On the Structure and Complexity of Worst-Case Equilibria*

Simon Fischer and Berthold Vöcking

RWTH Aachen, Computer Science 1,
52056 Aachen, Germany
{fischer, voecking}@cs.rwth-aachen.de

Abstract. We study an intensively studied resource allocation game introduced by Koutsoupias and Papadimitriou where n weighted jobs are allocated to m identical machines. It was conjectured by Gairing et al. that the fully mixed Nash equilibrium is the worst Nash equilibrium for this game w. r. t. the expected maximum load over all machines. The known algorithms for approximating the so-called "price of anarchy" rely on this conjecture. We present a counter-example to the conjecture showing that fully mixed equilibria cannot be used to approximate the price of anarchy within reasonable factors. In addition, we present an algorithm that constructs so-called *concentrated equilibria* that approximate the worst-case Nash equilibrium within constant factors.

1 Introduction

A central problem arising in the management of large-scale communication networks like the Internet is that of routing traffic through the network. Due to the large size of these networks, however, it is often impossible to employ a centralized traffic management. A natural assumption in the absence of central regulation is to assume that network users behave selfishly and aim at optimizing their own individual welfare. To understand the mechanisms in such non-cooperative network systems, it is of great importance to investigate the selfish behavior of users and their influence on the performance of the entire network.

In this paper, we investigate the price of selfish behavior under game theoretic assumptions, that is, we assume that each *agent* (i.e., user) is aware of the situation facing all other agents and aims at optimizing its own strategy. In particular, we investigate the structure of the network in a *Nash equilibrium*, i.e., a combination of mixed (randomized) strategies from which no users has an incentive to deviate. It is well known that such equilibria may be inefficient and do not always optimize the overall performance. We address the most basic case of a routing problem, a network consisting of m *identical parallel links* from an origin to a destination. There are n agents, each having an amount of traffic w_i to send from the origin to the destination. Each agent i sends the traffic using a

* Supported in part by the EU within the 6th Framework Programme under contract 001907 (DELIS) and by DFG grant Vo889/1-2.

X. Deng and Y. Ye (Eds.): WINE 2005, LNCS 3828, pp. 151–160, 2005.

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 aims at minimizing its individual cost, i. e., the expected load on the machine it selects.

Koutsoupias and Papadimitriou [10] proposed to investigate the price of uncoordinated individual decisions in terms of the worst-case *coordination ratio*, which is the ratio between the expected social cost in the worst possible Nash equilibrium and in the social optimum. Here, we define social cost as the maximal load of a machine. In [2] and [9] it has been shown that the coordination ratio for a system of n weighted jobs and m identical machines is bounded by $\Theta\left(\frac{\ln m}{\ln \ln m}\right)$. However, for a given instance of the game, i. e., for a given vector of weights, the ratio between worst social cost of a Nash equilibrium and the optimal social cost may be significantly smaller. It is therefore an interesting question how the cost of the worst Nash equilibrium can be computed for a given instance. It has been conjectured that the fully mixed Nash equilibrium in which every agent assigns the same probability to every machine is the worst possible [6, 5, 4]. If this conjecture was true, then computing the worst Nash equilibrium would be a trivial task and its social cost could be approximated arbitrarily well using the fully polynomial randomized approximation scheme (FPRAS) presented in [3].

In this paper, we show that the Fully Mixed Nash Equilibrium Conjecture does not hold. In fact, the ratio between the social cost of the fully mixed Nash equilibrium and the worst Nash equilibrium can be almost as bad as the coordination ratio. We then present a different kind of equilibrium that concentrates the large jobs on a few machines. These *concentrated equilibria* are as bad as the worst Nash equilibrium up to constant factors. They can be computed in linear time and hence we obtain the first constant-factor approximation for worst-case equilibria on identical machines.

1.1 The Game

Koutsoupias and Papadimitriou [10] introduced a resource allocation game in which n jobs of size $w_1, \ldots, w_n \geq 0$ shall be assigned to m identical machines. Each job is managed by a selfish agent. The set of *pure strategies* for task i is $[m] := \{1, \ldots, m\}$. Let $(j_1, \ldots, j_n) \in [m]^n$ be a combination of pure strategies, one for each task. The *load* of link j is defined as

$$\lambda_j \;=\; \sum_{j_k = j} w_k \ .$$

The *cost* for agent i is λ_{j_i}. Every agent aims at minimizing her cost. The *social objective* is to minimize the maximum cost over all agents or, equivalently, the maximum load over all machines.

Agents may also use *mixed strategies*, i. e., probability distributions on the set of pure strategies. Let p_i^j denote the probability that agent i assigns its job to link j. Then

$$\mathbb{E}\left[\lambda_j\right] \;=\; \sum_{i \in [n]} w_i \, p_i^j \ .$$

The social cost of a mixed strategy profile $\mathbf{P} = (p_i^j)$ is defined as

$$SC(\mathbf{P}) \;=\; \mathbb{E}\left[\max_{j\in[m]} \lambda_j\right] \;.$$

The *expected cost of task i on link j* is

$$c_i^j \;=\; w_i + \sum_{k\neq i} w_k\, p_k^j \;=\; \mathbb{E}\left[\lambda_j\right] + (1 - p_i^j)\, w_i \;.$$

A (mixed) strategy profile $\mathbf{P}$ defines 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 $q \in [m]$. A Nash equilibrium is called *fully mixed* if $p_i^j > 0$ for all $i \in [n]$, $j \in [m]$. The game under consideration admits a unique fully mixed Nash equilibrium $\mathbf{F}$ in which each job is assigned with probability $\frac{1}{m}$ to each machine [11].

1.2 The Conjecture

Mavronicolas and Spirakis [11] investigate the social cost of fully mixed Nash equilibria. The motivation for their study is the hope that the techniques for the analysis of fully mixed strategies can be appropriately extended to yield upper bounds on the social cost for general equilibria. This hypothesis is formalized in the following conjecture stated in [6, 5, 4].

Conjecture 1 (FMNE conjecture). The fully mixed Nash equilibrium $\mathbf{F}$ is the worst Nash equilibrium, that is,

$$SC(\mathbf{F}) \;\geq\; SC(\mathbf{P}) \;,$$

for every Nash equilibrium $\mathbf{P}$.

Several attempts have been made to prove the conjecture. For example, it was shown that the conjecture is true for the case $m = 2$ [4] and for the case that $\mathbf{P}$ refers only to pure equilibria [6]. Furthermore, it was shown that the conjecture holds in an approximate sense if $m = n$ [1, 6]. In [3], an FPRAS for the social cost of the fully mixed Nash equilibrium is presented. Further interesting discussions can be found in [6].

The FMNE conjecture seems to be intuitive and appealing since in case of its validity it would allow for an easy identification of the worst-case mixed Nash equilibrium, whereas the worst case pure Nash equilibrium is NP-hard to compute.

1.3 Outline and Contribution

In Section 2 we give a counter-example to the FMNE conjecture that shows that mixed Nash equilibria may have a social cost that is by a factor of $\Omega\left(\frac{\ln m}{\ln\ln m}\right)$ worse than the social cost of the fully mixed Nash equilibrium. This is indeed the worst possible.

In Section 3 we present a simple algorithm that constructs a constant factor approximation of the worst Nash equilibrium in linear time.

2 The Counterexample

We present a counterexample to the FMNE conjecture. More specifically, we show that there is a family of simple instances of the game for which there exists an equilibrium $\mathbf{P}$ with

$$SC(\mathbf{P}) \;=\; \Omega\left(SC(\mathbf{F}) \cdot \frac{\ln m}{\ln\ln m}\right) \; .$$

Let us remark that this is the worst possible ratio as it follows from the analyses in [2, 9] that the social cost of every Nash equilibrium can be at most $\mathcal{O}\left(\frac{\ln m}{\ln\ln m}\right)$ times the optimal social cost.

Theorem 1. *For every m, there exists an instance of the resource allocation game with m machines admitting a Nash equilibrium P with*

$$SC(\mathbf{P}) \;=\; \left(\frac{1}{4} - o(1)\right) \cdot \frac{\ln m}{\ln\ln m} \cdot SC(\mathbf{F}) \; .$$

The instance consist of $n = \mathcal{O}(f(m) \cdot m \ln m)$ jobs whose weights differ at most by a factor $\mathcal{O}(f(m) \cdot \ln m)$, where f denotes an arbitrary function in $\omega(1)$.

Proof. The counterexample uses only two different kinds of jobs: *Large jobs* of weight 1 and *small jobs* of weight $\frac{1}{k}$, $k \in \mathbb{N}$. Let $\ell \leq m$ denote the number of large jobs. The number of small jobs is $k(m-\ell)$. Thus, the total weight is m and the optimal assignment has social cost 1. We show that the fully mixed equilibrium has social cost close to optimal if the parameters k and ℓ are chosen appropriately.

Lemma 1. *If $k = \Omega(f(m) \cdot \ln m)$ and $\ell = \mathcal{O}(\sqrt{n}/f(m))$ then $SC(\mathbf{F}) \leq 2 + o(1)$.*

Proof. Recall that $\mathbf{F}$ assigns each job with probability $\frac{1}{m}$ to each of the machines.

- The assignment of the large jobs corresponds to a balls-and-bins experiment in which $\ell = \mathcal{O}(\sqrt{m}/f(m))$ balls are assigned uniformly at random to m bins. Fact 1 from the Appendix yields that for this experiment the expected number of balls in the fullest bin is $1 + o(1)$. Thus, the expected maximum load due to the large jobs is $1 + o(1)$, too.
- The assignment of the small jobs corresponds to a ball-and-bins experiment in which $k(m-\ell)$ balls are assigned uniformly at random to $m - \ell$ bins for $k = \Omega(f(m) \cdot \ln m)$. Fact 2 shows that for this experiment the expected number of balls in the fullest bin is $(1 + o(1)) \cdot k$. Since each ball corresponds to a job of weight $\frac{1}{k}$, the expected maximum load due to the small jobs is thus $1 + o(1)$ as well.

Combining the upper bounds for the small and the large jobs yieds that the maximum load over all machines is at most $2 + o(1)$ when taking into account all the jobs. □

Next we present a mixed Nash equilibrium whose maximum load is lower-bounded by a function in ℓ.

Lemma 2. *There exists a Nash equilibrium* $\mathbf{P}$ *with* $SC(\mathbf{P}) \geq (1 - o(1)) \cdot \frac{\ln \ell}{\ln \ln \ell}$.

Proof. We construct $\mathbf{P}$ in the following way. The small jobs are assigned using pure strategies. They are distributed evenly among the machines $1, \ldots, m - \ell$ such that each machine receives k small jobs. Hence, their load is fixed to 1. The large jobs are assigned to each of the remaining ℓ machines with probability $1/\ell$. Again, the expected load of these machines is 1. This is a Nash equilibrium since no job can improve by an unilateral move:

- For a small job i assigned to machine j_i, we have $c_i^{j_i} = 1$ and $c_i^j = 1 + 1/k$ for $j \neq j_i$.
- For a large job i, we have $c_i^j = 2 - 1/k$ if $j > m - \ell$ and $c_i^j = 2$ if $j \leq m - \ell$.

The social cost of this equilibrium equals the maximum occupancy of the balls-and-bins experiment where ℓ balls are assigned uniformly at random to ℓ bins. It is well-known that the maximum occupancy of this assignment is $(1 \pm o(1)) \cdot \frac{\ln \ell}{\ln \ln \ell}$ (see, e. g. [7]). □

The ratio between the bounds in Lemma 1 and 2 is maximized by choosing ℓ as large as possible under the constraints specified in Lemma 1. W.l.o.g., let $f(n) = \mathcal{O}(\ln n)$. We set $\ell = \Theta(\sqrt{m}/f(m))$. This way, $SC(\mathbf{P}) \geq (\frac{1}{2} - o(1)) \cdot \frac{\ln m}{\ln \ln m}$ and $SC(\mathbf{F}) \leq 2 + o(1)$. This completes the proof of Theorem 1. □

Let us remark that we can fine-tune the above example such that for $m = 14$ machines and $\ell = 3$ large jobs the expected maximum load of $\mathbf{P}$ is $17/9$ and the expected maximum load of $\mathbf{F}$ is $15/9 + 3/14 + \epsilon < 17/9$, where $\epsilon > 0$ can be made arbitrarily small by increasing the number of small jobs. Thus there is a counterexample to the FMNE conjecture with only 14 machines.

3 Approximating Worst-Case Equilibria

In this section we assume that jobs are ordered such that $w_1 \geq \cdots \geq w_n$. Also, w. l. o. g., we assume that the average load is 1, i. e. $\sum_{i=1}^{n} w_i = m$. Now, we define the quantities

$$M_i := \frac{e + w_i \ln(e + i)}{\ln(e + w_i \ln(e + i))} \qquad \text{and} \qquad M := \max_{i \in [n]} M_i, \tag{1}$$

where $e = 2.71\ldots$ is the Eulerian constant. We will see that the social cost of the worst Nash equilibrium of a given instance is $\Theta(M)$. In the next subsection we establish a lower bound by specifying an algorithm that outputs a Nash equilibrium of value $\Omega(M)$. Subsequently, we prove that an upper bound $\mathcal{O}(M)$ on the social cost of any equilibrium.

3.1 The Algorithm

We present an algorithm that constructs a Nash equilibrium that favors collisions between large jobs on few machines. It proceeds by partitioning the set of jobs into i large jobs and $n-i$ small jobs for a suitable index i. Then, all large jobs are assigned to machines $\{1,\ldots,k\}$ for a minimal k such that these machines are not overloaded, that is, the average load on these machines is at most 1. Additionally, small jobs are moved to machines $\{1,\ldots,k\}$ in order to guarantee that this produces a Nash equilibrium. The index i is chosen such that M_i is maximized. The pseudocode of algorithm Greedy-Nash is given below.

Algorithm 1. The Greedy-Nash algorithm

```
// find suitable threshold that separates large from small jobs
Choose i ∈ [n] such that it maximizes M_i.
// distribute largest jobs on first machines
Let W ← Σ_{j=1}^{i} w_j.
Choose k = ⌈W⌉.
p_j^l ← 1/k for j ∈ {1,...,i} and l ∈ {1,...,k}
// ensure that smaller jobs are satisfied, too
for all jobs j ∈ {i+1,...,n} in weight-decreasing order do
  if W/k ≤ (W_tot − W − w_j)/(m−k) then
    p_j^l ← 1/k for l ∈ {1,...,k}
    W ← W + w_j.
  else
    p_j^l ← 1/(m−k) for l ∈ {k+1,...,m}
  end if
end for
return ((p_i^j)_{i∈[n],j∈[m]}, M_i)
```

Note that the output of algorithm Greedy-Nash as described in the pseudocode has size $n \cdot m$. It can be represented in a compact way by specifying k and the set of jobs assigned to machines $\{1,\ldots,k\}$. This way, the algorithm has linear running time.

Intuitively, the proof of our lower bound M_i proceeds by merging the jobs such they have equal size and applying a Lemma on throwing $\Theta(k/w_i)$ balls with weight w_i into k bins.

Theorem 2. *Let M be defined as in Equation (1). A Nash equilibrium with expected maximal load $\Omega(M)$ can be constructed in time $\mathcal{O}(n)$ provided that the jobs are given in non-increasing order of weight.*

Proof. We first prove that the algorithm constructs a Nash equilibrium. We call the machines $1,\ldots,k$ *left machines* and the machines $k+1,\ldots,m$ *right machines*. Suppose that at the beginning of the for-loop all jobs $i+1,\ldots,n$ are assigned to the machines on the right with uniform probability $1/(m-k)$. Subsequently the algorithm may shift some of these jobs to the machines on the left. Before the loop starts, all jobs in $1, \ldots, i$ are satisfied because they only use the left machines and the expected load on every left machine is W/k whereas the loads on every right machine is $(W_{tot} - W)/(m - k) \geq W_i/k$ by our choice of k.

Note that it is an invariant of the loop that the total weight of jobs on left machines equals the value of the variable W. We have to show that after one pass of the loop job j is satisfied and no jobs in $\{1, \ldots, j-1\}$ become unsatisfied. Since job j goes to the group of machines on which all other jobs induce the smaller expected load, job j is obviously satisfied. If job j is assigned to the right machines the situation does not change and no other job can become unsatisfied. If job j is assigned to the left machines only other jobs on left machines can get unsatisfied. Assume that job $j' < j$ becomes unsatisfied. This job has weight $w_{j'} > w_j$ and this job being unsatisfied means that

$$\frac{W_{tot} - W - w_j}{m - k} < \frac{W + w_j - w_{j'}}{k} \leq \frac{W}{k}.$$

However, if this is the case, then job j would have been assigned to the right machines. Hence, the assignment returned by the algorithm is a Nash equilibrium.

We now show that this assignment has a social cost of at least $\Omega(M_i)$. For the time being, assume that $\sum_{j=1}^{i} w_j > 1$, that is $k \geq 2$ and the average load induced by jobs $1, \ldots, i$ on the left machines is at least $1/2$. For the purpose of the analysis we repeatedly split the jobs $1, \ldots, i-1$ into halves until their weight is in the range $[w_i, 2w_i]$. This way, the number of jobs with weight in $[w_i, 2w_i]$ is some number $i' \geq i$. In [9] it was shown that this inverse *ball fusion* does not increase the expected maximum load on the left machines. Finally, we reduce all job weights down to w_i again not increasing the expected maximum load. The average load on the left machines is still at least $1/4$. Then it follows from Fact 3 in the appendix, that now the expected maximum load is at least

$$\Omega\left(\frac{e + w_i \ln(e + k)}{\ln(e + w_i \ln(e + k))}\right)$$

This term gives a lower bound of $\Omega(M_i)$ if we assume that $w_i \geq 1/k$ as, in this case, $i' \cdot w_i \leq k$ implies $k \geq \sqrt{i'}$ which gives $\ln(e+k) = \Omega(\ln(e+i')) = \Omega(\ln(e+i))$.

We are left with two special cases in both of which we show that $M_i = \mathcal{O}(1)$ and hence is a trivial lower bound.

- $\sum_{j=1}^{i} w_j \leq 1$. In that case, our constructed Nash equilibrium has a trivial lower bound of 1. Furthermore,

  $$M_i \leq e + w_i \ln(e + i) \leq e + w_i \ln(e + 2/w_i) \leq e + 2 + w_1(e - 1) = \mathcal{O}(1)$$

  since $\ln(1 + \epsilon) \leq \epsilon$ for any $\epsilon > 0$.
- $w_i \leq 1/k$. Since $w_i \leq k/i$ we have $i \leq k/w_i \leq 1/w_i^2$ and hence $w_i \leq 1/\sqrt{i}$. Substituting this into M_i yields $M_i = \mathcal{O}(1)$.

Thus in all cases, the social cost is lower bounded by $\Omega(M_i)$. Since i is chosen to maximize M_i, it is also lower bounded by $\Omega(M)$. □

3.2 Upper Bound

The maximum load over all machines is equal to the maximum *height* over all jobs where the height of a job is defined as follows. We assume that jobs are thrown into the machines according to the Nash probability distribution one after another in non-increasing order of weight. The height of job i is the total weight of jobs on its machine at its insertion time. The important property of this definition is that the height of job i does not depend on the assignments of the jobs $1, \ldots, i-1$. More formally, we define indicator variables I_i^j where $I_i^j = 1$ if and only if ball $i \in [n]$ is assigned to machine $j \in [m]$. For any job $i \in [n]$, let j_i denote the machine that job i is assigned to and let

$$X_i = \sum_{k=i}^{n} I_k^{j_i} w_k$$

denote the height of this job. Obviously $H = \max_i\{X_i\}$, the maximal height over all jobs, is equivalent to the maximum load over the machines.

Theorem 3. *Let M be defined as in Equation (1). For any Nash equilibrium, it holds that $\mathbb{E}[H] \leq M$.*

Proof. Consider job i. For $\alpha \geq 1$, let $q = 2\,e\,\alpha\,M_i$. On every machine that job i assigns positive probability to, the expected total load induced by jobs $i+1, \ldots, n$ is upper bounded by 1 since we are at a Nash equilibrium, that is, $\mathbb{E}[X_i - w_i] \leq 1$. Applying a weighted Chernoff bound yields

$$\mathbb{P}[X_i - w_i \geq q] \leq \left(\frac{e}{q}\right)^{q/w_i} .$$

Observe that $q \geq e\sqrt{e + w_i \ln(e+i)}$ as $x/\ln x \geq \sqrt{x}$ for any $x \geq e$. Hence,

$$\begin{aligned}
\mathbb{P}[X_i - w_i \geq q] &\leq \left(\frac{1}{\sqrt{e + w_i \ln(e+i)}}\right)^{2e\alpha\,\ln(e+i)/\ln(e+w_i\ln(e+i))} \\
&= \left(\frac{1}{e}\right)^{e\,\alpha\,\ln(e+i)} \\
&\leq \frac{1}{i^{e\alpha}} \\
&\leq \frac{1}{i^2\,2^\alpha}
\end{aligned}$$

for $i \geq 2$. Applying a union bound we see that

$$\mathbb{P}[X - w_1 \geq \alpha\,M] \leq \mathbb{P}[\exists i : X_i - w_i \geq \alpha \cdot M_i] \leq 2^{-\alpha} \sum_{i=1}^{n} \frac{1}{i^2} \leq 2^{-\alpha} \cdot \frac{\pi^2}{6}$$

and hence

$$\begin{aligned}
\mathbb{E}[X] &\leq w_1 + \int_0^\infty \mathbb{P}[X - w_1 \geq t\,M]\,(M - w_1)\,dt \\
&\leq w_1 + (M - w_1)\left(\int_0^1 \mathbb{P}[X - w_1 \geq t\,M]\;dt + \int_1^\infty \mathbb{P}[X - w_1 \geq t\,M]\;dt\right) \\
&\leq w_1 + (M - w_1)\cdot\left(1 + \frac{\pi^2}{6}\int_1^\infty 2^t\,dt\right) \\
&= \mathcal{O}(M)
\end{aligned}$$

This finishes the proof of the theorem. □

4 Conclusions

We have shown that the fully mixed Nash equilibrium is not the worst-case equilibrium and does not even give a good approximation. In contrast, we have shown that concentrating large jobs on a few machines yields equilibria that approximate the worst-case within a constant factor. As these equilibria can be constructed in linear time we obtained the first constant factor approximation for the worst-case Nash equilibrium.

Our analysis is restricted to identical machines. The question whether worst-case equilibria can be approximated for the case of uniformly related machines remains open and is a challenging problem.

Acknowledgements. We thank Andreas Baltz for helpful comments on an early version of this manuscript.

References

1. Andreas Baltz and Anand Srivastav. On the asymptotic inefficiency of the fully mixed Nash equilibrium. Preprint, 2005.
2. Artur Czumaj and Berthold Vöcking. Tight bounds for worst-case equilibria. In *Proc. 13th Ann. ACM–SIAM Symp. on Discrete Algorithms (SODA)*, pages 413–420, 2002.
3. Dimitris Fotakis, Spyros Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *Proc. 29th Int. EATCS Coll. on Automata, Languages and Programming (ICALP)*, pages 123–134, Malaga, Spain, 2002.
4. Martin Gairing, Thomas Lücking, Marios Mavronicolas, and Burkhard Monien. Computing Nash equilibria for scheduling on restricted parallel links. In *Proc. 36th Ann. ACM. Symp. on Theory of Comput. (STOC)*, pages 613–622, 2004.
5. Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, and Paul G. Spirakis. Extreme Nash equilibria. In *Proc. 8th Italian Conference on Theoretical Computer Science*, pages 1–20, 2003.
6. Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, and Paul G. Spirakis. Structure and complexity of extreme Nash equilibria. *Theoretical Computer Science*, 2005. to appear.

7. 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.
8. Torben Hagerup and Christine Rüb. A guided tour of Chernoff bounds. *Information Processing Letters*, (33):305–308, 1990.
9. Elias Koutsoupias, Marios Mavronicolas, and Paul G. Spirakis. Approximate equilibria and ball fusion. *Theory Comput. Syst.*, 36(6):683–693, 2003.
10. Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proc. 16th Int. Symp. on Theoret. Aspects of Comput. Sci. (STACS)*, pages 404–413. Springer-Verlag, 1999.
11. Marios Mavronicolas and Paul G. Spirakis. The price of selfish routing. In *Proc. 33rd Ann. ACM. Symp. on Theory of Comput. (STOC)*, pages 510–519, 2001.

Appendix

The following facts about balls and bins processes have almost surely been shown somewhere else before. Due to space limitations, we leave the proofs for the full version.

Fact 1. *Let f denote any function in $\omega(1)$. If $n \leq \sqrt{m}/f(m)$ balls are assigned independently and uniformly at random to m bins. Then the expected number of balls in the fullest bin is $1 + o(1)$.*

Fact 2. *Let f denote any function in $\omega(1)$. If $n \leq m \cdot f(m) \cdot \ln m$ balls are assigned independently and uniformly at random to m bins. Then the expected number of balls in the fullest bin is $f(m) \cdot \ln m + O(\sqrt{f(m)} \cdot \ln m) = (1 + o(1)) \cdot f(m) \cdot \ln m$.*

Fact 3. *When n balls are thrown into m bins independently, uniformly at random, the expected number of balls in the fullest bin is*

$$\Omega\left(\frac{n/m + \ln(e + m)}{\ln(e + (m/n)\ln(e + m))}\right).$$

Thus, if the balls have weight $w = m/n$, so that the average load is 1, then the maximum weight over all bins is

$$\Omega\left(\frac{1 + w\ln(e + m)}{\ln(e + w\ln(e + m))}\right).$$