

Tradeoffs in worst-case equilibria

Baruch Awerbuch^a

^a*Dept. of Computer Science, Johns Hopkins University, Baltimore, MD, 21218.*

Yossi Azar^{b,1}, Yossi Richter^{b,2}

^b*School of Computer Science, Tel Aviv University, Tel Aviv, 69978, Israel.*

Dekel Tsur^c

^c*Department of Computer Science and Engineering, University of California San Diego, La Jolla, CA 92093*

Abstract

We investigate the problem of routing traffic through a congested network in an environment of non-cooperative users. We use the worst-case coordination ratio suggested by Koutsoupias and Papadimitriou to measure the performance degradation due to the lack of a centralized traffic regulating authority. We provide a full characterization of the worst-case coordination ratio in the restricted assignment and unrelated parallel links model. In particular, we quantify the tradeoff between the “negligibility” of the traffic controlled by each user and the worst-case coordination ratio. We analyze both pure and mixed strategies systems and identify the range where their performance is similar.

Email addresses: baruch@cs.jhu.edu (Baruch Awerbuch), azar@tau.ac.il (Yossi Azar), yo@tau.ac.il (Yossi Richter), dtsur@cs.ucsd.edu (Dekel Tsur).

¹ Research supported in part by the Israeli Ministry of industry and trade and by the Israel Science Foundation.

² Research supported in part by the Israeli Ministry of industry and trade.

1 Introduction

1.1 Overview:

In most communication networks it is infeasible to maintain one centralized authority to route traffic efficiently. As a result, users may decide individually how to route their traffic. Each user behaves selfishly in the sense that he wishes to minimize his transmission cost while being aware of the network congestion caused by other users. A system of users decisions is said to be in a Nash equilibrium if no user can benefit from changing his decision. Simple examples in game theory show that the performance of systems in Nash equilibrium, achieved by non-cooperative users, can be far from the global optimum. Recently, the question of quantifying the decrease in network performance caused by the lack of a centralized authority, received considerable attention among researchers. Koutsoupias and Papadimitriou [5,8] suggested to investigate the worst-case coordination ratio, which is the ratio between the worst possible Nash equilibrium and the global optimum, as a mean to understand the cost incurred due to the lack of a centralized regulating authority.

Problems of this type have been studied lately by two approaches. The worst-case coordination ratio in a network composed of m parallel related links was analyzed in [4,5,8], while making no assumptions on the relative amount of traffic controlled by each user. On the other hand, [11,12] obtained improved bounds for general networks while assuming that each user controls only a negligible fraction of the total traffic. We attempt to bridge between these two approaches, by analyzing the worst-case coordination ratio as a function of the relative fraction of the total traffic controlled by each user. Specifically, we quantify the required “negligibility” of each user’s traffic, needed to bound the worst-case coordination ratio by a constant.

We focus on two new models. First, we study the *restricted* assignment model (also called the subset model) which is defined as follows. The network consists of m parallel links. There are n users where user j has amount of traffic w_j , that can be transmitted through any link from a subset M_j of the m links. We consider both the *pure strategies* case where each user selects one link to transmit his traffic, and the more general case of *mixed strategies* where each user decides a probability distribution over his allowable links. In both cases each user is aware of the decisions made by other users. The cost incurred by each user is the total load on the link to which his traffic is assigned. Since all users behave selfishly, each wishes to minimize his cost by assigning his traffic to the least loaded link. The global objective however, is to minimize the load of the most loaded link. We note that even though this model is a simplification of real communication networks, it captures the essence of basic

networking problems as pointed out by [5,6,8,11].

Additionally, we study the more general model of *unrelated* links in which a task j ($j = 1, \dots, n$) is associated with an m -vector \vec{w}_j specifying its weight on each link. This model corresponds to a scenario in which the network includes diverse links, and the traffic generated by each task depends on the specific link used (e.g. network parameters such as latency, error rate, bandwidth etc. have different influence over the amount of traffic generated by different network applications). We analyze the worst-case coordination ratio as a function of the maximum stretch s in the system, where $s = \max_{j,i,l:w_{ij}<\infty} \frac{w_{ij}}{w_{lj}}$.

1.2 Our results:

We provide a full characterization of the worst-case coordination ratio for both pure and mixed strategies in the restricted assignment and unrelated links models. Specifically, we prove the following results:

- We show the following tight bounds for the restricted assignment model, as functions of the ratio r between the optimal assignment and the largest task (notice that $r \geq 1$):
 - For $1 \leq r \leq \log m$ we prove that the worst-case coordination ratio is bounded by $\Theta\left(\frac{\log m}{r \cdot \log(2 + \frac{\log m}{r})}\right)$ in the pure strategies case, and is bounded by $\Theta\left(\frac{\log m}{r \cdot \log \log(2 + \frac{\log m}{r})}\right)$ in the mixed strategies case.
 - For any $0 < \epsilon < 1$ and $r = \Omega\left(\frac{\log m}{\epsilon^2}\right)$ we show that the worst-case coordination ratio for both pure and mixed strategies systems is bounded by $1 + \epsilon$.

Note that general bounds of $\Theta\left(\frac{\log m}{\log \log m}\right)$ for pure strategies and $\Theta\left(\frac{\log m}{\log \log \log m}\right)$ for mixed strategies are obtained when $r = 1$, i.e. when making no assumptions on the largest task in the system. We also note that for $r = \Omega(\log m)$ the worst-case coordination ratio is bounded by a constant, for both pure and mixed strategies.

- In the unrelated links model we prove that the worst-case coordination ratio is bounded by $\Theta\left(s + \frac{\log m}{\log(2 + \frac{\log m}{s})}\right)$ in pure strategies systems, and $\Theta\left(\frac{s \cdot \log m}{\log \log m} + \frac{s \cdot \log m}{\log(s \cdot \log(2 + \frac{\log m}{s}))}\right)$ when mixed strategies are allowed.

The paper most closely related to ours is [4], which provides tight bounds for the *related* links model. In this model the links may have different speeds, each task is associated with a weight and can be assigned to any link, and the load of each link equals the total amount of weight assigned to it divided by its speed. Note that this is also a special case of the unrelated links model.

We follow the ideas in [4] and extend their results to the restricted assignment and unrelated links models, while formulating the influence of additional important parameters on the worst-case coordination ratio. Our proofs mainly rely on the techniques and ideas used in [4].

1.3 Related work:

Koutsoupias and Papadimitriou [5] initiated the study of worst-case coordination ratio in networks composed of m -parallel related links with possibly different speeds. They first investigated the case of two links and proved a worst-case coordination ratio of $3/2$ for the case of identical links, and $\phi = \frac{1+\sqrt{5}}{2}$ for links with possibly different speeds. They also obtained non-tight bounds for the general case of m parallel links. Mavronicolas and Spirakis [6] continued this line of research while focusing on the special case of fully-mixed strategies in which the probability of assigning any task to any link is non-zero. They proved that in this case the worst-case coordination ratio is bounded by $\Theta(\frac{\log m}{\log \log m})$ for both the identical links model and the general related links model where all tasks have equal weights and $m \leq n$. Czumaj and Vöcking [4] proved tight bounds for the m -parallel related links model, and showed that the worst-case coordination ratio is bounded by $\Theta(\frac{\log m}{\log \log m})$ in the identical links model and by $\Theta(\frac{\log m}{\log \log \log m})$ in the general related links model. Czumaj *et al.* [3] continued to study this problem and characterized the coordination and bicriteria ratios for different families of cost functions.

Roughgarden *et al.* [10,12] also examined the degradation in network performance due to unregulated traffic. Their model deals with general networks where users adopt *pure strategies* only, and the amount of traffic of each user is assumed to be a negligible fraction of the total traffic. The objective is to minimize the total latency. They proved that when the latency of each edge is a linear function of the edge congestion, any flow at Nash equilibrium has total latency at most $4/3$ of the optimal flow. Although for general latency functions the worst-case coordination ratio is unbounded, the following bicriteria result can be shown: for any network with continuous non-decreasing latency functions, a flow at Nash equilibrium has total latency no more than that of an optimal flow forced to route twice as much traffic. Roughgarden [11] also showed that the cost of unregulated traffic does not depend on the complexity of the network topology. He also studied the impact of latency functions belonging to specific classes. Roughgarden *et al.* [1,2,9] studied various ways to construct and price networks such that the cost incurred in unregulated traffic is minimized.

1.4 Paper structure:

The paper is organized as follows. Section 2 includes formal definitions and notations. The restricted assignment model is studied in Section 3. In Section 4 we analyze the unrelated links model.

2 Definitions and notations

The *restricted assignment model* is defined as follows: there are m parallel links and n users, where user j ($j = 1, \dots, n$) has a task with weight w_j , that can be assigned to any link from a subset M_j of the m links. We denote the largest task in the system by $w_{max} = \max_{1 \leq j \leq n} w_j$. Given an assignment of the tasks to the links, the load of each link is defined as the sum of the weights that are associated with the tasks being assigned to it. Given an instance of the problem we define the *global optimum* (denoted by OPT) to be the assignment of tasks to links that minimizes the maximum load of a link. We denote the ratio between the value of the optimal solution and the largest task by $r = \frac{OPT}{w_{max}}$. The *unrelated links model* is more general. Task j ($j = 1, \dots, n$) is associated with an m -vector \vec{w}_j , where w_{ij} indicates the weight of task j on link i .

We assume that the users are non-cooperative and each one wishes to minimize his own cost with no regard to the global optimum. We consider two types of users strategies systems:

- (1) **Pure strategies:** user j selects link $l_j \in M_j$ and assigns his task to it. Each user is aware of the choices made by all other users when making his decision.
- (2) **Mixed strategies:** user j selects a probability distribution $\{p_{ij}\}$ ($i \in M_j$) over the allowable set of links for task j . Each user is aware of the probability distributions selected by all other users.

For the remainder of this section we regard pure strategies as a special case of mixed strategies, and give our definitions in terms of mixed strategies systems. Given a system S of mixed strategies with probability distributions $\{p_{ij}^S\}$, we define the following random variables:

- A set of indicator random variables $\{X_{ij}^S\}$, where X_{ij}^S indicates whether task j is assigned to link i . By definition: $\Pr[X_{ij}^S = 1] = p_{ij}^S$.
- For each link i ($i = 1, \dots, m$) we define a random variable L_i^S , indicating the total load on the link: $L_i^S = \sum_{j=1}^n w_j \cdot X_{ij}^S$. We denote the maximum expected load by $\mu^S = \max_{1 \leq i \leq m} E[L_i^S]$.

- We define a random variable $L_{max}^S = \max_{1 \leq i \leq m} L_i^S$ to indicate the maximum link load, and denote its expectancy by $\mu_{max}^S = E[L_{max}^S]$. Clearly, $\mu_{max}^S \geq \mu^S$ (for pure strategies $\mu_{max}^S = \mu^S$).

For simplicity of notation, throughout the paper we omit the superscript S when meaning is clear from context.

Definition 1 *The expected cost of user j for assigning his task to link i in system S is defined as: $c_{ij}^S = E[L_i^S | X_{ij}^S = 1] = E[L_i^S] + (1 - p_{ij}^S)w_j$.*

Definition 2 *A system S is said to be in Nash Equilibrium if and only if for every task j and link i , $p_{ij}^S > 0$ only if $c_{ij}^S = \min_{1 \leq k \leq m} c_{kj}^S$.*

Definition 3 *The worst-case coordination ratio of an instance of the problem is defined as $R = \max_S \frac{\mu_{max}^S}{OPT}$, where the maximum is taken over all strategies systems S in Nash equilibrium.*

Throughout the paper we use the standard Gamma function defined by $\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$. We also frequently employ the following known equality for the inverse Gamma function: $\Gamma^{-1}(m) = (1 + o(1))(\frac{\log m}{\log \log m})$.

3 Worst-case equilibria in the restricted assignment model

In this section we provide tight bounds for worst-case equilibria in the restricted assignment model. We investigate both pure and mixed strategies. We first show upper bounds for the problem, and then provide matching lower bounds.

3.1 Upper bounds for restricted assignment

We begin by proving an upper bound on the maximum expected load in any system in Nash equilibrium. Recall that $r = \frac{OPT}{w_{max}}$; without loss of generality we assume that r is integral.

Lemma 4 *Let S be any system in Nash equilibrium. Then $\mu^S \leq (\beta + \frac{1}{r}) \cdot OPT$, for some β that satisfies the inequality $e(\frac{\beta}{e})^\beta \leq m^{\frac{1}{r}}$.*

PROOF. We order the links by non-increasing order of their expected loads. For every $k \geq 1$, define m_k to be the minimal integer such that $E[L_{m_k+1}] < k \cdot w_{max}$ ($m_k = m$ if there is no such integer). We define $l = \lfloor \frac{\mu^S}{w_{max}} \rfloor$. Note

that $m_l \geq 1$. The next claim, which is analogous to Claim 2.2 in [4], states an important property regarding the relation between m_k and m_{k+1} .

Claim 5 For every $k \geq 1$, $\frac{m_k}{m_{k+1}} \geq \frac{k+1}{r}$.

PROOF. Denote by J_{k+1} the set of tasks with positive probability to be assigned to any link from $[1, \dots, m_{k+1}]$, and denote the total weight of tasks from J_{k+1} by $W(J_{k+1}) = \sum_{j \in J_{k+1}} w_j$. Consider the way OPT schedules the tasks from J_{k+1} . We claim that OPT can not assign a task from J_{k+1} to a link with index larger than m_k . To prove this, let us assume, for contradiction, that OPT assigns task $j \in J_{k+1}$ to link $t > m_k$, and let $p_{qj} > 0$ for some link $q \leq m_{k+1}$. There follows:

$$\begin{aligned} c_{qj} &= E[L_q] + (1 - p_{qj})w_j \geq (k+1)w_{max} + (1 - p_{qj})w_j \geq (k+1)w_{max} \\ &\geq k \cdot w_{max} + (1 - p_{tj})w_j > c_{tj}, \end{aligned}$$

where the first inequality follows from the definition of q and the third inequality results from $w_j \leq w_{max}$ ($j = 1, \dots, n$). This contradicts the fact that the system S is in Nash equilibrium. Hence, OPT must assign all tasks from J_{k+1} to links in the range $[1, \dots, m_k]$. Let L_i^{OPT} denote the load on link i in the optimal assignment. Recall that $w_{max} = \frac{OPT}{r}$. We conclude that

$$m_k \cdot OPT \geq \sum_{i=1}^{m_k} L_i^{OPT} \geq W(J_{k+1}) \geq \sum_{i=1}^{m_{k+1}} E[L_i] \geq (k+1) \left(\frac{OPT}{r} \right) m_{k+1},$$

which yields the desired inequality. \square

We use Claim 5 iteratively together with the inequality $(n!/k!) \geq \frac{(n/e)^n}{(k/e)^k}$ to obtain the following:

$$m \geq m_r \geq \frac{l(l-1) \cdots (r+1)}{r^{l-r}} = \frac{l!}{r! \cdot r^{l-r}} \geq \frac{\left(\frac{l}{e}\right)^l}{\left(\frac{r}{e}\right)^r \cdot r^{l-r}}.$$

Substituting $l = \beta \cdot r$ we get: $e\left(\frac{\beta}{e}\right)^\beta \leq m^{1/r}$ and $\mu^S \leq (l+1)\frac{OPT}{r} = \left(\beta + \frac{1}{r}\right)OPT$. \square

3.1.1 Pure strategies

In the following theorem we show an upper bound on the worst-case coordination ratio as a function of the ratio r between the optimal solution and the largest task in the system.

Theorem 6 For pure strategies, when $1 \leq r \leq \log m$, $R = O\left(\frac{\log m}{r \cdot \log(2 + \frac{\log m}{r})}\right)$.

PROOF. Since $\mu_{max}^S = \mu^S$ for pure strategies, and given the inequality stated in Lemma 4, it suffices to bound the value of β in order to prove the theorem. Applying Stirling's formula to the inequality from Lemma 4 we have:

$$m^{\frac{1}{r}} \geq e \left(\frac{\beta}{e}\right)^\beta = (1 + o(1)) e \frac{\beta!}{\sqrt{2\pi\beta}},$$

and therefore $\beta \leq (1 + o(1)) \Gamma^{-1}(m^{\frac{1}{r}}) = O\left(\frac{\log m}{r \cdot \log(2 + \frac{\log m}{r})}\right)$. The theorem follows. \square

As a direct result from Theorem 6, we obtain the following general upper bounds for two extreme cases. In the first case we make no assumptions regarding the amount of traffic controlled by each user, namely, the largest task may be as large as the optimal solution, i.e. $r = 1$. In the second case the tasks are relatively small, i.e. $r = \Theta(\log m)$.

Corollary 7 For pure strategies, $R = O\left(\frac{\log m}{\log \log m}\right)$.

Corollary 8 For pure strategies, when $r = \Theta(\log m)$, we have $R = O(1)$.

Next, we study the worst-case coordination ratio in networks where each user controls a small fraction of the total traffic, i.e. $r = \Omega(\log m)$. The next theorem shows that in such networks the worst-case coordination ratio is close to 1.

Theorem 9 For any $0 < \epsilon < 1$, if $r = \Omega\left(\frac{\log m}{\epsilon^2}\right)$, then $R \leq 1 + \epsilon$.

PROOF. From Lemma 4 we know that $R \leq \beta + \frac{1}{r}$ where β satisfies the inequality: $e\left(\frac{\beta}{e}\right)^\beta \leq m^{1/r}$. When taking $r = \Omega\left(\frac{\log m}{\epsilon^2}\right)$ we obtain using Taylor expansion:

$$e^{O(\epsilon^2)} \geq m^{1/r} \geq e \left(\frac{\beta}{e}\right)^\beta = e^{\frac{1}{2}(\beta-1)^2 - O((\beta-1)^3)},$$

and the inequality yields $R \leq \beta + \frac{1}{r} \leq 1 + \epsilon$. \square

3.1.2 Mixed strategies

The following theorem gives an upper bound on the worst-case coordination ratio for the mixed strategies case. Its proof employs the techniques used in the proof of Theorem 1.1 in [4].

Theorem 10 For mixed strategies, when $1 \leq r \leq \log m$, $R = O\left(\frac{\log m}{r \cdot \log \log(2 + \frac{\log m}{r})}\right)$.

PROOF. For convenience we scale the tasks weights such that $w_{max} = 1$. Consider an arbitrary link i . Recall that $L_i = \sum_{j=1}^n w_j X_{ij}$, where $w_j \leq w_{max} = 1$ for $j = 1, \dots, n$, and $\mu = \max_{1 \leq i \leq m} E[L_i]$. We apply Hoeffding inequality³ and obtain for every $c > 1$:

$$\Pr[L_i \geq c \cdot \mu] \leq \left(\frac{e \cdot E[L_i]}{c \cdot \mu}\right)^{c \cdot \mu} \leq \left(\frac{e \cdot \mu}{c \cdot \mu}\right)^{c \cdot \mu} = \left(\frac{e}{c}\right)^{c \cdot \mu}.$$

We can now apply the union bound to obtain: $\Pr[L_{max} \geq c \cdot \mu] \leq m(e/c)^{c \cdot \mu}$. For every integer $\alpha > 10$ we can upper bound the expected maximum load by:

$$\begin{aligned} \mu_{max} &\leq \alpha \cdot \mu + \mu \cdot \sum_{k=\alpha}^{\infty} \Pr[L_{max} \geq k \cdot \mu] \leq \alpha \cdot \mu + m \cdot \mu \sum_{k=\alpha}^{\infty} (e/k)^{k \cdot \mu} \\ &\leq \alpha \cdot \mu + (2 + 1/\mu) \cdot m \mu \cdot (e/\alpha)^{\alpha \cdot \mu}, \end{aligned}$$

where the last inequality follows from the fact that the sequence is subgeometric. We can substitute α for $\Gamma^{-1}(m^{1/\mu}) + d$, for a sufficiently large constant d . We then have: $(e/\alpha)^{\alpha \cdot \mu} \leq 1/m$, hence $\mu_{max} \leq (\alpha + 2 + 1/\mu)\mu$. Since $\alpha = O\left(\frac{\log m}{\mu \cdot \log \frac{\log m}{\mu}}\right)$ there follows: $\mu_{max} = O\left(\frac{\log m}{\log \frac{\log m}{\mu}}\right)$. Since $OPT = r$, Lemma 4 together with the proof of Theorem 6 imply that $\mu = O\left(\frac{\log m}{\log(2 + \frac{\log m}{r})}\right)$, substituting μ we conclude that $\mu_{max} = O\left(\frac{\log m}{\log \log(2 + \frac{\log m}{r})}\right)$ and therefore $R = O\left(\frac{\log m}{r \cdot \log \log(2 + \frac{\log m}{r})}\right)$. \square

The following upper bounds for two extreme cases are derived directly from Theorem 10.

Corollary 11 In the restricted assignment problem $R = O\left(\frac{\log m}{\log \log \log m}\right)$.

Corollary 12 When $r = \Theta(\log m)$ we have $R = O(1)$.

The next theorem refers to networks where each user controls a small fraction of the total traffic.

³ We use the following version of Hoeffding inequality: Let $\{X_i\}_{i=1}^N$ be independent random variables with values in $[0, z]$, and let $X = \sum_{i=1}^N X_i$. Then $\Pr[X \geq t] \leq (e \cdot E[X]/t)^{t/z}$.

Theorem 13 For any $0 < \epsilon < 1$, if $r = \Omega\left(\frac{\log m}{\epsilon^2}\right)$, then $R \leq 1 + \epsilon$.

PROOF. For convenience we scale the tasks weights such that $w_{max} = 1$, hence $OPT = r$. We distinguish between two cases. First, assume that $\mu \geq OPT$. By the Hoeffding bound⁴, for every $1 \leq i \leq m$, and $0 < \epsilon < 1$:

$$\begin{aligned} \Pr[L_i \geq (1 + \epsilon/3)\mu] &\leq \Pr[L_i + (\mu - E[L_i]) \geq (1 + \epsilon/3)\mu] \\ &\leq e^{-\frac{\epsilon^2\mu}{27}} \leq m^{-3}, \end{aligned}$$

where the first inequality follows since $\mu - E[L_i] \geq 0$, the second inequality results from the fact that $E[L_i + (\mu - E[L_i])] = \mu$, and the last inequality follows since $\mu \geq OPT = r$. By applying the union bound, with high probability $L_{max} \leq (1 + \epsilon/3)\mu$ and therefore $\mu_{max} \leq (1 + \epsilon/3 + o(1))\mu$. By using Theorem 9 with $\epsilon/3$ we conclude that $\mu_{max} \leq (1 + \epsilon/3 + o(1))(1 + \epsilon/3)OPT \leq (1 + \epsilon)OPT$ and the theorem follows. The case of $\mu < OPT$ is handled by using the same analysis while replacing μ by OPT . \square

3.2 Lower bounds for restricted assignment

We begin by proving a tight lower bound for the pure strategies case, and then extend it to the mixed strategies case. Our constructions are similar to those used in the proof of Theorem 1.3 in [4]. We note that our lower bounds are proved even for unit weight tasks.

Theorem 14 For pure strategies, $R = \Omega\left(\frac{\log m}{r \cdot \log(2 + \frac{\log m}{r})}\right)$.

PROOF. We construct the following problem instance:

links: we allocate $l + 1$ link groups among the m links (the value of l will be determined later) such that in group $k = 0, \dots, l$ there are $\sqrt{m} \left\lfloor \frac{l!}{r^k} \cdot \frac{r^k}{k!} \right\rfloor$ links. Denote the number of links in group k by n_k .

tasks: we partition the tasks into l groups. In group $k = 1, \dots, l$ there are $k \cdot n_k$ unit weight tasks, each can be assigned to any link from groups $[k - 1, \dots, l]$.

⁴ We use the following Hoeffding bound which generalizes the Chernoff bound: Let $\{X_i\}_{i=1}^N$ be independent random variables (not necessarily from the same distribution), such that each X_i takes a value of 0 or $z_i \leq 1$. Then for $X = \sum_{i=1}^N X_i$, and $0 \leq \delta < 1$, $\Pr[X \geq (1 + \delta)E[X]] \leq e^{-\frac{\delta^2 E[X]}{3}}$.

Observe that $OPT \leq r + 1$ for this problem instance. The optimal solution assigns the tasks of group k ($k = 1, \dots, l$) to the links in group $k - 1$, at most $r + 1$ tasks per link. We define the following system of pure strategies, denoted by S : all tasks from group k ($k = 1, \dots, l$) are assigned to links from group k , k tasks per link.

Claim 15 *The system S is in Nash equilibrium.*

PROOF. Denote by M_k ($k = 1, \dots, l$) the set of links to which tasks from task group k can be assigned. Let j be a task from group k and consider the assignment of j to link i from link group k . Clearly, $c_{ij} = k$, and for each link $t \in M_k$ we have $c_{tj} \geq (k - 1) + 1 = k \geq c_{ij}$. Hence the system S is in Nash equilibrium. \square

We now turn to bound the coordination ratio. We should satisfy the inequality $\sqrt{m} \cdot \frac{l!}{r^l} \sum_{k=1}^l \frac{r^k}{k!} \leq m$. Since $e^r \geq \sum_{k=1}^l \frac{r^k}{k!}$, it will suffice if the following inequality holds:

$$*\sqrt{m} \cdot \frac{l!}{r^l} \cdot e^r = \sqrt{m} \cdot \frac{(1 + o(1))\sqrt{2\pi l} \cdot (\frac{l}{e})^l}{r^l} \cdot e^r \leq m.$$

By taking $l = \beta r$, we conclude that $R = \Omega(\frac{l}{r}) = \Omega(\beta) = \Omega\left(\frac{\log m}{r \log(2 + \frac{\log m}{r})}\right)$. \square

By substituting $r = O(\frac{\log m}{\epsilon^2})$ in the inequality (3.2) above we derive a tight lower bound (for pure and mixed strategies) for the case where each user controls a small fraction of the total traffic.

Corollary 16 *For any $0 < \epsilon < 1$, if $r = O(\frac{\log m}{\epsilon^2})$, then $R \geq 1 + \epsilon$.*

We can modify our construction from Theorem 14 to obtain a tight lower bound for the mixed strategies case.

Theorem 17 *For mixed strategies, $R = \Omega\left(\frac{\log m}{r \cdot \log \log(2 + \frac{\log m}{r})}\right)$.*

PROOF. We slightly modify the problem instance constructed in the proof of Theorem 14. Everything remains the same except task group l now contains $(l - 1) \cdot n_l$ tasks. Clearly, $OPT \leq r + 1$. We introduce the following system of mixed strategies, denote it by S :

- All tasks from group k ($k = 1, \dots, l - 1$) are assigned to links from group k , k tasks per link (the same as the construction in the proof of Theorem 14).

- Each task from group l has uniform distribution over the links from group l .

We first prove that the system S is in Nash equilibrium.

Claim 18 *The system S is in Nash equilibrium.*

PROOF. The cost of task j from group l on any link i in group l is: $c_{ij} = (l - 1) + (1 - 1/\sqrt{m})$. On the other hand, for any link t not in group l : $c_{tj} = (l - 1) + 1 > c_{ij}$. The proof concerning tasks from groups $[1, \dots, l - 1]$ is identical to the proof given in Claim 15. \square

In the following claim we determine a lower bound on μ_{max}^S .

Claim 19 $\mu_{max}^S = \Omega\left(\frac{\log m}{\log \log(2 + \frac{\log m}{r})}\right)$.

PROOF. Consider the assignment of tasks to links in group l . There are $(l - 1) \cdot \sqrt{m}$ unit weight tasks each with uniform distribution over the \sqrt{m} links. This corresponds to a model of throwing $(l - 1) \cdot \sqrt{m}$ balls uniformly at random to \sqrt{m} bins (see e.g. [7]). In this model the expected maximum occupancy is $\Omega\left(l + \frac{\log m}{\log((\log m)/l)}\right)$. In our case this lower bound corresponds to $\mu_{max}^S = \Omega\left(\frac{\log m}{\log \log(2 + \frac{\log m}{r})}\right)$. \square

Since $OPT \leq r + 1$, our lower bound on the worst-case coordination ratio in the case of mixed strategies follows directly from Claim 19. \square

4 Analysis of the unrelated links model

Recall that in the unrelated links model, a task j is associated with an m -vector $\vec{w}_j = (w_{1j}, \dots, w_{mj})$ specifying its weight on each link. Clearly, this model generalizes the restricted assignment model. We define the maximum stretch in the system as $s = \max_{j,i,l:w_{ij} < \infty} \frac{w_{ij}}{w_{lj}}$. In the next sections we show tight bounds for pure and mixed strategies as functions of s . Note that by increasing s the worst-case coordination ratio becomes arbitrarily large.

4.1 Upper bounds for the unrelated links model

We begin by proving an upper bound on the maximum expected load in any system in Nash equilibrium.

Lemma 20 *Let S be a system in Nash equilibrium. Then $\mu^S = O\left(s + \frac{\log m}{\log(2 + \frac{\log m}{s})}\right) \cdot OPT$.*

PROOF. We order the links by non-increasing order of their expected loads. For every $k \geq 1$, define m_k to be the minimal integer such that $E[L_{m_k+1}] < k \cdot OPT$ ($m_k = m$ if there is no such integer). Let $h = \lfloor \mu^S / OPT \rfloor$. The following claim shows a relation between m_k and m_{k+1} .

Claim 21 *For every $k \geq 1$, $\frac{m_k}{m_{k+1}} \geq \frac{k+1}{s}$.*

PROOF. Denote by J_{k+1} the set of tasks assigned by S with positive probability to a link from $[1, \dots, m_{k+1}]$. Suppose that there is a task $j \in J_{k+1}$ which is assigned by OPT to a link $t > m_k$, and let $q \leq m_{k+1}$ be a link to which j is assigned by S with positive probability. Then,

$$c_{qj} \geq (k+1) \cdot OPT \geq k \cdot OPT + (1 - p_{tj})w_{tj} > c_{tj},$$

contradicting the assumption that S is in Nash equilibrium. Therefore, all the tasks in J_{k+1} are assigned by OPT to links from $[1, \dots, m_k]$. Hence,

$$\begin{aligned} m_k \cdot OPT &\geq \sum_{i=1}^{m_k} L_i^{OPT} \geq \sum_{j \in J_{k+1}} \min_{1 \leq i \leq m} w_{ij} \geq \frac{\sum_{j \in J_{k+1}} \sum_{i=1}^{m_{k+1}} p_{ij}^S \cdot w_{ij}}{s} \\ &= \frac{\sum_{i=1}^{m_{k+1}} E[L_i^S]}{s} \geq \frac{(k+1) \cdot OPT}{s} \cdot m_{k+1}, \end{aligned}$$

and the claim follows. \square

Now, if $h \leq s$ then we are done. Otherwise, using Claim 21 we obtain that

$$m \geq m_s \geq \frac{h(h-1) \cdots (s+1)}{s^{h-s}} = \frac{h!}{s! \cdot s^{h-s}} \geq \left(\frac{h}{e^{1-s/h} s}\right)^h.$$

By applying the analysis used in the proof of Theorem 6, it follows that $h = O\left(\frac{\log m}{\log(2 + \frac{\log m}{s})}\right)$. \square

The next two theorems bound the worst-case coordination ratio for pure and mixed strategies. Theorem 22 follows directly from Lemma 20.

Theorem 22 *For pure strategies* $R = O\left(s + \frac{\log m}{\log(2 + \frac{\log m}{s})}\right)$.

Theorem 23 *For mixed strategies* $R = O\left(\frac{s \cdot \log m}{\log \log m} + \frac{s \cdot \log m}{\log(s \cdot \log(2 + \frac{\log m}{s}))}\right)$.

PROOF. The proof follows the same analysis used in the proof of Theorem 10, hence we indicate only the differences. We begin by scaling the tasks weights such that $OPT = 1$. Notice that as a consequence $w_{ij} \leq s$ for any $1 \leq i \leq m$ and $1 \leq j \leq n$. Using Hoeffding inequality we derive that $\Pr[L_i \geq c \cdot \mu] \leq (\frac{e}{c})^{\frac{c \cdot \mu}{s}}$ for every $c > 1$, and by the same analysis as in Theorem 10 we obtain that $\mu_{max} \leq \alpha \cdot \mu + (2 + 1/\mu) \cdot m \mu \cdot (e/\alpha)^{\frac{\alpha \cdot \mu}{s}}$, for every integer $\alpha > 10$. By taking $\alpha = \Gamma^{-1}(m^{s/\mu}) + d$, for some sufficiently large constant d , and plugging in the bound on μ from Lemma 20, the desired bound is obtained. \square

4.2 Lower bounds for the unrelated links model

In this section we show lower bounds that match the upper bounds proved in Section 4.1

Theorem 24 *For pure strategies,* $R = \Omega\left(s + \frac{\log m}{\log(2 + \frac{\log m}{s})}\right)$.

PROOF. We prove the theorem by constructing two instances of the problem. Without loss of generality we assume that s is integral.

We first assume that $s \leq \log m$. We give a construction similar to the one in Theorem 14: We partition the links into $Ks + 1$ groups, where $K \leq \Gamma^{-1}(m^{\frac{1}{3s}})$. For $k = 0, \dots, K - 1$ and $i = 1, \dots, s$, group number $ks + i$ contains $n_{ks+i} = \sqrt{m} \cdot \frac{(K!)^s}{(k!)^s (k+1)^i}$ links, and group 0 contains $n_0 = m - \sum_{l=1}^{Ks} n_l$ links. Note that

$$\sqrt{m} \cdot (K!)^s + \sum_{l=1}^{Ks} n_l \leq \sqrt{m} \cdot (K!)^s + \sqrt{m} \cdot \sum_{k=0}^{K-1} s \frac{(K!)^s}{(k!)^s} \leq \sqrt{m} \cdot (K!)^s (1 + es) \leq m,$$

so $n_0 \geq \sqrt{m} \cdot (K!)^s$.

The tasks are partitioned into Ks groups. For $k = 0, \dots, K - 1$ and $i = 1, \dots, s$, the $(ks + i)$ -th group contains $(k + 1) \cdot n_{ks+i}$ tasks. Each task in that group

has weight 1 on a link from group $ks + i - 1$, weight $s - (s - i)/(k + 1)$ on a link from group $ks + i$, and infinite weight on all other links.

The optimal assignment is to assign each task from the l -th group of tasks to a distinct link in the $(l - 1)$ -th group of links. Thus, $OPT = 1$. Now, consider the following system S of pure strategies: The tasks of the l -th group of tasks are evenly divided between the links of the l -th group of links. Clearly, the load on a link from the l -th group is exactly l for $l \geq 1$, and in particular, the maximum load is Ks .

The system S is in Nash equilibrium: For a task j from the l -th group that was assign to a link i , we have that $c_{ij} = l$, and for any link $k \neq i$, $c_{kj} \geq (l - 1) + 1 = l$. Therefore, the coordination ratio is $\Omega(Ks)$. Since we can take $K = \Omega\left(\frac{\log m}{s \cdot \log(2 + \frac{\log m}{s})}\right)$, it follows that $R = \Omega\left(\frac{\log m}{\log(2 + \frac{\log m}{s})}\right)$.

We now handle the case when $s = \Omega(\log m)$. Consider the following problem instance: there are m tasks, where task j has weight 1 on link j , weight s on link $(j + 1) \bmod m$, and infinite weight on any other link (assume the links are numbered $[0, \dots, m - 1]$).

Clearly, $OPT = 1$ by assigning task j to link j ($j = 0, \dots, m - 1$). Consider the system of pure strategies S where task j is assigned to link $((j + 1) \bmod m)$. Clearly, the system S is in Nash equilibrium. Moreover, the load on each link is s . Hence, the coordination ratio is $\Omega(s)$. \square

Theorem 25 *For mixed strategies, $R = \Omega\left(\frac{s \cdot \log m}{\log \log m} + \frac{s \cdot \log m}{\log(s \cdot \log(2 + \frac{\log m}{s}))}\right)$.*

PROOF. We prove the theorem by constructing two instances of the problem. Without loss of generality we assume that s is integral.

We first assume that $s \leq \log m$. We use the first problem instance from the proof of Theorem 24 with the following slight modification: Each task that belongs to group $1 \leq l = ks + i \leq Ks - 1$ has weight 2 on any link from group $l - 1$ and weight $2s - 2(s - i)/(k + 1)$ on any link from group l . Observe that $OPT = 2$, by assigning each task from the l -th group of tasks to a distinct link in the $(l - 1)$ -th group of links. Now, consider the following system S of mixed strategies: For $1 \leq l \leq Ks - 1$, the tasks of the l -th group are evenly divided between the links of the l -th group of links. All tasks belonging to group Ks have uniform distribution over the links of group Ks . It can be easily verified that the system S is in Nash equilibrium. Note that the assignment of tasks to links in group Ks corresponds to throwing $K \cdot \sqrt{m}$ balls (of size s) uniformly at random to \sqrt{m} bins. In this model the expected maximum occupancy is $\Omega\left(K + \frac{\log m}{\log((\log m)/K)}\right)$. Since we can take $K = \Omega\left(\frac{\log m}{s \cdot \log(2 + \frac{\log m}{s})}\right)$, it follows that

$$R \geq \mu_{max}^S/2 = \Omega\left(\frac{s \cdot \log m}{\log(s \cdot \log(2 + \frac{\log m}{s}))}\right).$$

Next, we handle the case when $s = \Omega(\log m)$, and construct the following problem instance:

links: we partition the m links into two groups, denoted M_1 and M_2 , each consisting of $m/2$ links.

tasks: we have two task groups, J_1 and J_2 , each consisting of $m/2$ tasks. The weight vectors of the tasks are defined as follows. Task j ($j = 0, \dots, m/2 - 1$) from group J_1 has weight s on all links of group M_1 , weight 1 on link j from group M_2 , and infinite weight on all other links. Task j ($j = 0, \dots, m/2 - 1$) from group J_2 has weight 2 on link j from group M_2 , weight $2s$ on all other links in M_2 , and infinite weight on all links of group M_1 .

Observe that $OPT = 3$, by assigning task j from group J_1 together with task j from group J_2 to link j in group M_2 . Now consider the following system of mixed strategies, denoted by S . Each task from group J_1 has uniform distribution over the links of group M_1 , and task j from group J_2 is assigned (with probability 1) to link $((j + 1) \bmod m/2)$ in group M_2 . It can be easily verified that the system S is in Nash equilibrium. Since the assignment of tasks from group J_1 to link group M_1 corresponds to throwing $m/2$ balls (of size s) to $m/2$ bins, we conclude that $R \geq \mu_{max}^S/3 = \Omega(\frac{s \cdot \log m}{\log \log m})$. \square

References

- [1] R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? In *Proc. 4th Annual ACM Conference on Electronic Commerce*, pages 98–107, 2003.
- [2] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proc. 35th ACM Symp. on Theory of Computing*, pages 521–530, 2003.
- [3] A. Czumaj, P. Krysta, and B. Vöcking. Selfish traffic allocation for server farms. In *Proc. 34th ACM Symp. on Theory of Computing*, pages 287–296, 2002.
- [4] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *Proc. 13th ACM-SIAM Symp. on Discrete Algorithms*, pages 413–420, 2002.
- [5] E. Koutsoupias and C.H. Papadimitriou. Worst-case equilibria. In *Proc. 16th Symp. on Theoretical Aspects of Comp. Science*, pages 404–413, 1999.
- [6] M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proc. 33rd ACM Symp. on Theory of Computing*, pages 510–519, 2001.

- [7] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.
- [8] C.H. Papadimitriou. Algorithms, games and the Internet. In *Proc. 33rd ACM Symp. on Theory of Computing*, pages 749–753, 2001.
- [9] T. Roughgarden. Designing networks for selfish users is hard. In *Proc. 42nd IEEE Symp. on Found. of Comp. Science*, pages 472–481, 2001.
- [10] T. Roughgarden. How unfair is optimal routing? In *Proc. 13th ACM-SIAM Symp. on Discrete Algorithms*, pages 203–204, 2002.
- [11] T. Roughgarden. The price of anarchy is independent of the network topology. *J. Comput. and Syst. Sci.*, 67(2):341–364, 2003.
- [12] T. Roughgarden and É. Tardos. How bad is selfish routing. *J. of the ACM*, 49(2):236–259, 2002.