully mixed case, the probabilities

$$p_i^\ell = \left(1 - \frac{m c_\ell}{\sum_{j=1}^{m} c_j}\right) \cdot \left(1 - \frac{\sum_{k=1}^{n} w_k}{(n-1) w_i}\right) + \frac{c_\ell}{\sum_{j=1}^{m} c_j} ,$$

for any user $i \in [n]$ and link $\ell \in [m]$, are Nash probabilities.

Propositions 11 and 13 together establish:

THEOREM 14. *Consider the fully mixed case. Then, for all users $i \in [n]$ and links $\ell \in [m]$,*

$$\left(1 - \frac{m c_\ell}{\sum_{j=1}^{m} c_j}\right) \cdot \left(1 - \frac{\sum_{k=1}^{n} w_k}{(n-1) w_i}\right) + \frac{c_\ell}{\sum_{j=1}^{m} c_j} \in (0,1)$$

if and only if there is a Nash equilibrium, which must be unique and has associated Nash probabilities

$$p_i^\ell = \left(1 - \frac{m c_\ell}{\sum_{j=1}^{m} c_j}\right) \cdot \left(1 - \frac{\sum_{k=1}^{n} w_k}{(n-1) w_i}\right) + \frac{c_\ell}{\sum_{j=1}^{m} c_j} ,$$

for any user $i \in [n]$ and link $\ell \in [m]$.

Proposition 14 implies that for the fully mixed case, Nash equilibrium can be checked for existence and evaluated (if existing) in time $\Theta(mn)$.

5. UNIFORM CAPACITIES & ARBITRARY TRAFFICS

In this section, we derive upper bounds on coordination ratio for the case of fully mixed strategies and the model of uniform capacities.

5.1 Preliminaries

Recall that in the model of uniform capacities, $c_\ell = c > 0$ for all links $\ell \in [m]$. We start with a characterization of Nash probabilities. (It follows immediately from Theorem 14.)

LEMMA 15. *Consider the fully mixed case for the model of uniform capacities. Then, there is a unique Nash equilibrium with Nash probabilities $p_i^\ell = 1/m$ for any user $i \in [n]$ and link $\ell \in [m]$.*

By Lemma 15, for each link ℓ, $\ell \in [m]$, $\mathbf{Pr}(\theta_\ell = \vartheta) = \binom{n}{\vartheta} \left(\frac{1}{m}\right)^\vartheta \left(1 - \frac{1}{m}\right)^{n-\vartheta}$, for each integer ϑ such that $0 \leq \vartheta \leq n$. Hence, each random variable θ_ℓ, $\ell \in [m]$, follows the binomial distribution with parameters n and $1/m$, so that $\mathcal{E}(\theta_1) = n/m$, $\ell \in [m]$.

5.2 Tail Lemma

We start by proving an upper bound on social cost under a certain assumption on the tail probability of $\max_{\ell \in [m]} \theta_\ell$.

PROPOSITION 16. *Consider the fully mixed case under the model of uniform capacities. Assume that, for a Nash equilibrium* $\mathbf{P}$*, there exists a function $\rho(m,n)$ such that for every link $\ell \in [m]$,*

$$\mathbf{Pr}\left(\max_{\ell \in [m]} \theta_\ell \leq \rho(m,n) \cdot \mathcal{E}(\theta_\ell)\right) \geq 1 - \frac{1}{n} .$$

Then,

$$\mathsf{SC}(\mathbf{w}, \mathbf{P}) \leq \frac{\max_{1 \leq k \leq n} w_k}{c} \cdot \left(\rho(m,n) \cdot \frac{n}{m} + 1\right) .$$

PROOF. We give an informal outline of our proof. We use the definitions for social cost and expectation to obtain that

$$\mathsf{SC}(\mathbf{w}, \mathbf{P}) \leq \frac{\max_{1 \leq k \leq n} w_k}{c} \sum_{\vartheta=0}^{n} \vartheta \, \mathbf{Pr}\left(\max_{\ell \in [m]} \theta_\ell = \vartheta\right) .$$

We then use the assumption and split the summation across $\rho(m,n)\mathcal{E}(\theta_{\ell_0})$, for any particular link $\ell_0 \in [m]$, to derive the claimed upper bound on $\mathsf{SC}(\mathbf{w}, \mathbf{P})$. □

Clearly, the function $\rho(m,n)$ in Lemma 16 describes the measure of the distribution of the random variable $\max_{\ell\in[m]}\theta_\ell$ in the *tail* of the distribution of *any* variable θ_ℓ, $\ell\in[m]$, that lies below its expectation $\mathcal{E}(\theta_\ell)=n/m$. So, call the function $\rho(m,n)$ a *tail function.*

Proposition 16 implies that to show an upper bound on social cost (and so, by Lemma 19, an upper bound on coordination ratio), it suffices to determine a suitable tail function $\rho(m,n)$. In the next section, we will do so in two particular instances.

5.3 Upper Bounds on Social Cost

In this section, we determine a suitable tail function under two particular cases on how m and n compare to each other. In doing so, we will envision our problem as a random allocation problem.

The Case $m=n$

We prove:

PROPOSITION 17. *Consider the fully mixed case under the model of uniform capacities. Assume that $m=n$. Then, for a Nash equilibrium $\mathbf{P}$,*

$$\mathsf{SC}(\mathbf{w},\mathbf{P}) \leq \frac{\max_{1\le k\le n} w_k}{c}\cdot\left(\frac{e\ln n}{\ln\ln n}+1\right).$$

PROOF. Since $\mathcal{E}(\theta_\ell)=1$ for all links $\ell\in[n]$, Lemma 4 that for each link $\ell\in[n]$,

$$\mathbf{Pr}\left(\max_{\ell\in[n]}\theta_\ell\le\frac{e\ \ln n}{\ln\ln n}\cdot\mathcal{E}(\theta_\ell)\right) \geq 1-\frac{1}{n}.$$

Thus, Proposition 16 applies with $\rho(m,n)=e\ln n/\ln\ln n$ to yield that

$$\mathsf{SC}(\mathbf{w},\mathbf{P}) \leq \frac{\max_{1\le k\le n} w_k}{c}\cdot\left(\frac{e\ln n}{\ln\ln n}+1\right),$$

as needed. □

The Case $m\le n/(24\ln n)$

We prove:

PROPOSITION 18. *Consider the fully mixed case under the model of uniform capacities. Assume that $m\le n/(24\ln n)$ for sufficiently large m and n. Then, for a Nash equilibrium $\mathbf{P}$,*

$$\mathsf{SC}(\mathbf{w},\mathbf{P}) \leq \frac{\max_{1\le k\le n} w_k}{c}\left(\frac{3}{2}\cdot\frac{n}{m}+1\right).$$

PROOF. Take any link $\ell\in[m]$. Since $\mathcal{E}(\theta_\ell)=n/m$, Lemma 5 implies that for any parameter $\beta\in(0,1)$,

$$\begin{aligned}\mathbf{Pr}(\theta_\ell>(1+\beta)\,\mathcal{E}(\theta_\ell)) &\leq \exp\left(-\frac{\beta^2}{2}\cdot\frac{n}{m}\right)\\ &\leq \exp\left(-\frac{\beta^2}{2}\cdot 24\ln n\right),\end{aligned}$$

Fix now $\beta=1/2$, so that

$$\begin{aligned}\mathbf{Pr}\left(\theta_\ell>\frac{3}{2}\cdot\frac{n}{m}\right) &\leq \exp(-2\ln n)\\ &= \frac{1}{n^2}.\end{aligned}$$

Thus,

$$\begin{aligned}\mathbf{Pr}\left(\max_{\ell\in[m]}\theta_\ell>\frac{3}{2}\cdot\frac{n}{m}\right) &\leq \mathbf{Pr}\left(\bigwedge_{\ell\in[m]}\left(\theta_\ell>\frac{3}{2}\cdot\frac{n}{m}\right)\right)\\ &\leq \mathbf{Pr}\left(\bigvee_{\ell\in[m]}\left(\theta_\ell>\frac{3}{2}\cdot\frac{n}{m}\right)\right)\\ &\leq \sum_{\ell\in[m]}\mathbf{Pr}\left(\theta_\ell>\frac{3}{2}\cdot\frac{n}{m}\right)\\ &= m\cdot\frac{1}{n^2}\\ &\leq \frac{n}{24\ln n}\cdot\frac{1}{n^2}\\ &< \frac{1}{n}.\end{aligned}$$

Thus, Proposition 16 applies with $\rho(m,n)=3/2$ to yield that

$$\mathsf{SC}(\mathbf{w},\mathbf{P}) \leq \frac{\max_{1\le k\le n} w_k}{c}\left(\frac{3}{2}\cdot\frac{n}{m}+1\right),$$

as needed. □

5.4 Lower Bound on Social Optimum

We prove:

LEMMA 19. *For the model of uniform capacities,*

$$\mathsf{SO}(\mathbf{w}) \geq \frac{1}{c}\cdot\max\left\{\frac{\sum_{1\le k\le n} w_k}{m},\ \max_{1\le k\le n} w_k\right\}.$$

PROOF. Clearly, in the optimal allocation of traffics to links, some link must receive traffic no less than the average (over all links) traffic, and some link must receive the maximum (over all users) traffic. This implies the claim. □

5.5 Recap

We are now ready to show our results for the fully mixed case and under the model of uniform capacities. Assuming $m=n$, we use Proposition 17 and Lemma 19 we obtain:

THEOREM 20. *Consider the fully mixed case for the model of uniform capacities. Assume that $m=n$. Then,*

$$\mathsf{CR} \leq \frac{e\ln n}{\ln\ln n}+1.$$

Assuming now that $m\le n/(24\ln n)$, we use Proposition 18 and Lemma 19 to obtain:

THEOREM 21. *Consider the fully mixed case for the model of uniform capacities. Assume that $m\le n/(24\ln n)$ for sufficiently large m and n. Then,*

$$\mathsf{CR} < \left(\frac{3}{2}+\varepsilon\right)\cdot\mathsf{max/min}(\mathbf{w}).$$

for any arbitrary constant $\varepsilon>0$.

6. NON-UNIFORM CAPACITIES & IDENTICAL TRAFFICS

In this section, we use properties of Nash equilibria shown in Section 4 to derive bounds on coordination ratio for the case of fully mixed strategies and the model of non-uniform capacities. Furthermore, we assume that all traffics are *identical*; that is, for all users i, $w_i = w$ for some constant $w > 0$.

For each link ℓ, $\ell \in [m]$, denote $\widetilde{c_\ell} = c_\ell / (\sum_{j=1}^{n} c_j)$, the *reduced capacity* of link ℓ. (Clearly, $\sum_{\ell \in [m]} \widetilde{c_\ell} = 1$.)

6.1 Preliminaries

We start with a simple fact characterizing the existence and uniqueness of Nash equilibria in this case; this is a direct consequence of Theorem 14.

LEMMA 22. *Consider the fully mixed case for the model of non-uniform capacities. Assume that all traffics are identical. Then, for all links ℓ, $\ell \in [m]$,*

$$\widetilde{c_\ell} \quad \in \quad \left(\frac{1}{m+n-1}, \; \frac{n}{m+n-1} \right) ,$$

if and only if there exists a Nash equilibrium, which must be unique and has associated Nash probabilities

$$p_i^\ell \quad \in \quad \frac{(n+m-1)\,\widetilde{c_\ell} - 1}{n-1} ,$$

for any user $i \in [n]$ and link $\ell \in [m]$.

PROOF. By Theorem 14, there exists a Nash equilibrium, which must be unique, if and only if for all users $i \in [n]$ and links $\ell \in [m]$,

$$(1 - m\widetilde{c_\ell}) \cdot \left(1 - \frac{1}{n-1} \cdot \frac{nw}{w}\right) + \widetilde{c_\ell} \quad \in \quad (0,1) .$$

Solving for $\widetilde{c_\ell}$ yields that

$$\widetilde{c_\ell} \quad \in \quad \left(\frac{1}{m+n-1}, \; \frac{n}{m+n-1} \right) ,$$

as needed. Also, by Theorem 14, the associated Nash probabilities are

$$\begin{aligned} p_i^\ell &= (1 - m\widetilde{c_\ell}) \cdot \left(1 - \frac{1}{n-1} \cdot \frac{nw}{w}\right) + \widetilde{c_\ell} \\ &= \frac{(n+m-1)\,\widetilde{c_\ell} - 1}{n-1} , \end{aligned}$$

for all users $i \in [n]$ and links $\ell \in [m]$, as needed. ☐

Lemma 22 describes the Nash probabilities for the case of identical traffics and under the model the model of non-uniform capacities; Lemma 15 represents its analog for the case of arbitrary traffics and under the model of uniform capacities. However, these two lemmas stand in sharp contrast to each other, since Lemma 15 shows the existence of a Nash equilibrium in *all* possible cases, while Lemma 22 provides conditions under which a (still unique) Nash equilibrium may exist. Thus, Lemmas 15 and 22 reveal an essential difference with respect to Nash equilibria between the case of uniform capacities and identical traffics, and the case of non-uniform capacities and arbitrary traffics, respectively.

Further on, we remark that Lemma 22 shows that each Nash probability is now *independent* of the particular user and depends only on the link; to emphasize this property, we will write α_ℓ to denote p_i^ℓ for any user $i \in [n]$ and link $\ell \in [m]$. Thus, for each link ℓ, $\ell \in [m]$,

$$\mathbf{Pr}\,(\theta_\ell = \vartheta) \quad = \quad \binom{n}{\vartheta} \alpha_\ell^{\vartheta} (1-\alpha_\ell)^{n-\vartheta} .$$

Hence, θ_ℓ follows the binomial distribution with parameters n and α_ℓ, $\ell \in [m]$, so that $\mathcal{E}(\theta_\ell) = \alpha_\ell n$.

6.2 Tail Lemma

We prove an upper bound on social cost under a certain assumption on the tails of the probability distributions of the random variables θ_ℓ, $\ell \in [m]$.

PROPOSITION 23. *Consider the fully mixed case under the model of non-uniform capacities. Assume that all traffics are equal to w. Assume that for each link ℓ, $\ell \in [m]$, there exists a function $\rho_\ell(m,n)$ so that*

$$\mathbf{Pr}\left(\bigwedge_{\ell \in [m]} \left(\theta_\ell \leq \rho_\ell(m,n) \cdot \max\{1, \mathcal{E}(\theta_\ell)\}\right) \right) \quad > \quad 1 - \frac{\delta}{n} ,$$

for some constant $\delta > 0$. Then, in a Nash equilibrium $\mathbf{P}$,

$$\mathsf{SC}\,(\mathbf{w},\mathbf{P}) \quad \leq \quad \frac{w\,(m+n-1)}{\sum_{\ell \in [m]} c_\ell} \cdot \left(\max_{\ell \in [m]} \rho_\ell(m,n) \cdot \frac{1}{n-1} + \delta \right) .$$

PROOF. We give an informal outline of our proof. We start by observing that

$$\mathsf{SC}\,(\mathbf{w},\mathbf{P}) \quad = \quad w \cdot \mathcal{E}\left(\max_{\ell \in [m]} \frac{\theta_\ell}{c_\ell} \right) .$$

To bound from above the expectation of the (discrete) random variable $\max_{\ell \in [m]} \theta_\ell / c_\ell$, we partition its sample space according to the event that for *all* links $\ell \in [m]$,

$$\theta_\ell \quad \leq \quad \rho_\ell(m,n) \max\left\{1, \mathcal{E}\left(\theta_\ell\right)\right\} ,$$

to obtain that

$$\begin{aligned} &\mathsf{SC}\,(\mathbf{w},\mathbf{P}) \\ &\leq \quad w \cdot \left(\max_{\ell \in [m]} \rho_\ell(m,n) \cdot \max\left\{ \max_{\ell \in [m]} \frac{1}{c_\ell}, \; \max_{\ell \in [m]} \frac{\mathcal{E}(\theta_\ell)}{c_\ell} \right\} \right. \\ &\qquad \left. + \frac{\delta}{\min_{\ell \in [m]} c_\ell} \right) . \end{aligned}$$

We then use Lemma 22 and the expressions for the Nash probabilities α_ℓ, $\ell \in [m]$, to establish the claim. ☐

6.3 Upper Bound on Social Cost

In this section, we determine a suitable tail function for each link $\ell \in [m]$, under the assumption that $m \leq n$. Taking the maximum of these tail functions will yield, via Lemma 23, an upper bound on social cost. In our analysis, we will still envision our problem as an *occupancy problem* [19]. We show:

PROPOSITION 24. *Consider the fully mixed case under the model of non-uniform capacities. Assume that all traffics are identical and equal to w. Assume also that $m \leq n$. Then, in a Nash equilibrium $\mathbf{P}$,*

$$\begin{aligned} &\mathsf{SC}\,(\mathbf{w},\mathbf{P}) \\ &\leq \quad \frac{w}{\sum_{\ell \in [m]} c_\ell} \cdot \left(\frac{e \ln n}{\ln \ln n} \cdot \frac{n(2n-1)}{n-1} + (2e+1) \cdot (2n-1) \right) . \end{aligned}$$

PROOF. We give an informal outline of our proof. We will determine an appropriate tail function $\rho_\ell(m,n)$ for each link $\ell \in [m]$ so that

$$\mathbf{Pr}\left(\bigwedge_{\ell\in[m]} (\theta_\ell \leq \rho_\ell(m,n) \cdot \max\{1, \mathcal{E}(\theta_\ell)\})\right) \quad > \quad 1 - \frac{\delta}{n},$$

for some appropriate constant $\delta > 0$, whence Lemma 23 applies. To do so, we consider separately each of two possible cases as to how $\mathcal{E}(\theta_\ell)$ compares to 1, for each link $\ell \in [m]$, and we use Lemmas 6 and 7.

Set now $\rho_\ell(m,n) = \frac{e \ln n}{\ln\ln n}$, for each link $\ell \in [m]$. We use the case analysis above to show that

$$\mathbf{Pr}\left(\bigwedge_{\ell\in[m]} (\theta_\ell \leq \rho_\ell(m,n) \cdot \max\{1, \mathcal{E}(\theta_\ell)\})\right) \quad > \quad 1 - \frac{2e+1}{n}.$$

Then, Lemma 23 establishes the claim. $\square$

6.4 Lower Bound on Social Optimum

We continue with a lower bound on social optimum.

LEMMA 25. *Consider the fully mixed case under the model of uniform capacities. Assume that all traffics are identical and equal to w. Then,* $\mathsf{SO}(\mathbf{w}) \geq \frac{nw}{\sum_{e\in[m]} c_e}$.

PROOF. We give an informal outline of our proof. The proof considers an alternative routing model where each individual traffic may be "split" among more than one links; clearly, the social optimum is no less than the value it attains in this model. This splitting assumption is modeled by introducing a *split fraction* $\phi_\ell \in [0,1]$ for each link $\ell \in [m]$, representing the fraction of the total traffic nw received by link ℓ. Thus, the social optimum for this model is the minimum, over all possible choices of split fractions ϕ_ℓ for links, of the maximum, over all links, of $\phi_\ell(nw)/c_\ell = (nw)\phi_\ell/c_\ell$. The claim follows then easily from the observation that there exists a link $\ell \in [m]$ such that $\phi_\ell \geq \widetilde{c_\ell}$ (using also the fact that $\sum_{\ell\in[m]} \widetilde{c_\ell} = 1$). $\square$

6.5 Recap

We are now ready to show our result for the fully mixed case under the model of non-uniform capacities and assuming identical traffics. We use Proposition 24 and Lemma 25 to show:

THEOREM 26. *Consider the fully mixed case for the model of non-uniform capacities. Assume that all traffics are identical. Assume also that $m \leq n$. Then,*

$$\mathsf{CR} \quad \leq \quad \frac{2e\ln n}{\ln\ln n} + 2(2e+1).$$

7. DISCUSSION

In this work, we studied the problem of selfish routing in a noncooperative network consisting of a set of m parallel links, within a game-theoretic framework suggested in a recent work by Koutsoupias and Papadimitriou [25]. We used coordination ratio [25] as a measure of performance loss in noncooperative networks due to lack of coordination.

We identified a set of stochastic equilibrium equations (one per each of n users), called the minimum expected latency cost equations, that describe the Nash equilibrium of the system. In turn, we particularized these equations to study in depth a natural special case of Nash equilibria, namely fully mixed strategies, where each user may ship its traffic over any link. For the case of fully mixed strategies, we proved asymptotically tight upper bounds of $\Theta(\ln n/\ln\ln n)$ on coordination ratio over two interesting instances: all links have the same capacity while users' traffics may vary, and symmetrically, all users carry the same traffic while links' capacities may now vary. The techniques we developed seem to bridge the gap between computer science and game theory for the specific instances of the selfish routing problem we study.

Although the case of fully mixed strategies is one out of the exponentially many possible Nash equilibria, we feel that it does encapsulate the difficulty of the whole problem. We believe that the techniques we developed for studying the case of fully mixed strategies can be appropriately extended to the general case of Nash equilibria to yield corresponding upper bounds on coordination ratio in the general case. We leave this extention as the most obvious open problem suggested by our work. Still for the case of fully mixed strategies and under the model of uniform capacities, it would be very interesting to prove bounds on coordination ratio for values of the ratio m/n in the interval $(1/24\ln n, 1)$ or larger than 1, thus complementing Theorems 20 and 21.

A wide avenue of other significant problems that remain tantalizingly open include the consideration of more general network topologies, latency cost functions, and performance measures within the game-theoretic framework for selfish routing adopted in this article. We hope that our work provides a solid, initial ground for settling these issues.

Acknowledgments

We thank Elias Koutsoupias and Christos Papadimitriou whose paper *"Worst-case Equilibria"* [25] has inspired our work.

8. REFERENCES

[1] E. Altman, "Flow Control Using the Theory of Zero-Sum Markov Games," *IEEE Transactions on Automatic Control,* Vol. 39, pp. 814–818, April 1994.

[2] E. Altman, T. Basar, T. Jimenez and N. Shimkin, "Competitive Routing in Networks with Polynomial Cost," *Proceedings of the 19th IEEE Conference on Computer Communications (IEEE INFOCOM 2000),* March 2000.

[3] D. Angluin and L. G. Valiant, "Fast Probabilistic Algorithms for Hamiltonian Circuits and Matchings," *Journal of Computer and System Sciences,* Vol. 18, pp. 155–193, 1979.

[4] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela and M. Protasi, *Complexity and Approximation – Combinatorial Optimization Problems and Their Approximability Properties,* Springer, 1999.

[5] Y. Azar, A. Z. Broder, A. R. Karlin and E. Upfal, *Balanced Allocations, Proceedings of the 26th Annual ACM Symposium on the Theory of Computing,* pp. 593–602, May 1994.

[6] T. Basar and G. J. Olsder, *Dynamic Noncooperative Game Theory,* SIAM Series in Classics in Applied Mathematics, Philadelphia, 1999.

[7] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis,* Cambridge University Press, 1998.

[8] J. J. Carrahan, P. A. Russo, K. Kitami and R. Kung, “Intelligent Network Overview,” *IEEE Communications Magazine,* Vol. 31, pp. 30–36, March 1993.

[9] H. Chernoff, ”A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the Sum of Observations,” *Annals of Mathematical Statistics,* Vol. 23, pp. 493–509, 1952.

[10] S. Deering and R. Hinden, *Internet Protocol Version 6 Specification,* Internet Draft, IETF, March 1995.

[11] P. Dubey, “Inefficiency of Nash Equilibria,” *Mathematics of Operations Research,* Vol. 11, No. 1, pp. 1–8, February 1986.

[12] A. A. Economides and J. A. Silvester, “Multi-objective Routing in Integrated Services Networks: A Game Theory Approach,” *Proceedings of the IEEE Conference on Computer Communications (IEEE INFOCOM 1991),* pp. 1220–1225, 1991.

[13] J. Feigenbaum, C. Papadimitriou and S. Shenker, “Sharing the Cost of Multicast Transactions,” *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing,* pp. 218–227, Portland, Oregon, May 2000.

[14] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes,* Oxford Science Publications, Second Edition, 1992.

[15] M. Habib, C. McDiarmid, J. Ramirez-Alfonsin and B. Reed (eds.), *Probabilistic Methods in Algorithmic Discrete Mathematics,* Vol. 16 in *Algorithms and Combinatorics,* Springer, 1998.

[16] M.-T. T. Hsiao and A. A. Lazar, “Optimal Decentralized Flow Control of Markovian Queueing Networks with Multiple Controllers,” *Performance Evaluation,* Vol. 13, No. 3, pp. 181–204, 1991.

[17] N. L. Johnson and S. Kotz, *Urn Models and Their Applications,* John Wiley & Sons, 1977.

[18] R. Karp, E. Koutsoupias, C. H. Papadimitriou and S. Shenker, “Optimization Problems in Congestion Control,” *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science,* October 2000.

[19] V. F. Kolchin, V. P. Chistiakov and B. A. Sevastianov, *Random Allocations,* V. H. Winston, New York, 1978.

[20] Y. Korilis and A. Lazar, ”On the Existence of Equilibria in Noncooperative Optimal Flow Control,” *Journal of the ACM,* Vol. 42, No. 3, pp. 584–613, 1995.

[21] Y. A. Korilis, A. A. Lazar and A. Orda, ”Architecting Noncooperative Networks,” *IEEE Journal on Selected Areas in Communications,* Vol. 13, No. 7, pp. 1241–1251, 1995.

[22] Y. Korilis, A. A. Lazar and A. Orda, “Achieving Network Optima Using Stackelberg Routing Strategies,” *IEEE/ACM Transactions on Networking,* Vol. 5, No. 1, pp. 161–173, February 1997.

[23] Y. A. Korilis, A. A. Lazar and A. Orda, ”Capacity Allocation under Noncooperative Routing,” *IEEE Transactions on Automatic Control,* Vol. 42, No. 3, pp. 309–325, 1997.

[24] Y. A. Korilis, A. A. Lazar and A. Orda, “Avoiding the Braess Paradox in Non-Cooperative Networks,” *Journal of Applied Probability,* Vol. 36, pp. 211–222, 1999.

[25] E. Koutsoupias and C. H. Papadimitriou, ”Worst-case Equilibria,” *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science,* G. Meinel and S. Tison eds., pp. 404–413, Vol. 1563, Lecture Notes in Computer Science, Springer-Verlag, Trier, Germany, March 1999.

[26] A. A. Lazar, A. Orda and D. E. Pendarakis, “Virtual Path Bandwidth Allocation in Multi-User Networks,” *IEEE/ACM Transactions on Networking,* Vol. 5, pp. 861–871, December 1997.

[27] L. Libman and A. Orda, “The Designer’s Perspective to Atomic Noncooperative Networks,” *IEEE/ACM Transactions on Networking,* Vol. 7, No. 6, pp. 875–884, December 1999.

[28] M. Mavronicolas, A. Mouskos and P. Spirakis, ”The Cost of Lack of Coordination in Distributed Network Routing,” unpublished manuscript, April 2000.

[29] R. D. McKelvey and A. McLennan, “Computation of Equilibria in Finite Games,” in *Handbook of Computational Economics,* Vol. 1, pp. 87–142, 1996.

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

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

[32] A. Orda, R. Rom and N. Shimkin, “Competitive Routing in Multiuser Communication Networks,” *IEEE/ACM Transactions on Networking,* Vol. 1, pp. 510–521, October 1993.

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

[34] G. Owen, *Game Theory,* Third Edition, Academic Press, 1995.

[35] C. H. Papadimitriou and M. Yannakakis, “On the Approximability of Trade-offs and Optimal Access of Web Sources,” *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science,* October 2000.

[36] T. Roughgarden and É. Tardos, ”How Bad is Selfish Routing?” *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science,* October 2000.

[37] Z. Zhang and C. Douligeris, “Convergence of Synchronous and Asynchronous Greedy Algorithms in a Multiclass Telecommunications Environment,” *IEEE Transactions on Computers,* Vol. 40, pp. 1277–1281, August 1992.