Tight Bounds for Worst-Case Equilibria

Artur Czumaj*
Department of Computer Science
New Jersey Institute of Technology
czumaj@cis.njit.edu

Berthold Vöcking†
Max-Planck-Institut für Informatik
Saarbrücken, Germany
voecking@mpi-sb.mpg.de

Abstract

The coordination ratio is a game theoretic measure that aims to reflect the price of selfish routing in a network. We show that the worst-case coordination ratio on m parallel links (of possibly different speeds) is

$$\Theta\left(\frac{\log m}{\log\log\log m}\right) .$$

Our bound is asymptotically tight and it entirely resolves an open question posed recently by Koutsoupias and Papadimitriou [3].

1 Introduction

In an attempt to measure the efficiency of *non-cooperative network systems* (like the Internet), Koutsoupias and Papadimitriou [3] proposed to investigate the relationship between the systems in which each user is aware of the situation facing all other users and trying to optimize its own strategy (that is, being in a *Nash equilibrium*) and optimal strategies in such systems. Similar questions have been already considered in game theory (see, e.g., [7]) and they aim at providing some knowledge about systems having no coordination, in which each user may behave in a greedy or anarchic way. It is well known that Nash equilibria do not always optimize the overall performance (see, e.g., the Prisoner's dilemma [8]). Therefore, in order to understand the phenomenon of non-cooperative systems Koutsoupias and Papadimitriou started investigations of the behavior of the worst-case *coordination ratio*, which is the ratio between the worst possible Nash equilibrium and the social (i.e., overall) optimum. In other words, this analysis seeks the price of uncoordinated individual utility-maximizing decisions ("the price of anarchy"). Koutsoupias and Papadimitriou [3, 7] proposed to investigate the coordination ratio for routing problems in which a set of several agents want to send some traffic from sources to destinations.

In this paper, we address the most basic case of a routing problem. As proposed by Koutsoupias and Papadimitriou [3], we consider a network that consist of m *parallel links* $1, 2, \ldots, m$ from an origin to a destination, all with possibly different speeds $s_1, \ldots, s_m$. There are n agents $1, 2, \ldots, n$, each having an amount of traffic w_i to send from the origin to the destination (we assume throughout the paper that all w_i are non-negative). Each agent i, $i \in \{1, \ldots, n\}$, sends the traffic using a possibly randomized *mixed strategy*, with p_i^j denoting the probability that agent i sends the entire traffic w_i to a link j. We assume the agents are *selfish* in the sense that each of them tries to minimize its individual cost. Assuming each agent is aware of the strategies of the other agents and behaves in a non-cooperative and selfish way, the system will come to a *Nash equilibrium*, i.e., a combination of typically randomized choices (mixed strategies) from which no agent has an incentive to deviate.

We notice that the model considered in this paper is a simplification of the problems arising in real networks. However, as pointed out in [3, 5, 7], this model seems to be appropriated to describe several basic networking problems. We believe that understanding the ratio between worst possible Nash quilibrium and the social optimum in simple situations is necessary for making rigorous analyzes in more complicated networks. Readers interested in more detailed exposition of this model and in its applications are referred to [3, 5, 7, 9].

The model. We define now our model formally trying to follow the notation used by Koutsoupias and Papadimitriou [3].

The routing model described above can be formally defined as an allocation problem with m independent links with speeds $s_1, \ldots, s_m$ and n independent tasks with weights $w_1, \ldots, w_n$. The goal is to allocate the tasks to the links to minimize the maximum load of the links in the system.

We use the standard notation $[N]$ to denote set $\{1, \ldots, N\}$. The set of *pure strategies* for task i is therefore $[m]$ and a *mixed strategy* is a distribution on this set.

* Supported in part by NSF grant CCR-0105701 and by SBR grant No. 421090.

† Supported in part by DFG grant Vo889/1-1 and the IST program of the EU under contract number IST-1999-14186 (ALCOM-FT).

Let $(j_1, \ldots, j_n) \in [m]^n$ be a combination of pure strategies, one for each task, its *cost* for task i is

$$\sum_{j_k=j_i} \frac{w_k}{s_{j_i}} ,$$

which is the time needed for link j_i chosen by task i to complete all tasks allocated to that link[1].

Similarly, for a combination of pure strategies $(j_1, \ldots, j_n) \in [m]^n$, the *load* of link j is defined as

$$\sum_{j_k=j} \frac{w_k}{s_j} .$$

Given n tasks of length $w_1, \ldots, w_n$ and m links of speed $s_1, \ldots, s_m$, let $\mathtt{opt}$ denote the *social optimum*, that is, the minimum cost of a pure strategy:

$$\mathtt{opt} := \min_{(j_1,\ldots,j_n)\in[m]^n} \max_{j\in[m]} \sum_{i:j_i=j} \frac{w_i}{s_j} .$$

For example, if all links have the same unit speed ($s_j = 1$ for every $j \in [m]$) and all weights are the same ($w_i = 1$ for every $i \in [n]$), then the social optimum is $\lceil \frac{n}{m} \rceil$. Furthermore, it is easy to see that in any system $\mathtt{opt} \geq \frac{\max_i w_i}{\max_j s_j}$. In general, however, computing the social optimum is $\mathcal{NP}$-hard even for identical speeds (see [3]).

Let p_i^j denote the probability that agent i sends the entire traffic w_i to a link j, and let ℓ_j denote the *expected load* on link j, that is,

$$\ell_j := \sum_{i\in[n]} \frac{w_i \, p_i^j}{s_j} .$$

For a task i, the *expected cost of task* i *on link* j (or its *finish time* when its load w_i is allocated to link j [3]) is equal to

$$c_i^j := \frac{w_i}{s_j} + \sum_{t\neq i} \frac{w_t \, p_t^j}{s_j} = \ell_j + (1-p_i^j)\,\frac{w_i}{s_j} .$$

DEFINITION 1.1. **(Nash equilibrium)** *The probabilities* $(p_i^j)_{i\in[n],j\in[m]}$ *define a* Nash equilibrium *if and only if any task* i *will assign non-zero probabilities only to links that minimize* c_i^j*, that is,* $(p_i^j) > 0$ *implies* $c_i^j \leq c_i^q$*, for every* $q \in [m]$.

In words, a Nash equilibrium is characterized by the property that there is no incentive for any task to change its strategy. As an example, we observe that in the system considered above in which all links have the same unit speed and all weights are the same, the uniform probabilities $p_i^j = \frac{1}{m}$ for all $j \in [m]$ and $i \in [n]$ define a system in a Nash equilibrium.

For the rest of the paper, we fix an arbitrary Nash equilibrium, that is, fix the probabilities $(p_i^j)_{i\in[n],j\in[m]}$ that define a Nash equilibrium. Let us consider the randomized allocation strategies in which each task i is allocated to a single link chosen independently at random according to the probabilities p_i^j, that is, task i is allocated to link j with probability p_i^j. Let C_j, $j \in [m]$, be the random variable indicating the *load of link* j in our random experiment. We observe that C_j is the weighted sum of independent 0-1 random variables J_i^j, $\mathbf{Pr}[J_i^j = 1] = p_i^j$, such that $C_j = \frac{1}{s_j} \sum_{i=1}^n w_i \cdot J_i^j$.

Let c denote the *maximum expected load* over all links, that is,

$$c := \max_{j\in[m]} \ell_j .$$

Notice that $\mathbf{E}[C_j] = \ell_j$, and therefore $c = \max_{j\in[m]} \mathbf{E}[C_j]$.

Finally, we define the *social cost* C to be the expected maximum load (instead of maximum expected load), that is,

$$C := \mathbf{E}[\max_{j\in[m]} C_j] .$$

Observe that $c \leq C$ and possibly $c \ll C$. Recall that $\mathtt{opt}$ denotes the *social optimum* (i.e., the minimum cost of a pure strategy). In this paper our main focus is on estimating the *coordination ratio* which is the worst-case ratio

$$R := \max \frac{C}{\mathtt{opt}} ,$$

where the maximum is over all Nash equilibria.

We shall use also standard notation to denote by $\Gamma(N)$ the *Gamma (factorial) function* (which for any natural number N is defined by $\Gamma(N+1) = N!$ and for an arbitrary real number $x > 0$ is defined as $\Gamma(x) = \int_0^\infty t^{x-1}\, e^{-t}\, dt$, see also `http://functions.wolfram.com/GammaBetaErf/Gamma/`). We shall use frequently the inverse of the Gamma function, $\Gamma^{(-1)}(N)$, where for our applications we shall use the fact that $\Gamma^{(-1)}(N) = x$ such that $\lfloor x \rfloor! \leq N - 1 \leq \lceil x \rceil!$. We notice that it is well known that $\Gamma^{(-1)}(N) = \frac{\log N}{\log\log N}(1+o(1))$ and that $(\alpha/e)^\alpha = N$ for [2] $\alpha = \Gamma^{(-1)}(N) + \Theta(1)$.

1.1 Previous results.

Koutsoupias and Papadimitriou [3] initiated the study of the worst-case coordination ratio and show the following results for networks consisting of m parallel links:

[1] In the original formulation of Koutsoupias and Papadimitriou [3], an additional additive term L^{j_i} was used. However, since in all papers we are aware of all analyzes assumed $L^{j_i} = 0$, we skipped that term in our presentation. We want to point out, however, that our bounds are not affected by these additive terms.

[2] To see this, observe, for example, Stirling's formula for $N! = (N/e)^N \cdot \sqrt{2\pi N} \cdot (1 + o(1))$, from which one can easily derive the required inequality.

- For two identical links the worst-case coordination ratio is exactly $\frac{3}{2}$.
- For two links (not necessarily identical, that is, with possibly different speeds) the worst-case coordination ratio is $\phi = \frac{1+\sqrt{5}}{2}$.
- For m identical links the worst-case coordination ratio is $\Omega(\frac{\log m}{\log \log m})$ and it is at most $3 + \sqrt{4\, m\, \ln m}$.
- The worst-case coordination ratio for any number of tasks and m (not necessarily identical) links is $\mathcal{O}(\sqrt{\frac{s_1}{s_m} \sum_{j=1}^{m} \frac{s_j}{s_m}} \sqrt{\log m})$, where s_j is the speed of link j, and $s_1 \geq s_2 \geq \cdots \geq s_m$.

Mavronicolas and Spirakis [5] greatly extended some of the bounds above and show the following results in the so-called *fully-mixed model*[3]:

- For m identical links in the fully-mixed Nash equilibrium the worst-case coordination ratio is $\Theta(\frac{\log m}{\log \log m})$ (see also [1]).
- For m (not necessarily identical) links and n identical weights in the fully-mixed Nash equilibrium, if $m \leq n$, then the worst-case coordination ratio is $\Theta(\frac{\log n}{\log \log n})$.

We emphasize that besides the very special case of $m = 2$ parallel links, no asymptotically tight results have been known even for systems with identical links. In particular, even the main conjecture from the work of Koutsoupias and Papadimitriou [3] that the worst-case coordination ratio for m *identical links* is $\Theta(\frac{\log m}{\log \log m})$ has remained unproved prior to our work.

1.2 New results.

Our first result is an upper bound for the worst-case coordination ratio.[4]

THEOREM 1.1. *The coordination ratio for m parallel links is bounded from above by*

$$\mathcal{O}\left(\min\left\{\frac{\log m}{\log\log\log m},\ \frac{\log m}{\log\left(\frac{\log m}{\log(s_1/s_m)}\right)}\right\}\right) ,$$

where it is assumed that the speeds satisfy $s_1 \geq \cdots \geq s_m$.

In particular, the worst-case coordination ratio for m parallel links is

$$\mathcal{O}\left(\frac{\log m}{\log\log\log m}\right) .$$

[3]The *fully-mixed model* is a special class of Nash equilibria in which all p_i^j are non-zero.

[4]To simplify the notation, throughout the entire paper, for any non-negative real x we shall use $\log x$ to denote $\log x = \max\{\log_2 x, 1\}$.

The theorem follows directly from the following two lemmas.

LEMMA 1.1. *The maximum expected load c satisfies*

$$c = \mathrm{opt} \cdot \mathcal{O}\left(\min\left\{\frac{\log m}{\log\log m},\ \log\left(\frac{s_1}{s_m}\right)\right\}\right) .$$

where it is assumed that the speeds satisfy $s_1 \geq \cdots \geq s_m$.

LEMMA 1.2. *The social cost C satisfies*

$$C = \mathrm{opt} \cdot \mathcal{O}\left(\frac{\log m}{\log\left(\frac{\mathrm{opt}\cdot\log m}{c}\right)}\right) .$$

Observe that in the special case when all links are identical, the coordination ratio is $\mathcal{O}\left(\frac{\log m}{\log\log m}\right)$ by Theorem 1.1. Recently, and independently of our work, also Koutsoupias et al. [4] obtained the same upper bound. However, in this special case we get a much stronger bound that is actually tight up to an additive constant.

THEOREM 1.2. *For m identical links the worst-case coordination ratio is at most*

$$\Gamma^{(-1)}(m) + \Theta(1) = \frac{\log m}{\log\log m} \cdot (1 + o(1)) .$$

This bound improves upon the result due to Mavronicolas and Spirakis [5], not only by extending the class of Nash equilibria for which the upper bound holds, but also by tightening the result up to a constant additive factor. Indeed, as it was observed by Koutsoupias and Papadimitriou [3] and by Mavronicolas and Spirakis [5], one can obtain a lower bound for the worst-case coordination ratio for m identical links by considering the system in which all $n = m$ tasks have $p_i^j = \frac{1}{m}$ for every $i, j \in [n]$; the classical result of Gonnet [2] implies that in such a system the worst-case coordination ratio is $\Gamma^{(-1)}(m) - \frac{3}{2} + o(1)$.

Furthermore, we prove that our analysis of the upper bound in Theorem 1.1 is asymptotically tight.

THEOREM 1.3. *The coordination ratio for m parallel links is lower bounded by*

$$\Omega\left(\min\left\{\frac{\log m}{\log\log\log m},\ \frac{\log m}{\log\left(\frac{\log m}{\log(s_1/s_m)}\right)}\right\}\right) .$$

In particular, the worst-case coordination ratio for m parallel links is

$$\Omega\left(\frac{\log m}{\log\log\log m}\right) .$$

In fact, we will show, analogously to the upper bound, that, for every positive real r and every $S \geq 1$, there exists a set of m links with $\frac{s_1}{s_m} = S$ having a Nash equilibrium satisfying (i) $\mathrm{opt} = r$,

$$(ii) \quad c = \mathrm{opt} \cdot \Omega\left(\min\left\{\frac{\log m}{\log\log m}, \log\left(\frac{s_1}{s_m}\right)\right\}\right) \quad ,$$

and

$$(iii) \quad C = \mathrm{opt}{\cdot}\Omega\left(\frac{\log m}{\log\left(\frac{\mathrm{opt}\cdot\log m}{c}\right)}\right) \quad .$$

Combining Theorems 1.1 and 1.3 we obtain asymptotically tight bound for the worst-case coordination ratio for m parallel links.

2 Upper Bound: Proof of Theorem 1.1

In this section we prove Lemmas 1.1 and 1.2, from which Theorem 1.1 directly follows.

2.1 Proof of Lemma 1.1. Fix an arbitrary Nash equilibrium, that is, fix the probabilities $(p_i^j)_{i\in[n],j\in[m]}$ that define a Nash equilibrium. Without loss of generality, assume $s_1 \geq s_2 \geq \cdots \geq s_m$. Let us normalize (scale) the weights of all tasks such that $\mathrm{opt} = 1$. Under this normalization, we have to show that $c = \mathcal{O}(\frac{\log m}{\log\log m})$ and $c = \mathcal{O}\left(\log\left(\frac{s_1}{s_m}\right)\right)$. We prove these bounds in two separate Lemmas 2.1 and 2.2.

LEMMA 2.1. $c \leq \Gamma^{(-1)}(m) + 1 = \frac{\log m}{\log\log m}\,(1 + o(1))$.

Proof. For $k \geq 1$, define j_k to be the smallest index in $\{0, 1, \ldots, m\}$ such that $\ell_{j_k+1} < k$ or, if no such index exists, $k = m$. Let us observe that the following properties hold:

- for every $k \geq 1$ with $0 < j_k \leq m$, all links $j \leq j_k$ have load at least k, and
- for every $k \geq 1$ with $0 \leq j_k < m$, link $j_k + 1$ has load less than k.

Let $c^* = \lfloor c - 1 \rfloor$. We will show that $j_1 \geq c^*!$. Combining this inequality with the obvious constraint $j_1 \leq m$ will imply an appropriate upper bound on c.

In order to estimate j_1, we start with estimating j_{c^*}. Observe that link 1 does not need to be the link with highest expected load. The following claim, however, shows that ℓ_1 is close to c^*.

CLAIM 2.1. $j_{c^*} \geq 1$*, and hence* $\ell_1 \geq c^*$.

Proof. For the purpose of contradiction, assume $j_{c^*} = 0$. This implies that $\ell_1 < c^* \leq c - 1$. Let q denote the link with the maximum expected load. Then $\ell_1 + 1 < c = \ell_q$.

We observe that all the tasks that have positive probability on q must have weight larger than s_1. Indeed, if one such a task i had weight $w_i \leq s_1$, then it would have expected cost on link 1 to be at most $\ell_1 + \frac{w_i}{s_1} \leq \ell_1 + 1 < \ell_q$, which contradicts the assumption that all tasks with positive probability on q have minimal expected cost on q in Nash equilibrium.

Thus, we have shown the existence of a task i of weight $w_i > s_1$. Recall that link 1 is the fastest link and its speed is s_1. Therefore the cost of allocating task i to any link will lead to the cost that is bigger than 1. Consequently, $\mathrm{opt} > 1$, which contradicts our initial assumption that $\mathrm{opt} = 1$. This completes the proof of Claim 2.1. □

The next claim gives a lower bound on j_k in terms of j_{k+1}.

CLAIM 2.2. *For* $k \geq 1$, $j_k \geq (k+1)\, j_{k+1}$.

Proof. Let T be the set of tasks in the system that have positive probability on at least one of the links in $\{1, \ldots, j_{k+1}\}$. Fix an optimal allocation strategy OPTStr . We distinguish between two different ways of how OPTStr might allocate the tasks in T to the links.

Case 1: Suppose OPTStr allocates at least one of the tasks in T to a link j, $j > j_k$. We will show that this implies $\mathrm{opt} > 1$ and hence contradicts our assumptions.

Let W_T denote the minimum weight of the tasks in T. We first derive a lower bound on W_T. The expected load of the links in $\{1, \ldots, j_{k+1}\}$ (and hence of all links in T) is at least $k + 1$. The expected load of link $j_k + 1$ is less than k. Therefore, the requirement of Nash equilibria yields $W_T > s_{j_k+1}$. But this implies that allocating a task from T to link j gives cost at least $\frac{W_T}{s_j} \geq \frac{W_T}{s_{j_k+1}} > 1$ which yields $\mathrm{opt} > 1$ and hence a contradiction .

Case 2: Now let us assume OPTStr allocates all tasks in T to the links in $\{1, \ldots, j_k\}$. We will show that this implies $j_k \geq (k+1)\, j_{k+1}$.

Let $W_{\Sigma T}$ denote the sum of the weights of the tasks in T. On one hand, we observe that $W_{\Sigma T}$ is lower bounded by the sum of the expected weight on the links $\{1, \ldots, j_{k+1}\}$, that is,

$$\begin{aligned} W_{\Sigma T} = \sum_{i\in T} w_i &\geq \sum_{i\in T} w_i \cdot \sum_{j=1}^{j_{k+1}} p_i^j = \sum_{j=1}^{j_{k+1}} \sum_{i\in T} w_i\, p_i^j \\ &= \sum_{j=1}^{j_{k+1}} \sum_{i=1}^{n} w_i\, p_i^j = \sum_{j=1}^{j_{k+1}} \ell_j\, s_j \quad . \end{aligned}$$

Therefore, since $\ell_j \geq k+1$ for all $j \leq j_{k+1}$, we obtain

$$W_{\Sigma T} \geq \sum_{j=1}^{j_{k+1}} \ell_j\, s_j \geq (k+1) \sum_{j=1}^{j_{k+1}} s_j \quad .$$

On the other hand, since we assumed that $\mathtt{opt} = 1$ and OPTStr allocates all tasks in T to the links $\{1, \ldots, j_k\}$, we obtain

$$W_{\Sigma T} \leq \sum_{j=1}^{j_k} s_j \ .$$

Combining these inequalities gives $\sum_{j=1}^{j_k} s_j \geq (k+1) \sum_{j=1}^{j_{k+1}} s_j$. Since the sequence of link speeds is non-increasing, this implies that $j_k \geq (k+1)\, j_{k+1}$. This completes the proof of Claim 2.2. □

Finally, we combine the two claims above and obtain

$$j_1 \geq (c^*)! \cdot j_{c^*} \geq (c^*)! \ .$$

By definition, $j_1 \leq m$. Consequently $(c^*)! \leq m$, which implies $c \leq \Gamma^{(-1)}(m) + 1 = \frac{\log m}{\log\log m}(1 + o(1))$. This completes the proof of Lemma 2.1. □

Next, we prove another upper bound for the maximum expected cost c.

LEMMA 2.2. $c = \mathcal{O}(\log(s_1/s_m))$.

Proof. The following claim shows that the speeds of the links $j_1, j_2, \ldots$ increase in a geometric fashion.

CLAIM 2.3. *For* $1 \leq k \leq c - 3$, $s_{j_{k+2}+1} \geq 2\, s_{j_k+1}$.

Proof. Fix an optimal strategy OPTStr . Notice that every link $j' \leq j_{k+2}$ has cost $\ell_{j'} \geq (k+2) > 1 = \mathtt{opt}$. Therefore, OPTStr has to allocate at least one of the tasks that have positive probability on one of the links $1, \ldots, j_{k+2}$ to the links $j_{k+2} + 1 \ldots, m$. (Observe that in Section 2.1 it is shown that $j_{\lfloor c-1 \rfloor} \geq 1$. Hence, the existence of link $j_{k+2} \geq j_{\lfloor c-1 \rfloor}$ is guaranteed.) Clearly, such a task can have weight at most $s_{j_{k+2}+1}$ because otherwise the cost of OPTStr would be larger than $\mathtt{opt}$. Therefore, there exists a link $j \in \{1, \ldots, j_{k+2}\}$ and a task i of weight $w_i \leq s_{j_{k+2}+1}$ with $p_i^j > 0$.

On one hand, the expected cost of task i on link j in the Nash equilibrium is at least $k+2$ because, for $j \leq j_{k+2}$, we have $c_i^j \geq \ell_j \geq k+2$. On the other hand, the expected cost of task i on link $j_k + 1$ is $c_i^{j_k+1} < k + w_i/s_{j_k+1}$. Now, the Nash equilibrium property implies that the cost of task i on link j is not larger than on $j_k + 1$. Consequently, $k + 2 \leq k + w_i/s_{j_k+1} \leq k + s_{j_{k+2}+1}/s_{j_k+1}$. Clearly, this inequality implies that $2 s_{j_k+1} \leq s_{j_{k+2}+1}$ and hence, Claim 2.3 is shown. □

Claim 2.3 says that in a Nash equilibrium the speeds decrease geometrically with the cost. This implies that

$$s_m \leq s_{j_1} \leq 2^{-(c-5)/2} \cdot s_{c-1} \leq 2^{-(c-5)/2} \cdot s_1 \ .$$

Thus, $c \leq 2\,\log(s_1/s_m) + \mathcal{O}(1)$. This completes the proof of Lemma 2.2. □

We conclude the proof of Lemma 1.1 by observing that it follows directly from Lemmas 2.1 and 2.2. □

2.2 Proof of Lemma 1.2. Without loss of generality, let us assume that $\mathtt{opt} \geq 1$ and that $s_1 \geq s_2 \geq \cdots \geq s_m$. Then, we have to show that $C = \mathcal{O}\left(\frac{\log m}{\log(\log m\ \mathtt{opt}/c)}\right)$. Recall that C_j is a random variable describing the load on link j. We have $\mathbf{E}[C_j] = \ell_j \leq c$ and $C = \mathbf{E}[\max_{j \in [m]} C_j]$. Thus, we need to show, for every $j \in [m]$, that it is unlikely that C_j deviates much from its expectation. For this purpose, we will use a Hoeffding bound. In order to apply this bound, we need to show that the weights of the tasks assigned to link j cannot be much larger than s_j. This is shown in the next lemma.

LEMMA 2.3. *For every link* j *and every task* i *with* $p_i^j \in (0, \frac{1}{4}]$, $w_i \leq 12\, s_j\, \mathtt{opt}$.

Proof. Previously, in the proof of Lemma 1.1, we defined indices j_r provided $\mathtt{opt} = 1$. Now, we extend this definition to for arbitrary $\mathtt{opt}$ in natural way: for $k \geq 1$, we define j_k as the smallest index in $\{0, 1, \ldots, m\}$ such that $\ell_{j_k+1} < k \cdot \mathtt{opt}$, or, $k = m$, if no such index exists.

With this modification, we first observe that Claim 2.3 holds without any change. Therefore, we apply Claim 2.3 to show that $w_i \leq 12\, s_j\, \mathtt{opt}$ for $p_i^j \in (0, \frac{1}{4}]$. First, assume that $j \in \{j_k + 1, \ldots, j_{k-1}\}$ for some $k \leq c - 3$. Then, on one hand, the expected cost of task i on link j is

$$c_i^j = \ell_j + (1 - p_i^j)\frac{w_i}{s_j} \geq (k-1) \cdot \mathtt{opt} + \frac{3\,w_i}{4\,s_j} \ ,$$

because $\ell_j \geq (k-1)\,\mathtt{opt}$ and $1 - p_i^j \geq \frac{3}{4}$. On the other hand, the expected cost of task i on link $j_{k+2} + 1$ is

$$\begin{aligned} c_i^{j_{k+2}+1} &\leq \ell_{j_{k+2}+1} + \frac{w_i}{s_{j_{k+2}+1}} \\ &\leq (k+2) \cdot \mathtt{opt} + \frac{w_i}{2\,s_j} \ , \end{aligned}$$

by applying $\ell_{j_{k+2}+1} \leq (k+2)\,\mathtt{opt}$, $s_{j_{k+2}+1} \geq 2\, s_{j_k+1}$ (Claim 2.3), and $s_{j_k+1} \geq s_j$. Since we assume a Nash equilibrium, the cost of task i on link j cannot be larger than the cost of task i on link $j_{k+2} + 1$. Consequently, $(k-1)\,\mathtt{opt} + \frac{3\,w_i}{4\,s_j} \leq (k+2)\,\mathtt{opt} + \frac{w_i}{2\,s_j}$, which implies $w_i \leq 12\, s_j\, \mathtt{opt}$.

It remains to investigate the case $j \leq j_k$, where $k = \lfloor c - 3 \rfloor$. We observe that the expected cost of task i on the fastest link s_1 is at most $c \cdot \mathtt{opt} + \mathtt{opt} = (c+1) \cdot \mathtt{opt}$. The expected cost of task i on link j, however, is at least $k \cdot \mathtt{opt} + \frac{3\,w_i}{4\,s_j} \geq (c-4) \cdot \mathtt{opt} + \frac{3\,w_i}{4\,s_j}$. Hence, in this case $w_i \leq \frac{20}{3}\, s_j\, \mathtt{opt}$. This proves Lemma 2.3. □

Now we will apply the lemma in order to show that for every $j \in [m]$, it is unlikely that C_j deviates much from its expectation. First, let us focus on a single link. Fix a link j, $j \in [m]$. Let $T_j^{(1)}$ denote the set of tasks with $p_i^j \in (0, \frac{1}{4}]$ and $T_j^{(2)}$ the set of tasks with $p_i^j \in (\frac{1}{4}, 1]$. Furthermore, let $C_j^{(1)}$ and $C_j^{(2)}$ denote random variables that describe the cost on link j only counting tasks in $T_j^{(1)}$ and $T_j^{(2)}$, respectively. Clearly, only tasks with $p_i^j > 0$ can be allocated to link j. Hence, $C_j = C_j^{(1)} + C_j^{(2)}$.

First, let us consider the tasks in $T_j^{(1)}$ only. Recall that C_j is defined as the weighted sum of independent 0-1 random variables J_i^j, $\mathbf{Pr}[J_i^J = 1] = p_i^j$, such that $C_j = \sum_{i=1}^{n} \frac{w_i}{s_j} \cdot J_i^j$. Thus, $C_j^{(1)}$ is a weighted sum of independent 0-1 random variables as well. Notice that $c \geq \mathrm{opt}$. Next, by Lemma 2.3, we can upper bound the maximum weight in this sum by $max_{i \in T_j^{(1)}} \frac{w_i}{s_j} \leq 12\,\mathrm{opt}$. Hence, we can apply a Hoeffding bound[5] to obtain

$$\begin{aligned}
\mathbf{Pr}[C_j^{(1)} \geq \alpha\, c] \quad &\leq \quad \left(\frac{e \cdot \mathbf{E}[C_j^{(1)}]}{\alpha \cdot c} \right)^{\alpha\, c/(12\,\mathrm{opt})} \\
&\leq \quad \left(\frac{e \cdot \ell_j}{\alpha \cdot c} \right)^{\alpha\, c/(12\,\mathrm{opt})} \\
&\leq \quad (e/\alpha)^{\alpha\, c/(12\,\mathrm{opt})} \quad ,
\end{aligned}$$

for every $\alpha > 1$. Now, let us consider the tasks in $T_j^{(2)}$. Since $p_i^j \geq \frac{1}{4}$ for every $i \in T_j^{(2)}$, we immediately obtain $C_j^{(2)} \leq 4\,\mathbf{E}[C_j^{(2)}] \leq 4\,c$. As a consequence,

$$\mathbf{Pr}[C_j \geq (\alpha{+}4)\, c] \;\leq\; (e/\alpha)^{\alpha\, c/(12\,\mathrm{opt})} \;, \text{ for every } \alpha > 1 \;.$$

Until now we focused on a single, fixed link. Summing over all m links, by the union bound, the probability that the maximum cost $\mathcal{L} = \max_{j \in [m]} C_j$ does exceed $(\alpha + 4)\, c$ can be bounded from above by $m \cdot (e/\alpha)^{\alpha\, c/(12\,\mathrm{opt})}$. Recall that C is defined to be the expectation of the maximum cost over all links. Hence, for every integer $\lambda > 1$ we can estimate C as follows.

$$\begin{aligned}
C \;=\; \mathbf{E}[\mathcal{L}] \quad &\leq \quad (4 + \lambda)\, c + \sum_{\alpha=\lambda}^{\infty} \mathbf{Pr}[\mathcal{L} \geq (\alpha + 4)\, c] \\
&\leq \quad (4 + \lambda)\, c + m \cdot \sum_{\alpha=\lambda}^{\infty} (e/\alpha)^{\alpha\, c/(12\,\mathrm{opt})} \\
&\leq \quad (4 + \lambda)\, c + 2 \cdot m \cdot (e/\lambda)^{\lambda\, c/(12\,\mathrm{opt})} \quad ,
\end{aligned}$$

where the last inequality holds for $\lambda \geq r_0$, where r_0 is a sufficiently large constant. Finally, if we substitute $\lambda = \Gamma^{(-1)}(m^{12\,\mathrm{opt}/c}) + k_0$ for a suitably large constant k_0, we obtain $(e/\lambda)^{\lambda\, c/(12\,\mathrm{opt})} \leq \frac{1}{m}$. As a consequence,

$$C \;\leq\; (\lambda + 6)\, c \;.$$

Finally, we observe that in our case $\lambda = \Theta\left(\frac{\mathrm{opt} \cdot \log m}{c \cdot \log(\mathrm{opt}\, \log m / c)}\right)$, to conclude that

$$C \quad = \quad \mathrm{opt} \cdot \mathcal{O}\left(\frac{\log m}{\log\left(\frac{\mathrm{opt}\, \log m}{c}\right)} \right) \quad . \qquad \square$$

2.3 Extension of Lemma 1.2 for identical links: Proof of Theorem 1.2.

It is easy to simplify the proof and to improve Lemma 1.2 when all links are identical, that is, all s_j are the same. In that case, one can assume without loss of generality that $\mathrm{opt} = 1$ and $s_j = 1$ for every $j \in [m]$. Let T^j, $j \in [m]$, denote the set of all tasks i with $p_i^j > 0$. Given that, we can show the following lemma.

LEMMA 2.4. *In the systems with identical links it holds that* $p_i^j w_i \geq \ell_j - 1$ *for all* $j \in [m], i \in T^j$.

Proof. We use similar arguments as in the proof of Lemma 2.3. The cost of task i on link j is $c_i^j = \ell_j + (1 - p_i^j)\, w_i$. Let q be any link with $\ell_q \leq \frac{1}{m} \sum_{r \in [j]} \ell_r$. Clearly, $\ell_q \leq 1$ and hence $c_i^q \leq 1 + w_i$. Now, the lemma follows from the requirement $c_i^j \leq c_i^q$ of Nash equilibria. $\square$

We consider two separate cases. Suppose first that $\ell_j - 1 \geq 2/\Gamma^{(-1)}(m)$. Notice that Lemma 2.4 implies that $\ell_j = \sum_{i \in T^j} p_i^j w_i \geq |T^j|\,(\ell_j - 1)$. Since we have $\ell_j \leq 2$ and $1/(\ell_j - 1) \leq \Gamma^{(-1)}(m)/2$, we obtain $|T^j| \leq \ell_j/(\ell_j - 1) \leq \Gamma^{(-1)}(m)$. This inequality immediately implies $C_j \leq \Gamma^{(-1)}(m)$ because in this case at most $\Gamma^{(-1)}(m)$ tasks have positive probability on link j. The other case we have to consider is when $\mathbf{E}[C_j] = \ell_j \leq 1 + 2/\Gamma^{(-1)}(m)$. Here, applying Hoeffding bound in the same way as it is done in the proof of Lemma 1.2 yields $C_j \leq \Gamma^{(-1)}(m) + \mathcal{O}(1)$, with probability at least $1 - \frac{1}{m}$. This easily implies Theorem 1.2. $\square$

3 Lower Bound: Proof of Theorem 1.3

This section is devoted to the proof of Theorem 1.3, which states that our upper bounds proven in the previous section are essentially tight.

Our analysis follows a course similar to the one for the upper bound in the previous section. First, we will describe a mixed strategy in Nash equilibrium with $\mathrm{opt} = \Theta(1)$ and $c = \Theta\left(\frac{\log m}{\log \log m}\right)$. Then, we apply a stochastic analysis

[5] In this paper we use the following standard version of Hoeffding bound: *Let* $X_1, \ldots, X_N$ *be independent random variables with values in the interval* $[0, z]$ *for some* $z > 0$*, and let* $X = \sum_{i=1}^{N} X_i$*, then for any* t *it holds that* $\mathbf{Pr}[\sum_{i=1}^{N} X_i \geq t] \leq (e \cdot \mathbf{E}[X]\;/\;t)^{t/z}$.

showing $C = c \cdot \Theta\left(\frac{\log m}{\log((\log m)/c)}\right)$. Finally, we will take into account also the speeds of the links in our construction. Combining these bounds yields the theorem.

In fact, our construction can be easily generalized to show that for every positive real r and every $S \geq 1$, there exists a set of m links with $\frac{s_1}{s_m} = S$ having a Nash equilibrium satisfying $c = \Omega\left(\mathrm{opt} \cdot \min\left\{\frac{\log m}{\log\log m},\ \log\left(\frac{s_1}{s_m}\right)\right\}\right)$, $C = \mathrm{opt} \cdot \Omega\left(\frac{\log m}{\log\left(\frac{\mathrm{opt}\cdot\log m}{c}\right)}\right)$, and $\mathrm{opt} = r$.

3.1 Lower Bound for Pure Strategies. We start by defining a pure strategy $\mathcal{S}$ that we will transform afterwards into a mixed strategy $\mathcal{S}'$. Without loss of generality, let $\sqrt{m}$ be an integer. We consider $K+1$ groups of links $0, 1, \ldots, K$, for a suitable K to be defined later. The groups are defined as follows:

- for $1 \leq k \leq K$, the number of links in group k is equal to $\sqrt{m} \cdot \frac{K!}{k!}$ (notice that for $1 \leq k < K$ the number of links in group k is exactly $(k+1)$ times the number of links in group $k+1$),
- the number of links in group 0 is at least $\sqrt{m} \cdot K!$,
- for $0 \leq k \leq K$, the speed of the links in group k is 2^k,
- for $0 \leq k \leq K$, for each link in group k, there are exactly k tasks of weight 2^k each having probability one to be allocated to this link.

In our construction K can be chosen to be any positive integer that satisfies $\sqrt{m} \cdot \sum_{k=0}^{K} \frac{K!}{k!} \leq m$. Thus, in particular, our analysis can be carried over for all K satisfying $\sqrt{m} \cdot K! \cdot e \leq m$, and hence, for all $K \leq \Gamma^{(-1)}(\sqrt{m}/e) - 1$.

LEMMA 3.1. *Strategy $\mathcal{S}$ satisfies the following properties:*

1. *the maximum load is* $c = K$,
2. *the social optimum is* $1 \leq \mathrm{opt} \leq 2$, *and*
3. *the system is in Nash equilibrium.*

Proof.

1. This property follows from the fact that if a link j is in group k then its load is $C_j = \frac{k\,2^k}{s_j} = k$.
2. The social optimum cost 2 can be achieved, for example, by allocating all tasks "assigned" to the links in group k, $k \geq 1$, to the links in group $k-1$. Observe that there are exactly $k \cdot \sqrt{m} \cdot \frac{K!}{k!}$ tasks assigned initially to the links in group k, $k > 0$ (and zero tasks assigned to the links in group 0) and each such a task has weight 2^k. On the other hand, there is at least the same number $\sqrt{m} \cdot \frac{K!}{(k-1)!}$ of links in group $k-1$, each link with speed 2^{k-1}. Therefore, if we allocate each task from group k to a single link in group $k-1$, then since the weight of each task is 2^k and the speed of each link is 2^{k-1}, the cost of every link in the system is at most 2. Hence, the social optimum is at most 2.

   To see the lower bound for opt, let us observe that any task i in group K has weight $w_i = 2^K$ and the fastest link j has speed $s_j = 2^K$. This implies that the social optimum cannot be smaller than $\frac{w_i}{s_j} = 1$.
3. Let us take any task i that is allocated to link r in group $k \geq 1$ and let j be any link, $j \neq r$, in group t, $0 \leq t \leq K$. In order to prove that the system is in a Nash equilibrium, we must prove only that $c_i^j \geq c_i^r$. Observe that $c_i^r = k$ and $c_i^j = \ell_j + \frac{w_i}{s_j} = t + 2^{k-t}$. As $t + 2^{k-t} \geq k$ for any non-negative t and r, none of the tasks allocated to r has an incentive to migrate to another link. Therefore, by Definition 1.1, the system is in a Nash equilibrium. □

3.2 Lower Bound for Mixed Strategies. Clearly, since the strategy $\mathcal{S}$ is pure, we have $c = C$. Now, our aim is to slightly modify the allocation of tasks to obtain a mixed strategy $\mathcal{S}'$ for which $C = c \cdot \Theta\left(\frac{\log m}{\log((\log m)/c)}\right)$.

We focus our attention on group K. Let L denote the set of links in this group. L contains $\sqrt{m}$ links. Each of these links has speed 2^K, and to each link we have assigned exactly K tasks of size 2^K each. Let T denote the set of these tasks. The cardinality of this set is $\sqrt{m} \cdot K$. Now, we change the pure strategy $\mathcal{S}$ into a mixed strategy $\mathcal{S}'$ by setting $p_i^j = \frac{1}{\sqrt{m}}$ for every $i \in T, j \in L$. We observe the following properties for our new mixed strategy $\mathcal{S}'$.

LEMMA 3.2. *Strategy $\mathcal{S}'$ satisfies the following properties:*

1. *the maximum load is* $c = K$.
2. *the social optimum is* $1 \leq \mathrm{opt} \leq 2$,
3. *the system is in Nash equilibrium, and*
4. *the social cost is* $C = \Omega\left(\frac{\log m}{\log((\log m)/K)}\right)$.

Proof.

1. The maximum load c is the same as for strategy $\mathcal{S}$.
2. The value of opt is unaffected by the modification of the probabilities.
3. We have to check that the tasks in T do not have smaller expected costs on other links than on the links in L. Observe that the expected cost of these tasks on L slightly increased from K to $K + 1 - \frac{1}{\sqrt{m}} \leq K + 1$. However, for every link $j \notin L$ in group $t < K$ and any

$i \in T$, we have $c_i^j = \ell_j + \frac{w_i}{s_j} = t + 2^{K-t} \geq K+1$, where the last inequality holds for any two integers t and K. Consequently, the system is in a Nash equilibrium.

4. To observe this property, we notice that the allocation of the tasks in T to the links in L corresponds to the allocation problem of throwing $\sqrt{m} \cdot K$ balls uniformly at random into $\sqrt{m}$ bins (see, e.g., [6]). In expectation, it is known that the expected maximum occupancy in this allocation problem is $\Theta\left(K + \frac{\log m}{\log((\log m)/K)}\right)$, which is $\Theta\left(\frac{\log m}{\log((\log m)/K)}\right)$ in our case. Since the sizes of the tasks in T correspond to the speeds of the links in L, this bound on the maximum occupancy directly implies a lower bound on the social cost. □

Thus, by Lemma 3.2, for every positive integer $K \leq \Gamma^{(-1)}(\sqrt{m}/e) - 1 = \frac{\log m}{2 \log\log m}(1 + o(1))$, there exists a set of m links and a Nash equilibrium with $\log(s_1/s_m) = K$ (unless $K = 1$, in which case $s_1 = s_m$), $1 \leq \mathtt{opt} \leq 2$, $c = K$, and

$$C = \Omega\left(\frac{\log m}{\log\left(\frac{\mathtt{opt} \cdot \log m}{c}\right)}\right) .$$

Moreover, we can easily extend this construction to hold for arbitrary positive values of $\mathtt{opt}$. This completes the proof of Theorem 1.3. □

References

[1] P. G. Bradford. On worst-case equilibria for the Internet. Manuscript, 2001.

[2] G. Gonnet. Expected length of the longest probe sequence in hash code searching. *Journal of the Association for Computing Machinery*, 28(2):289–304, 1981.

[3] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In C. Meinel and S. Tison, editors, *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, Trier, Germany, March 4–6, 1999. Springer-Verlag, Berlin.

[4] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Personal communication, 2001.

[5] M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 510–519, Hersonissos, Crete, Greece, July 6–8, 2001. ACM Press, New York, NY.

[6] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.

[7] C. H. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, Hersonissos, Crete, Greece, July 6–8, 2001. ACM Press, New York, NY.

[8] C. H. Papadimitriou and M. Yannakakis. On complexity as bounded rationality. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing (STOC)*, pages 726–733, Montreal, Quebec, Canada, May 23–25, 1994. ACM Press, New York, NY.

[9] T. Roughgarden and É. Tardos. How bad is selfish routing? In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 93–102, Redondo Beach, CA, November 12–14, 2000. IEEE Computer Society Press, Los Alamos, CA.