Approximate Equilibria and Ball Fusi

*

ELIAS KOUTSOUPIAS
University of California at Los Angeles, USA & University of Athens,

MARIOS MAVRONICOLAS
University of Cyprus, Cyprus

PAUL SPIRAKIS
University of Patras, Greece & Computer Technology Institute, Gre

Abstract

We consider *selfish routing* over a network consisting of m parallel *lin* We assume a collection of n *users,* each employing a *mixed strategy,* which a probability distribution over links, to control the routing of its own assign *traffic*. The *expected latency* through a link is taken to be the expectation the traffic assigned to the link; all users sharing a link incur the same *expect latency cost,* taken to be the link's expected latency. In an *approximate eq librium,* each user selfishly routes its traffic so that the expected latency any link is at most a constant multiple of the *optimum* (i.e., the *least possib* maximum latency had global regulation been available.

We introduce and study the *approximate coordination ratio,* defined the ratio of the maximum (over all links) expected latency in the *worst* po sible approximate equlibrium, over the *least* possible maximum latency h global regulation been available, as a measure of the cost of lack of co dination among the users. The load balancing aspect of our selfish routi problem immediately implies a lower bound of $\Omega\left(\frac{\ln m}{\ln \ln m}\right)$ on approxima coordination ratio. The main result of our work provides a correspondi upper bound to establish that this lower bound is *tight*.

To show the upper bound, we introduce and analyze a novel variant the classical *balls and bins* problem, in which balls with arbitrary *weights* a placed into bins according to arbitrary probability distributions; this varia may be of independent interest. A central milestone to our analysis of t variant of the balls and bios problem we consider is a new probabilistic to that we introduce here and call *ball fusion*; this tool is used to reduce t variant of the problem where balls bear weights to the classical version (w no weights).

*This work has been partially supported by the IST Program of the European Union und numbers IST-1999-14186 (ALCOM-FT) and IST-2001-33116 (FLAGS), by the Joint Progr entific and Technological Collaboration between Greece and Cyprus (research project Effic Performance of Distributed Systems), and by funds for the promotion of research at Un Cyprus.

Keywords

Networking, Routing, Game Theory, Nash Equilibrium, Approcimate Equilibri
Fusion

1 Introduction

1.1 Motivation-Framework

We consider a *selfish routing problem* in communication networks, whe
from a *source* to a *destination* are to be established by a collection of *non-c*
entities called *users*. However, different users may wish to optimize co
different objectives of performance and demand. Such networks are he
called *noncooperative* (cf. [22]). Such noncooperative and antagonistic
apply to various modern networking environments, where a *single* perf
objective, to be achieved via cooperation among the users, is no longe
propriate assumption. The objective of our work is to study the inher
due to the lack of a central authority to monitor and regulate network c
according to global objectives.

A natural framework in which to study such multiobjective proble
selfish objectives is (noncooperative) *Game Theory* [23, 24]. Several no
equilibria for noncooperative games have been defined and studied in the
literature of Game Theory; the most famous of them is Nash equilibriu
nally defined in John Nash's seminal paper [19]. Roughly speaking, the s
chosen by the players in a game constitute a *Nash equilibrium* if no playe
better by unilaterally adopting some other strategy.

Several variants and refinements of Nash equilibria have been studi
literature of Game Theory in order to best model the appropriate solut
cept for specific games (or classes of them). Examples include the *Sta*
equilibrium [28], where a distinguished *leader* among players in the gar
a powerful position, while the rest of the players, called *followers,* act ra
a *perfect equilibrium* [18, 27] is obtained as a result of a limiting proce
converges to suitable *mixed strategies* of the players; finally, a *saddle-po*
librium [20] is specially tailored for two-person, zero-sum finite game
work, we introduce and analyze a very general, new class of equilibria, v
call *approximate equilibria*; the definition is tailored to the specific selfish
game we consider in this paper.

Unlike all previously known classes of equilibria, our approximate e
satisfy a very weak law: namely, that the expected latency induced on
be at most a constant times the *optimal* (i.e., the least possible) maximun
that could be achieved by a global algorithm. This law represents a very
requirement to impose on any reasonable solution to the (specific) selfish
problem we study. We argue that Nash equilibria are a special class of

mate equilibria (for the specific game we consider).

In corespondence to approximate equilibria, we introduce *approxima dination ratio,* as a measure of performance loss due to lack of coordina measure generalizes the *coordination ratio* measure introduced in corresp to Nash equilibria, in a very recent paper by Koutsoupias and Papadimitr Roughly speaking, the approximate coordination ratio is the ratio betwee *cial cost* (specifically, the *maximum expected latency* in the setting we c in the *worst* possible approximate equilibrium, over the *social optimum,* the best "offline" global cost (specifically, the *least possible* maximum links) latency in our setting) had all information been available to a cer work authority that regulates traffic. (The definition of approximate coor ratio trades Nash equilibria in the definition of coordination ratio by Kou and Papadimitriou [13] for approximate equilibria.)

In this work, we follow Koutsoupias and Papadimitriou [13] (and its up work [15]) to restrict attention on the simplest case of a network cons m parallel *links*. Each of n users fixes a *mixed strategy,* which is a probab tribution over links. We assume that the *latency* over each link is equal to traffic assigned to the link. (This is known as the model of *uniform spee* Thus, the problem for this simple network is a selfish *resource allocati* lem.

1.2 Contribution

We view our selfish routing problem as an instance of the random exper independently placing n balls into m bins at random. This is the classi *and bins* problem (cf. [10]), which has been studied extensively both in its form and with algorithmic extensions; (see, for example, the recent su Mitzenmacher *et al.* [16]). A central result in the theory of the balls problem is that if one places *uniformly* and *independently* m *identical* b m bins, the expected maximum number of balls in a bin is $\Theta\left(\frac{\log n}{\log \log}\right.$ benefit from the well developed theory of the balls and bins problem, we r users and links as balls and bins, respectively. However, since user traffic identical in the problem we study, balls are not identical either; instead, th with *weights* representing the users' traffics, while the probabilities use random experiment are now arbitrary (in order to account for any set o strategies for the users). The only property we retain in the revised versio balls and bins problem we consider is that balls are still placed independe the bins. Thus, in total, we analyze a variant of the balls and bins proble balls with arbitrary weights are placed into bins according to arbitrary pr distributions; the bounds shown by our analysis immediately translate to on approximate coordination ratio.

Our analysis consists of two major steps. In the first step, called *ba*

(Section 3.1), we reduce the case of arbitrary weights to the case where all
are within a factor of 2 to each other. To do so, we "fuse" the two (c
smallest balls together to form a new larger ball with weight equal to
of the two. Moreover, we assign new probabilities to the resulting ball
that the expected weight assigned to each link is preserved. We proceed
(Lemma 1) that, roughly speaking, the social cost of the resulting game i
than the social cost of the original game (before applying ball fusion). W
this fusion procedure until all remaining balls have weights withing a fa
to each other.

In the second step of our analysis (Section 3.2), we consider the spe
where all balls have *identical* weights; the social cost for this case is no m
a multiplicative factor of 2 times the social cost for the case where all b
weights within a factor of 2, resulting at the end of the previous step. We
ply standard techniques for estimating tails and Chernoff bounds [5] to s
(roughly speaking), for this case, the social cost is at most $O(\ln m / \ln \ln$
the maximum expected number of balls that is placed into any link (Le
Putting together the two steps yields that, for balls with arbitrary weight
any arbitrary probability distribution, the expected maximum is at most O
times the maximum expected number of balls that is placed into any li
orem 1). The remaining link needed for completing the proof of our ma
(Theorem 2), namely that the approximate coordination ratio is $O(\ln m /$
is just the property we required in the definition of approximate equilibria
maximum expected latency be at most a constant multiple of the optimu

1.3 Related Work and Comparison

Our work joins the recent studies by Koutsoupias and Papadimitriou [
by Mavronicolas and Spirakis [15], of performance degradation caused b
of regulation in a congested network. The strongest known results fo
mance degradation bounds in Nash equilibria settings were derived by
and Voecking [6]. Their results cover also the case of arbitrary link spe
though in the Nash equilibria setting. Our technique is completely differ
their techniques in being purely combinatorial, it does not require any
structure or property that are specific to Nash equilibria, and it yields a
general result. A recent work by Berebrink *et al.* [4] studies also a pro
placing balls with arbitrary weights into bins. They consider, however, d
form allocations and study protocols that try to improve the load bala
many rounds of communication.

We note that similar results, and results for links with speeds as w
recently derived in a work by Czumaj and Vöcking [6] hat followed the
appearance of our work [8]. However, in their work, they focus their atter
techniques to a special case of approximate equilibria, namely the wel
Nash equilibria. In contrast, our technique, that we code by the name *ba*

completely bypasses the use of any properties that are specific to Nash eq and thus derives a far more general result; n addition, it is much simpler a intuitive than the techniques used in [6], and also it is of independent int

Game-theoretic models, concepts and techniques have been employed in the context of various networking problems such as *flow control* [1, 1 *ing* [12, 13, 22], *bandwidth allocation* [14], *Web access* [25], *multicasting congestion control* [9]. For algorithmic applications of game theory to t discipline of computer science, see [26] for a gathering of recent success

1.4 Road Map

The rest of this paper is organized as follows. Section 2 provides definiti it summarizes some background material and preliminary results. Our ma is shown in Section 3. We conclude, in Section 4, with a discussion of o and some open problems.

2 Definitions, Background and Preliminaries

Throughout, denote for any integer $m \geq 2$, $[m] = \{1, \ldots, m\}$. For an ev a sample space, denote $\mathbf{Pr}(E)$ the probability of event E occurring. For a variable X, denote $\mathcal{E}(X)$ the *expectation* of X. Denote e the base of th logarithm.

As in [13] and [15], we consider a *network* consisting of a set of m *links* $1, 2, \ldots, m$ from a *source* node to a *destination* node. Each of n *users* $1, 2, \ldots, n$, or *users* for short, wishes to route a particular amount along a (non-fixed) link from source to destination; denote w_i the *traffic* $i \in [n]$. 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* f is a probability distribution on the set of links. We often use subscripts and superscripts for links. A *set of pure strategies,* one per user, is rep by an n-tuple $\langle \ell_1, \ell_2, \ldots, \ell_n \rangle \in [m]^n$; a *set of mixed strategies,* one pe represented by an $m \times n$ *probability matrix* $\mathbf{P}$ of mn probabilities p_i^ℓ, $i \in$ $\ell \in [m]$, where p_i^ℓ is the probability that user i selects link ℓ. Clearly, $\mathbf{P}$ For a probability matrix $\mathbf{P}$, define *indicator variables* $I_i^\ell \in \{0, 1\}$, $i \in$ $\ell \in [m]$, such that $I_i^\ell = 1$ if and only if $p_i^\ell > 0$. The setting is remin classical *random allocation* problems, where n *balls* (with *weights*) ar into m *bins* at random (see, e.g., [10]). Thus, we will interchangeably terms users and balls, and links and bins, respectively, in our discussion.

We assume that the *latency* for traffic w through link $\ell \in [m]$ equals corresponds to the model of *uniform speeds or capacities,* introduced in studied in [13, 15]. In the model of *non-uniform capacities,* the latency f w through link ℓ equals w/c_ℓ, where $c_\ell > 0$ is the *capacity* of link ℓ.)

For any set of pure strategies $\langle \ell_1, \ell_2, \ldots, \ell_n \rangle$, the *latency cost for use* denoted λ_i, is $(\sum_{k:\ell_k=\ell_i} w_k)$; that is, the latency cost for user i is the of the link it chooses. For any set of mixed strategies $\mathbf{P}$, denote θ_ℓ the *traffic* on link ℓ; clearly, $\theta_\ell = \sum_{i=1}^{n} p_i^\ell w_i$.

We are interested in special classes of mixed strategies called *app* equilibria. In such equilibria the probability matrix $\mathbf{P}$ has the property any link , its expected latency is at most a constant times the *optimum* lat latter is the latency that could be achieved if a central scheduler was sc all the loads to links.

Associated with a traffic vector $\mathbf{w}$ and a set of mixed strategies $\mathbf{P}$ is t *cost* [13, Section 2], denoted $\mathsf{C}(\mathbf{w}, \mathbf{P})$, which is the expectation, over all choices of the users, of the maximum (over all links), of the latency through a link; thus,

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

On the other hand, the *social optimum* [13, Section 2] associated with a tra tor $\mathbf{w}$, denoted OPT, is the *minimum possible* maximum (over all links) l traffic through a link; thus, $\mathsf{OPT} = \min_{\langle \ell_1, \ell_2, \ldots, \ell_n \rangle \in [m]^n} \max_{\ell \in [m]} \sum_{k:\ell}$

The *approximate coordination ratio* is the *maximum* value, over a vectors $\mathbf{w}$ and approximate equilibria $\mathbf{P}$, of the ratio $\mathsf{C}(\mathbf{w}, \mathbf{P})/\mathsf{OPT}$.

The following result from [13] relates the expected traffic on each the social cost. This result shows that the Nash equilibria are a special ca equilibria considered here. It is not difficult to construct approximate equ that are *not* Nash equilibria.

Proposition 1 (Koutsoupias and Papadimitriou [13]) *Take any a Nash rium* $\mathbf{P}$. *Then, for any link* $\ell \in [m]$,

$$\theta_\ell \quad \leq \quad \left(2 - \frac{1}{m}\right) \cdot \mathsf{OPT}\,.$$

3 Main Result

In this section we derive general bounds of the expected maximum when randomly n balls independently into m bins. In the most-studied occupan lem, identical balls are placed uniformly into the bins. Although many va the problem have been extensively studied (see [16] for an excellent sur general case —in which the balls have arbitrary weights and the probabi arbitrary— has not been considered before. It is clear that we cannot s about the expected maximum $\mathsf{C}(\mathbf{w}, \mathbf{P})$ in the general case. For examp all balls fall with probability one into the first bin or when the weight of

dwarfs the rest, then $\mathsf{C}(\mathbf{w},\mathbf{P})$ can be as high as $\sum_i w_i$. But we will sh
if we exclude these two pathological cases, that is, when there are no ve
balls and the expected weights in the bins are balanced, then, up to a
factor, the worst case happens when all balls are identical. We show:

Theorem 1 *For every traffic vector* $\mathbf{w}$ *and probability matrix* $\mathbf{P}$*,*

$$\mathsf{C}(\mathbf{w},\mathbf{P}) \quad \in \quad \max\left\{\max_{i\in[n]} w_i, \max_{\ell\in[m]} \theta_\ell\right\} \cdot O\left(\frac{\ln m}{\ln\ln m}\right).$$

We start with an informal outline of our proof. Suppose we replace t
with their sum and assign weights to the sum so that the expected traffic
bin remains the same. What will happen to $\mathsf{C}(\mathbf{w},\mathbf{P})$? Naturally, given the
correlation between the two balls, we expect $\mathsf{C}(\mathbf{w},\mathbf{P})$ to increase or re
same. We will show that this is indeed the case. With this in mind, w
the process of replacing the two smaller balls with their sum until all ba
weights within a factor of 2 to each other. Thus, we reduce the problem to
of identical balls (within a factor of 2) with non-uniform probabilities.
apply standard techniques of tail estimation to prove the claimed bound.
continue with the details of the formal proof.

3.1 Ball Fusion

Given the $n \times m$ probability matrix $\mathbf{P}$, define the $(n-1)\times m$ probabilit
$\widehat{\mathbf{P}}$ such that $\widehat{p}_i^\ell = p_i^\ell$ if $i < n-1$, and

$$\widehat{p}_{n-1}^\ell \quad = \quad \frac{p_{n-1}^\ell w_{n-1} + p_n^\ell w_n}{w_{n-1}+w_n}.$$

Thus, all rows of $\widehat{\mathbf{P}}$ are identical to the corresponding ones of $\mathbf{P}$, excep
last one, each entry of which is a linear combination of the two corres
entries in the last two rows of $\mathbf{P}$. Similarly, define $\widehat{\mathbf{w}}$ to be identical to
first $n-2$ entries while its last entry is the sum of the last two entries o
$\widehat{w}_i = w_i$ for $i = 1,\ldots,n-2$ and $\widehat{w}_{n-1} = w_{n-1} + w_n$. The following l
crucial:

Lemma 1 $\mathsf{C}(\mathbf{w},\mathbf{P}) \le C(\widehat{\mathbf{w}},\widehat{\mathbf{P}})$ *for all* $\mathbf{w}$ *and* $\mathbf{P}$.

To prove the lemma, consider that we place the balls into the bins sequ
The processes for $\mathsf{C}(\mathbf{w},\mathbf{P})$ and $\mathsf{C}(\widehat{\mathbf{w}},\widehat{\mathbf{P}})$ are identical for the first $n-$
This suggests that we should study the case when the bins are not initiall
Thus we will prove a (slight) generalization of the lemma in which the b
some initial weight. Let $\mathbf{L}(L_1,\ldots,L_m)$ be the initial weights of the bin
$\mathsf{C}(\mathbf{w},\mathbf{P},\mathbf{L})$ be the expected maximum when we place balls with weigh
probabilities $\mathbf{P}$. It is straightforward how to extend (1) when the links ha
loads $\mathbf{L}$.

Lemma 2 $\mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L}) \leq \mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L})$ *for all* $\mathbf{w}$ *and* $\mathbf{P}$.

Proof. We first show the lemma for the case $n = 2$ and the general case immediately from it.

We want to show that for every $\mathbf{w}$, $\mathbf{P}$, and $\mathbf{L}$, $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L}) - \mathsf{C}(\mathbf{w},$ nonnegative. If we fix $\mathbf{L}$, $\mathbf{w}$, and the probabilities p_1^j, $j = 1, \ldots, m$, the ex $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L}) - \mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L})$ is a linear function of the $\widehat{p}_2^j$s, $j = 1, \ldots, m$, a seen from the definition of social cost. Thus, subject to the condition $\sum_j$ the expression $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L}) - \mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L})$ is minimized when one of th 1 and the rest are 0. Here is an alternative way to see this: when we f and the probabilities of the first ball, the strategy that minimizes $\mathsf{C}(\widehat{\mathbf{w}},$ $\mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L})$ is to place the second ball in the bin with the expected n value. Similarly for the $\widehat{p}_1^j$s. In other words, for every $\mathbf{w}$ and $\mathbf{L}$, $\mathsf{C}(\widehat{\mathbf{w}},$ $\mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L})$ is minimized by "pure strategies", that is, according to $\mathbf{P}$, ea placed into a fixed bin.

But for pure strategies, it can be directly verified that $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L}) - \mathsf{C}($ is nonnegative as follows: Assume that according to $\mathbf{P}$, the ball w_{n-1} i in bin j_{n-1} and the ball w_n is placed in bin j_n. If $j_{n-1} = j_n$ ther $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L}) = \mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L})$. Otherwise, $\mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L}) = \max(\max_j L_j,$ $w_{n-1}, L_{j_n} + w_n)$. On the other hand, according to $\widehat{\mathbf{P}}$, a ball with weight w_n is placed into bin j_{n-1} with probability $w_{n-1}/(w_{n-1}+w_n)$ and to bi the remaining probability. Then $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L}) = \frac{w_{n-1}}{w_{n-1}+w_n} \max(\max_j L_j,$ $w_{n-1} + w_n) + \frac{w_n}{w_{n-1}+w_n} \max(\max_j L_j, L_{j_n} + w_{n-1} + w_n)$. From th be directly verified that $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L}) - \mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L})$ is nonnegative. A sim to see this is to consider that $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L})$ is the *expected maximum* of th iment with probabilites $\widehat{\mathbf{P}}$ but $\mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L})$ is the *maximum expected* of t experiment.

This settles the lemma when $n = 2$. The case $n > 2$, f mediately if we consider placing the balls in order. To formaliz need to define $\mathbf{L}'$ the vector of weights after placing the first $n -$ $w_1, \ldots, w_{n-2}$. The probability distribution for $\mathbf{L}'$ is the same for both $\mathsf{C}($ and $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L})$. Let also $\mathbf{w}_{n-1}$ and $\mathbf{P}_{n-1}$ be the weights and prob of the last 2 balls. Then $\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}, \mathbf{L}) = \sum_{\mathbf{L}'} \mathbf{Pr}(\mathbf{L}')\mathsf{C}(\widehat{\mathbf{w}}_{n-1}, \widehat{\mathbf{P}}_{n-1}$ $\sum_{\mathbf{L}'} \mathbf{Pr}(\mathbf{L}')\mathsf{C}(\mathbf{w}_{n-1}, \mathbf{P}_{n-1}, \mathbf{L}') = \mathsf{C}(\mathbf{w}, \mathbf{P}, \mathbf{L})$ where the inequality cor the special case of 2 balls.

Remark: n the above proof we used essentially only the fact that th holds for pure strategies. Thus for any cost function C for which the lem for pure strategies, it also holds for mixed strategies. In particular, *the len holds for the capacitated case* (in which the weight of bin $\ell \in [m]$ is di its capacity c_ℓ).

Using Lemma 1, we can transform a set of non-identical balls into
almost identical ones as follows. Let $h = 2\max_i w_i$. Keep replacing
lighter balls with their sum while the resulting ball is no heavier than h. I
that the resulting weights $\widehat{w}_1, \ldots, \widehat{w}_k$ are all within a factor of 2 and in p
in the interval $[h/2, h]$—if a ball has weight less than $h/2$ it can be fused
with the (original) ball of maximum weight $h/2$.

If we now keep the probabilities of the balls the same and incre
weight to h the value of $\mathsf{C}(\mathbf{w}, \mathbf{P})$ at most doubles. More precisely, $\mathsf{C}(\mathbf{v}$
$\mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}) \leq h \cdot \mathsf{C}(\mathbf{1}, \widehat{\mathbf{P}})$, where $\mathbf{1}$ is the n'-dimensional vector of all 1's,
$n' \leq n$ (n' is the number of resulting balls).

3.2 Identical Balls

We now consider the case where all n balls are identical and have w
Roughly speaking, we will establish that in this case the social cost is
than $O(\ln m / \ln \ln m)$ times the maximum (over all links) of the expect
ber of balls in any particular bin. Note that this may not hold when the nu
balls is small; for example, if there is only one ball which is placed unifor
the m bins, then, the social cost (expectation of the maximum number
in a bin) is 1, while the maximum expected number of balls in any parti
is $1/m$, which is substantially smaller. Thus, to fix this problem, we wi
the social cost (expected maximum) by either the maximum expected nu
balls in a link or by the weight of each ball. Formally, we show:

Lemma 3 *For any arbitrary probability matrix* $\mathbf{P}$*,*

$$\mathsf{C}(\mathbf{1}, \mathbf{P}) \leq \left(\frac{2e \ln m}{\ln \ln m} + 1\right) \cdot \max\left\{\max_{\ell \in [m]} \theta_\ell, 1\right\}.$$

Proof. The proof is not substantially different from the well-studied v
uniform probabilities. Roughly speaking, we focus on a fixed bin $\ell \in [m]$
show that the probability of the weight exceeding $\max\{1, \theta_\ell\}$ by a facto
than $\Theta(\ln m / \ln \ln m)$ is small (specifically, $1/m^2$). The claim then fo
averaging according to the definition of social cost.

For each pair of a user $i \in [n]$ and link $\ell \in [m]$, define the (bina
dom variable X_i^ℓ by $\mathbf{Pr}\left(X_i^\ell = 1\right) = p_i^\ell$ and $\mathbf{Pr}\left(X_i^\ell = 0\right) = 1 -$
$X^\ell = \sum_{i=1}^n X_i^\ell$ and $\theta_\ell = \mathcal{E}(X^\ell)$. For any parameter $\alpha > 0$, we upp
the probability $\mathbf{Pr}(X^\ell \geq \alpha \max\{\theta_\ell, 1\})$, using the fact that X^ℓ is
of independent Bernoulli variables and apply Chernoff bounds to ob
$\mathbf{Pr}\left(X^\ell \geq \alpha\theta_\ell\right) \leq \left(\frac{e^{\alpha-1}}{\alpha^\alpha}\right)^{\theta_\ell} \leq \left(\frac{e}{\alpha}\right)^{\alpha\theta_\ell}$. We now choose appropriatel
rameter α so that this probability does not exceed $1/m^2$. We proceed
analysis.

- Assume first that $\theta_\ell \geq 1$. Then, choose $\alpha = \frac{2e \ln m}{\ln \ln m}$ $\mathbf{Pr}\left(X^\ell \geq \alpha \max\{\theta_\ell, 1\}\right) = \mathbf{Pr}\left(X^\ell \geq \alpha\theta_\ell\right) \leq \left(\frac{e}{\alpha}\right)^{\alpha\theta_\ell} \leq \left(\frac{e}{\alpha}\right)^{\alpha}$
- Assume now that $\theta_\ell < 1$. Then, choose $\alpha = \frac{1}{\theta_\ell}\frac{2e \ln m}{\ln \ln m}$ $\mathbf{Pr}\left(X^\ell \geq \alpha \max\{\theta_\ell, 1\}\right) = \mathbf{Pr}\left(X^\ell \geq \alpha\right) \leq \mathbf{Pr}\left(X^\ell \geq \alpha\right)$ $\left(\frac{e}{\alpha}\right)^{\alpha\theta_\ell} \leq \left(\frac{e}{\alpha\theta_\ell}\right)^{\alpha\theta_\ell} \leq \frac{1}{m^2}$.

So, in both cases, $\mathbf{Pr}\left(X^\ell \geq \frac{2e \ln m}{\ln \ln m} \max\{\theta_\ell, 1\}\right) \leq \frac{1}{m^2}$. If we cons all m bins, the probability that there exists a bin $\ell \in [m]$ such that X^ℓ $\frac{2e \ln m}{\ln \ln m} \max\{\theta_\ell, 1\}$ is at most $1/m$. We can now bound the expected m $\mathsf{C}(\mathbf{1}, \mathbf{P})$: with probability at most $1 - 1/m$, the expected maximum is $\frac{2e \ln m}{\ln \ln m} \max\{\theta_\ell, 1\}$, and with the remaining probability, the expected m is at most n. Taking into account that $n/m \leq \max_\ell \theta_\ell$, we get that

$$\mathsf{C}(\mathbf{1}, \mathbf{P}) \leq \left(\frac{e \ln m}{\ln \ln m} + 1\right) \cdot \max\left\{\max_{\ell \in [m]} \theta_\ell, 1\right\},$$

as needed.

3.3 Proof of Main Theorem

The proof of Theorem 1 is now complete: we showed

$$\mathsf{C}(\mathbf{w}, \mathbf{P}) \leq \mathsf{C}(\widehat{\mathbf{w}}, \widehat{\mathbf{P}}) \leq \mathsf{C}((2 \max_i w_i) \cdot \mathbf{1}, \widehat{P}) = O(\frac{\ln m}{\ln \ln m}) \max(\max_i w$$

where $\widehat{\mathbf{w}}$ and $\widehat{\mathbf{P}}$ are the weights and probabilities of the balls resulting fro the original balls.

We can now combine Theorem 1 with the definition of approximate e to derive the desired upper bound on the approximate coordination ratio definition on each link $\ell \in [m]$ is at most $k\mathsf{OPT}$ where OPT is the social c and k is the constant involved in the definition of approximate equilibria

Furthermore, since OPT is bounded below by the maximum weight,

$$\max(\max_i w_i, \max_{\ell \in [m]} \theta_\ell) \leq \max(\mathsf{OPT}, k\mathsf{OPT})$$

from which we derive our main result $\mathsf{C}(\mathbf{w}, \mathbf{P}) = O(\frac{\ln m}{\ln \ln m})\, \mathsf{OPT}$. Thu lows:

Theorem 2 *The approximate coordination ratio of selfish routing on m links and any number of agents is* $O(\ln m / \ln \ln m)$.

4 Discussion and Directions for Further Resea

We have given a method (i.e., ball fusion) which leads to tight bound social cost of any arbitrary set of mixed strategies of n agents. We hav also that only a very weak property of such strategies (namely, the appr equilibria property) is needed in addition , in order to get tight bound approximate coordination ratio of such selfish allocations.

An interesting research direction is to examine cases where the pro approximate equilibria is relaxed even more so that the expected latency is just a function of OPT. Our method gives immediately some bounds case. Finally, it would be interesting to study the relative efficiency betwe equilibria and approximate equilibria. (Although this work and the work maj and Vöcking [6] establish identical (up to constant factors) upper bo their individual efficiencies, as captured by approximate coordination r coordination ratio, respectively, it is not yet known how the two compare other.)

References

[1] E. Altman. Flow Control Using the Theory of Zero-Sum Markov *IEEE Transactions on Automatic Control*, Vol. 39, No. 4, pp. 814–8 1994.

[2] Y. Azar, A. Z. Broder, A. R. Karlin and E. Upfal. Balanced Allo *SIAM Journal on Computing*, Vol. 29, No. 1, pp. 180–200, Septemb

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

[4] P. Berenbrink, F. Meyer auf der Heide and K. Schröder. Allocating V Jobs in Parallel. *Proceedings of the 9th Annual ACM Symposium on Algorithms and Architectures*, pp. 302–310, June 1997.

[5] H. Chernoff. A Measure of Asymptotic Efficiency for Tests of a H sis Based on the Sum of Observations. *Annals of Mathematical S* Vol. 23, No. 4, pp. 493–509, December 1952.

[6] A. Czumaj and B. Vöcking, Tight Bounds for Worst-case Equilibr *cedings of the 13th Annual ACM-SIAM Symposium on Discrete Alg* pp. 413–420, January 2002.

[7] J. Feigenbaum, C. Papadimitriou and S. Shenker. Sharing the Cost ticast Transactions. *Proceedings of the 32nd Annual ACM Sympo Theory of Computing*, pp. 218–227, May 2000.

[8] E. Koutsoupias, M. Mavronicolas and P. Spirakis. A Tight Bound
dination Ratio. Technical Report 010029, Department of Computer
University of California at Los Angeles, 2001.

[9] R. Karp, E. Koutsoupias, C. H. Papadimitriou and S. Shenker. O
tion Problems in Congestion Control. *Proceedings of the 41st Annu
Symposium on Foundations of Computer Science*, pp. 66–74, Octob

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

[11] Y. Korilis and A. Lazar. On the Existence of Equilibria in Noncoo
Optimal Flow Control. *Journal of the ACM*, Vol. 42, No. 3, pp. 5
May 1995.

[12] Y. Korilis, A. Lazar and A. Orda. Architecting Noncooperative N
IEEE Journal on Selected Areas in Communications, Vol. 13, N
1241–1251, September 1995.

[13] E. Koutsoupias and C. H. Papadimitriou. Worst-case Equilibria. *
*ings of the 16th Annual Symposium on Theoretical Aspects of Comp
ence*, G. Meinel and S. Tison eds., pp. 404–413, Vol. 1563, Lecture
Computer Science, Springer-Verlag, March 1999.

[14] A. A. Lazar, A. Orda and D. E. Pendarakis. Virtual Path Bandwid
cation in Multi-User Networks. *IEEE/ACM Transactions on Net*
Vol. 5, No. 6, pp. 861–871, December 1997.

[15] M. Mavronicolas and P. Spirakis. The Price of Selfish Routing. *Pro
of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 5
July 2001.

[16] M. Mitzenmacher, A. W. Richa, and R. Sitaraman. The Power of T
dom Choices: A Survey of Techniques and Results. *Handbook of R
ized Computing*, Kluwer Academic Publishers, 2001.

[17] R. Motwani and P. Raghavan. *Randomized Algorithms*, Cambridge
sity Press, 1995.

[18] R. B. Myerson. Refinement of the Nash Equilibrium Concept. *Jo
Game Theory*, Vol. 7, pp. 73–80, 1978.

[19] J. F. Nash. Non-cooperative Games. *Annals of Mathematics*, Vol. 5
pp. 286–295, 1951.

[20] J. von Neumann and O. Morgenstern. *Theory of Games and Econo
havior*. (second edition), Princeton University Press, Princeton, NJ

[21] N. Nisan and A. Ronen. Algorithmic Mechanism Design. *Proce*
the 31st Annual ACM Symposium on Theory of Computing, pp. 1
May 1999.

[22] A. Orda, R. Rom and N. Shimkin. Competitive Routing in Multius
munication Networks. *IEEE/ACM Transactions on Networking*
No. 5, pp. 510–521, October 1993.

[23] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*, T
Press, 1994.

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

[25] C. H. Papadimitriou and M. Yannakakis. On the Approximability o
offs and Optimal Access of Web Sources. *Proceedings of the 41s*
IEEE Symposium on Foundations of Computer Science, pp. 86–92,
2000.

[26] Web page `http://www.cs.berkeley.edu/~c`
`tos/games/cs294.html` for the course *CS294: Sem*
Algorithmic Aspects of Game Theory. taught by C. H. Papad
University of California at Berkeley, Spring 2001.

[27] R. Selten. Reexamination of the Perfectness Concept for Equilibriu
in Extensive Games. *International Journal of Game Theory*, Vo
25–55, 1975.

[28] H. von Stackelberg. *The Theory of the Market Economy*. Oxford U
Press, Oxford, England, 1934.

Elias Koutsoupias is affiliated with the Department of Computer Science, Uni
California at Los Angeles, Los Angeles, USA, and with the Department of Inf
University of Athens, Athens, Greece. Email: `elias@cs.ucla.edu`

Marios Mavronicolas is affiliated with the Department of Computer Science, U
of Cyprus, Nicosia, Cyprus. Email: `mavronic@ucy.ac.cy`

Paul Spirakis is affiliated with the Department of Computer Engineering and Inf
University of Patras, Patras, Greece, and with the Computer Technology Institu
Greece. Email: `spirakis@cti.gr`