On the Packing of Selfish Items*

Vittorio Bilò[1,2]

[1] Dipartimento di Matematica
Università di Lecce
Provinciale Lecce-Arnesano, P.O. Box 193
73100 Lecce, Italy
email: {vittorio.bilo@unile.it}

[2] Dipartimento di Informatica
Università di L'Aquila
Via Vetoio, Coppito
63100 L'Aquila, Italy
email: {bilo@di.univaq.it}

Abstract

In the non cooperative version of the classical Minimum Bin Packing problem, an item is charged a cost according to the percentage of the used bin space it requires. We study the game induced by the selfish behavior of the items which are interested in being packed in one of the bins so as to minimize their cost. We prove that such a game always converges to a pure Nash equilibrium starting from any initial packing of the items, estimate the number of steps needed to reach one such equilibrium, prove the hardness of computing good equilibria and give an upper and a lower bound for the price of anarchy of the game. Then, we consider a multidimensional extension of the problem in which each item can require to be packed in more than just one bin. Unfortunately, we show that in such a case the induced game may not admit a pure Nash equilibrium even under particular restrictions. The study of these games finds applications in the analysis of the bandwidth cost sharing problem in non cooperative networks.

1. Introduction

The intriguing reality surrounding all the researchers working on the field of algorithms and computational complexity is that of continuously coping with the scarceness of resources while attempting to give a proper solution to practical or theoretical problems. All the results achieved during the years have led to the flourishing of many theories and models now standing as worldwide recognized landmarks in the field. In one of his talks, Papadimitriou [19] called these approaches as *compromises*. The first compromise was due to the scarceness of computational time, thus giving life to the theory of approximation algorithms. The second one came from the scarceness of information on the problem which is asked to be solved (*online problems*) which led to the definition of competitive analysis.

Recently, a third compromise has arisen, thus giving life to a new challenging and intriguing research direction, due to the introduction of classical aspects of Economics and Game Theory such as selfishness and rational behavior of the agents handling the variables characterizing a given problem. Differently from the past, in fact, some of the variables of the problem may not be willing to implement the (optimal, approximate or competitive) solution computed with the classical techniques, for several reasons. This means that now we have to cope with the total lack (not only scarceness) of cooperation among the entities involved in our problems. The reasons justifying this new approach come from the affirmation of huge unregulated networks (as, for instance, the Internet) whose users are not interested in the optimization of some global or social function, but only in satisfying their requirements at the minimum cost. Under these assumptions, classical network optimization problems are to be modeled as non cooperative strategic games and their Nash equilibria are analyzed and compared with respect to the social optimum.

The problem we consider in this paper is the classical Minimum Bin Packing problem with the constraint that the items to be packed are handled by selfish agents. All the bins have the same fixed cost and the cost of a bin is shared among all the items it contains according to the normalized fraction of the bin they use. More formally, if we denote the height of item i as a_i, the j-th bin as B_j and the sum of the heights of the items packed into B_j as H_j, we have that the cost

*This work was partially supported by the European Union under COST Action 293 GRAAL.

1-4244-0054-6/06/$20.00

paid by i for using B_j is $cost(i, j) = \frac{a_i}{H_j}$. Suppose that item i is packed into B_j. Since i wants to minimize its cost, it will migrate from B_j each time it will detect another bin $B_{j'}$ such that $cost(i, j') < cost(i, j)$. This inequality holds for any j' such that $H_{j'} + a_i > H_j$, thus an item will migrate each time it will detect a bin in which it fits better with respect to the unused space. The social cost that we want to minimize is the number of used bins.

The bin packing problem has been one of the first optimization problems being considered under both the first and the second compromise. Thus it seems to be a very appropriate case study for appreciating the loss in optimality of the solutions that can be achieved when new compromises need to be considered.

Moreover, although this is a state-of-the-art problem and for such a reason its study under these new assumptions is clearly self-motivated, this problem has direct practical implications and can be used to model the following interesting scenario. Consider two nodes, a source and a destination, connected by a potentially infinite number of links and a set of users wishing to establish a connection between the two nodes having a certain bandwidth. All the links can carry the same fixed bandwidth at the same fixed cost. The cost of each link is shared among its users according to the same protocol described above for the selfish bin packing problem, that is according to the normalized level of usage. For such a reason, users, which are assumed to be selfish, want to route their traffic on the most exploited link. If we suppose that the establishment of a new link causes some inefficiencies or extra-costs in the whole system, we have that the best performances are achieved when the overall number of used links is minimized. Extending the underlying model from a set of parallel links between two nodes to a generic multigraph clearly gives rise to a multidimensional bin packing game, where an item may require to be packed in more than one bin.

The main algorithmic issues coming from the study of non cooperative games are the following: (1) Proving the existence of pure equilibria [4][1]; (2) Proving the convergence to a Nash equilibrium starting from any initial combination of the agents' choices and estimating the convergence time; (3) Finding a generic Nash equilibrium or an equilibrium having particular properties (for instance, the one minimizing the global cost); (4) Measuring the *price of anarchy*, that is the ratio between the worst Nash equilibrium and the *social optimum*, that is the optimal solution that could be achieved if all players cooperated.

[1]The existence of randomized or mixed equilibria is stated by Nash's Theorem [18].

The price of anarchy [13], in particular, is the notion that captures the fundamental aspects of the third compromise as a new loss in terms of distance from the optimal solution because of the lack of coordination among the users.

During the last six years, there have been lots of results in the study of games induced by selfish routing in non cooperative networks.

The first model (the *task allocation model*), introduced by Koutsoupias and Papadimitriou in [13], evaluates the link congestion in a network consisting of two nodes connected by a set of parallel links. This model essentially reduces to the problem of allocating selfish tasks to parallel machines, i.e., the classical Minimum Multiprocessor Scheduling problem analyzed under the third compromise. The analysis of the price of anarchy for this model has been improved in [17] and definitively characterized in [12, 2]. As to pure Nash equilibria, Mehlhorn proved the existence of at least one such equilibrium (Theorem 1 in [6]). In the same paper, it is also shown that computing the best and the worst pure equilibrium is NP-hard, while a polynomial-time algorithm determining a pure equilibrium is given. The latter result has been improved in [5], where an efficient algorithm is presented computing a pure equilibrium starting from any allocation of tasks without increasing its social cost. As a consequence, by applying this algorithm to the solution returned by the PTAS for the centralized problem due to Hochbaum and Shmoys [11], it is possible to obtain pure equilibria whose social cost is arbitrarily close to the social optimum. Finally, in [3] it is shown that the non cooperative game induced by the selfish behavior of the tasks is always able to converge to a pure equilibrium[2] and the number of steps needed for convergence in several different environments are determined. Other results on this model can be found in [10, 16, 7].

Another well-studied model (the *flow routing model*) has been introduced by Roughgarden and Tardos in [25] where the cost of using a link is not only a function of the load, but it is a latency function. In this model we are given a graph and a set of communication requests between different nodes and it is assumed that each user controls a negligible amount of traffic so that it can be modeled as a flow. The authors prove existence and uniqueness[3] of pure equilibria and bound the price of anarchy for linear and non linear latency functions. Other interesting results on this model can be found in [21, 22, 23, 24].

The most important difference between these two models is that the connection requests issued by the

[2]Note that this result improves on Mehlhorn's Theorem.
[3]With respect to the social cost.

agents are unsplittable in the task allocation model, while they are splittable in the routing flow model, (see [1] for a very good survey on selfish routing). The latter property often induces "easier" games. In fact, every game induced by splittable requests is equivalent to a congestion game, a well-known class of games always possessing pure equilibria [20], while in the case of unsplittable requests in [14] it has been shown that no pure equilibria may exist even under monotonicity of the delay function. Finally, some recent results have been achieved by analyzing the resulting model obtained by combining some of the aspects of the above two ones in [9, 8, 15].

We study the non cooperative Minimum Bin Packing problem and the properties of its pure Nash equilibria. We assume that each item represents an unsplittable request for connection. Our results are omnicomprehensive in the sense that they cover the majority of the algorithmic aspects listed before. In particular, we show that the game always converges to a pure Nash equilibrium, estimate the number of steps needed to converge, prove the hardness of computing the best pure Nash equilibrium and provide an upper and a lower bound on the price of anarchy of the game. We also consider a multidimensional extension of the problem in which an item can require to be packed in more than just one bin. Unfortunately, we show that in such a case the induced game may not admit a pure Nash equilibrium even when there are only two different types of items or all the items require the same number of bins.

2. Existence Proof and Convergence Time

In this section we provide the proof of existence of at least one pure Nash equilibrium for the Minimum Bin Packing game and estimate the number of steps needed to converge starting from any initial packing. We denote as σ the input sequence of items, as $n = |\sigma|$ the number of items, as a_{MIN} the minimum height of an item and as A the sum of the heights of all the items belonging to σ, that is $A = \sum_{i=1}^{n} a_i$.

Theorem 2.1 *The bin packing game always converges to a pure Nash equilibrium.*

Proof. It is quite easy to show the convergence of the game by reasoning on the effects caused by an improving migration performed by an item on the whole system. In fact, let C be a generic configuration of bins obtained by listing them in non increasing order starting from the most filled bin to the most empty one. If C is not an equilibrium, then there exist two bins B_j and $B_{j'}$ and an item i packed into B_j such that i fits into $B_{j'}$ and $cost(i, j') < cost(i, j)$. Let us denote by C' the new configuration obtained by letting i migrate from B_j to $B_{j'}$ and then reordering the bins. Since the new choice performed by i only affects H_j and $H_{j'}$, it results that C' is lexicographically strictly greater than C. The sequence of possible configurations that can be generated is clearly upper-bounded by the configuration $C^* = \{\underbrace{1, \ldots 1}_{\lfloor A \rfloor}, A - \lfloor A \rfloor\}$, hence the sequence of improvements always converges to a local maximum after a finite number of migrations. $\diamondsuit$

It is worth noting here that such a proof holds for any possible cost sharing function charging an item i a cost depending only on a_i and on the amount of occupied space in the bin chosen by i even if we relax the assumption that all the bins have the same cost. In such a general case our bin packing game can be seen as a particular variant of the task allocation model whose underlying environment is given by at most n unrelated machines. Thus, it is possible to use some of the results in [3] on the convergence time to a pure Nash equilibrium for unrelated machines.

Since we assume that all the bins have the same cost, our model is similar to the task allocation model performed on identical machines. However, the particular nature of our cost sharing function makes our game significantly more complicated, thus it seems hard to obtain similar results to those achieved in [3].

In order to bound the convergence time in our model, we define a suitable potential function which proves to be useful in the case in which all the heights a_i are rational numbers, i.e., $a_i = \frac{num_i}{den_i}$, where num_i and den_i are integers for any $1 \leq i \leq n$. Let M be the minimum common multiplier of the values den_i, $1 \leq i \leq n$.

Theorem 2.2 *The bin packing game converges to a pure Nash equilibrium in at most $\frac{Mn}{2a_{MIN}}$ improving steps in the case in which all items have rational heights.*

Proof. Given any configuration of the bins $C^t = \{B_1, \ldots, B_{k(t)}\}$ at step t, we define the following potential function $\Phi(t) = 2^{\sum_{j=1}^{k(t)} H_j^2}$. If at step t item i performs an improving step by migrating from B_j to $B_{j'}$, we have that $\frac{\Phi(t+1)}{\Phi(t)} = 2^{(H_j - a_i)^2 + (H_{j'} + a_i)^2 - H_j^2 - H_{j'}^2} = 2^{2a_i(H_{j'} + a_i - H_j)} \geq 2^{2a_i/M} \geq 2^{2a_{MIN}/M}$. The potential function of the system is lower-bounded by 1 and upper-bounded by 2^n. Since at each improving step the potential increases by a multiplicative factor of at least $2^{2a_{MIN}/M}$, we have that it will reach its maximum in at most $\log_{2^{2a_{MIN}/M}} 2^n = \frac{\log 2^n}{\log 2^{2a_{MIN}/M}} = \frac{Mn}{2a_{MIN}}$. $\diamondsuit$

By exploiting the same potential function, we can also derive a bound on the number steps needed to reach an approximate pure Nash equilibrium, defined as follows.

Definition 2.3 *A configuration of bins C is an ϵ-pure Nash equilibrium, with $0 < \epsilon \leq 1$, if for any pair of bins $B_j, B_{j'} \in C$ and any item a_i packed into B_j it holds $cost(i,j') \geq cost(i,j) \cdot \epsilon$.*

Clearly, it follows from the definition that a 1-pure Nash equilibrium is a pure Nash equilibrium. Strictly speaking, an ϵ-pure Nash equilibrium is a configuration of bins in which no item possesses an ϵ-improving step, where an ϵ-improving step performed by item i is a migration from B_j to $B_{j'}$ such that $cost(i,j') < cost(i,j) \cdot \epsilon$. According to our particular cost sharing function, item i can perform an ϵ-improving step, migrating from B_j to $B_{j'}$, if $(H_{j'} + a_i)\epsilon > H_j$.

Theorem 2.4 *The bin packing game converges to an ϵ-pure Nash equilibrium in at most $\frac{n}{2(1-\epsilon)a_{MIN}^2}$ ϵ-improving steps for any $0 < \epsilon < 1$.*

Proof. If at step t item i performs an ϵ-improving step by migrating from B_j to $B_{j'}$, we have that $\frac{\Phi(t+1)}{\Phi(t)} = 2^{(H_j - a_i)^2 + (H_{j'} + a_i)^2 - H_j^2 - H_{j'}^2} = 2^{2a_i(H_{j'} + a_i - H_j)} = 2^{2a_i[(1-\epsilon)(H_{j'}+a_i) + \epsilon(H_{j'}+a_i) - H_j]} > 2^{2a_i(1-\epsilon)(H_{j'}+a_i)} \geq 2^{2(1-\epsilon)a_i^2} \geq 2^{2(1-\epsilon)a_{MIN}^2}$. By using the same arguments as in the previous theorem the claim follows. $\diamondsuit$

Finally, it is worth noting that the number of used bins in any sequence of migrations is always non increasing, since no item can achieve an improvement by migrating to an empty bin. This means that there always exists a solution, among the optimal ones, which is a pure Nash equilibrium and moreover, if a central authority would be able to enforce an initial solution having an approximation ratio equal to c, then the social cost of the resulting pure Nash equilibrium would be at most c times the social optimum. The above discussion proves the following result.

Theorem 2.5 *It is NP-hard to compute the best pure Nash equilibrium for the bin packing game.*

3. Bounding the Price of Anarchy

We now prove an upper bound on the price of anarchy of our game and also provide a very close lower bound.

Theorem 3.1 *Any pure Nash equilibrium for the bin packing game uses at most $\lceil \frac{5}{3} Opt \rceil + 1$ bins, where Opt is the number of bins used by an optimal centralized solution.*

Proof. Let C be the configuration of bins yielded by a generic pure Nash equilibrium obtained by listing the bins in non increasing order. We first remark that all the bins are filled for more than a half, except for at most one bin, otherwise an improving step would trivially exist.

Suppose that the average occupation on the first Opt bins in C is at least 3/4. This means that the "wasted" space in the first Opt bins in C with respect to the optimal solution is at most $\frac{1}{4}Opt$. This quantity is spread among all the additional bins used by C. Since they are filled for at least more than one half of their height (except for at most one bin), at most $\lceil \frac{1}{2} Opt \rceil$ bins are needed to pack such left-over items, thus the number of bins in C is at most $Opt + \lceil \frac{1}{2}Opt \rceil = \lceil \frac{3}{2} Opt \rceil < \lceil \frac{5}{3} Opt \rceil + 1$.

Now suppose that the average occupation on the additional bins in C without considering the least loaded one is at least 5/8. This implies that the "wasted" space in the first Opt bins in C is at most $\frac{3}{8}Opt$ and is spread among additional bins having an average occupation of at least 5/8. It follows that the number of bins in C is at most $Opt + \lceil \frac{8}{5}\frac{3}{8} Opt \rceil = \lceil \frac{8}{5} Opt \rceil < \lceil \frac{5}{3} Opt \rceil + 1$.

It is left to consider the case in which the average occupation on the first Opt bins in C is less than 3/4 and the one on the additional bins without considering the least loaded one is less than 5/8. Let us denote by F the set of all the first Opt bins in C which are filled for less than 3/4. If $|F| = 1$, using the same argument as in the first case we obtain that the "wasted" space in the first Opt bins in C with respect to the optimal solution is at most $\frac{1}{4}(Opt - 1) + \frac{1}{2} = \frac{1}{4}Opt + \frac{1}{4}$, thus the number of bins in C becomes at most $Opt + \lceil \frac{1}{2}Opt + \frac{1}{2} \rceil \leq \lfloor \frac{3}{2} Opt \rfloor + 1 < \lceil \frac{5}{3} Opt \rceil + 1$. If $|F| > 1$, then any of the $|F| - 1$ least filled bins in F cannot contain an item i such that $a_i \leq 1/4$ because otherwise i could perform an improving step by migrating to the most filled bin in F. Moreover, they cannot contain an item i such that $1/4 \leq a_i \leq 3/8$ because otherwise i could migrate to the second least loaded additional bin. Such a bin, in fact, must have an available space of at least 3/8 and of at most 1/2 by hypothesis. As a consequence, all the bins in F except for one must contain a single item. Let $f = |F|$ and g be the number of additional bins, i.e., the difference between the number of bins used by C and Opt. Using a similar argument it is not difficult to see that all the g additional bins cannot contain an item i such that $0 < a_i \leq 3/8$, thus we have identified $f + g - 1$ items of height greater than $1/2$ which is clearly a lower bound on the value of the optimal solution, that is $f + g - 1 \leq Opt$. This means that the number of bins which are filled for more than $3/4$ are $Opt - f \geq g - 1$. Thus the overall occupation in the first Opt bins in C is at least $\frac{3}{4}(g-1) + \frac{Opt-g+1}{2} = \frac{g+2\cdot Opt-1}{4}$, which means that the "wasted" space in the first Opt bins in C is at most $Opt - \frac{g+2\cdot Opt-1}{4}$. It follows that the number of bins in C is at most $Opt + 2(Opt - \frac{g+2\cdot Opt-1}{4}) = 2Opt - \frac{(g-1)}{2}$. Since such a number is also upper-bounded by the value $Opt + g$, we have that the number of bins in C is at most $\lceil \frac{5}{3}Opt \rceil + 1$ when $g = \frac{2Opt+1}{3}$. $\diamondsuit$

Theorem 3.2 *The price of anarchy for the bin packing game cannot be better than* $8/5$ *even when the social optimum goes to infinity.*

Proof. Consider the set of items listed in Figure 1a. The configuration of the 8 bins represented in Figure 1b is a pure Nash equilibrium since no item can perform an improving migration. In fact, no item in the last 5 bins can migrate elsewhere because there is not enough space to pack it. The items in B_3 are in an almost entirely filled bin and it can be easily seen that they cannot lower their cost. Finally, each of the items in the first two bins can migrate only in one of the last 5 bins, but none of them would achieve an improvement. The optimal solution, showed in Figure 1c, uses only 5 bins, hence, since this configuration can be replicated an arbitrarily number of times and still remains an equilibrium, the thesis follows. $\diamondsuit$

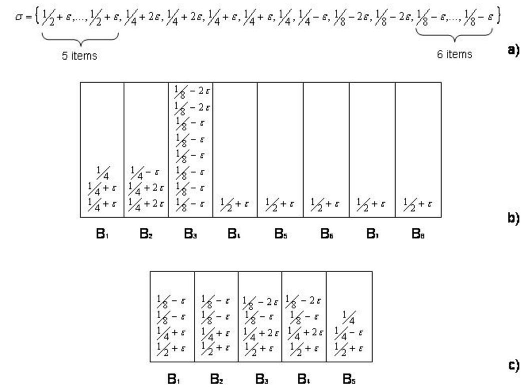

Figure 1. a) The set of items σ. b) The worst pure Nash equilibrium. c) The social optimum.

4. A Multidimensional Extension

In this section we consider the natural extension of our game to the case in which the bins are divided into different classes and each item requires to be packed in a set of bins each belonging to a different class. More formally, let $\mathcal{B}$, with $|\mathcal{B}| = m$, be the set of different classes of bins, each item $i \in \sigma$ has a set of possible requirements $R(i)$, defined by a function $R : \sigma \rightarrow 2^{\mathcal{P}}$, where $\mathcal{P} \subseteq 2^{\mathcal{B}}$. All the bins belonging to the same class have the same height and the same cost and we want to pack each item by choosing one of its possible requirements so as to minimize the number of used bins.

In practice, this problem models a network $G = (V, E, c)$, with $c : E \rightarrow I\!R^+$, in which different pairs of nodes $\{x, y\}$ want to communicate by using one of the possible paths between x and y. By modeling each edge of G with a class of bins, that is setting $\mathcal{B} = E$, and each request of communication with an item $i \in \sigma$, we have that $R(i = \{x, y\})$ becomes a set of paths connecting x to y. Again, in this setting, each item is charged a cost which is equal to the sum of the cost of the used bins according to the level of their usage. We prove that this extended model gives rise to games which may not admit a pure equilibrium even under special restrictions.

Theorem 4.1 *The extended bin packing game may not admit a pure Nash equilibrium even for the case in which all bins have the same cost and there are only two different types of items.*

Proof. Consider the game induced by the graph depicted in Figure 2, in which each depicted edge represents a chain of k real edges, where k is the number associated with each depicted edge and set $\sigma = \{a_1 = a_2 = \frac{1}{4}, a_3 = a_4 = \frac{1}{2}\}$. The set of paths Π_i that can be chosen from item i, thus representing i's set of strategies, is defined as follows:

$\Pi_1 = \{p_1^1 = e_{(1,2)}e_{(2,3)}, p_1^2 = e_{(1,7)}e_{(7,8)}e_{(8,3)}\}$,
$\Pi_2 = \{p_2^1 = e_{(1,2)}e_{(2,3)}, p_2^2 = e_{(1,3)}\}$,
$\Pi_3 = \{p_3^1 = e_{(1,2)}e_{(2,6)}, p_3^2 = e_{(1,5)}e_{(5,6)}\}$,
$\Pi_4 = \{p_4^1 = e_{(1,3)}e_{(3,4)}, p_4^2 = e_{(1,5)}e_{(5,7)}e_{(7,8)}e_{(8,4)}\}$.

Since there are 4 items each of which has 2 strategies, we have to consider 16 different combinations of choices. We show in Figure 4 that none of them

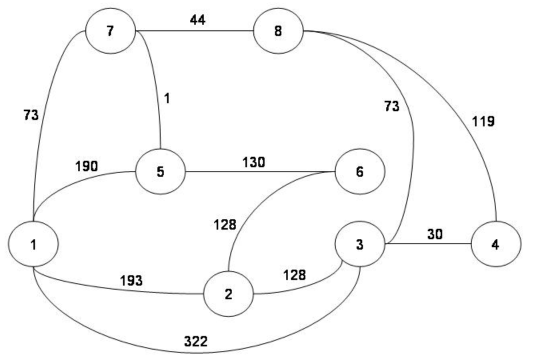

Figure 2. A communication graph inducing a game not admitting a pure Nash equilibrium even when there are only two types of items.

represents a Nash equilibrium. ◇

Clearly, the case in which there is only one type of item can be easily reduced to a congestion game thus always admitting a pure Nash equilibrium. Hence the previous theorem states that any non trivial instance of the extended bin packing game may not converge to a pure Nash equilibrium. Other special instances of the game can be obtained by allowing particular topologies for the underlying communication graph G from which the game generates. We prove in the following theorem the non existence of a pure Nash equilibrium even when all the paths that can be chosen by all the items have the same number of edges.

Theorem 4.2 *The extended bin packing game may not admit a pure Nash equilibrium even for the case in which all bins have the same cost and all the allowed paths in the underlying graph are of the same length.*

Proof. We use the underlying communication graph depicted in Figure 3 and set $\sigma = \{a_1 = a_2 = \frac{1}{4}, a_3 = a_4 = \frac{1}{2}, a_5 = a_6 = \frac{1}{10}, a_7 = \frac{1}{9}, a_8 = 0.207\}$. The set of paths Π_i that can be chosen from item i are defined as follows:

$\Pi_1 = \{p_1^1 = e_{(1,2)}e_{(2,9)}e_{(9,16)}e_{(16,3)}, p_1^2 = e_{(1,7)}e_{(7,8)}e_{(8,3)}\}$,
$\Pi_2 = \{p_2^1 = e_{(1,2)}e_{(2,9)}e_{(9,16)}e_{(16,4)}, p_2^2 = e_{(1,11)}e_{(11,4)}\}$,
$\Pi_3 = \{p_3^1 = e_{(1,2)}e_{(2,6)}, p_3^2 = e_{(1,5)}e_{(5,10)}e_{(10,6)}\}$,
$\Pi_4 = \{p_4^1 = e_{(1,11)}e_{(11,4)}, p_4^2 = e_{(1,5)}e_{(5,7)}e_{(7,8)}e_{(8,4)}\}$,
$\Pi_5 = \{e_{(1,12)}e_{(12,5)}e_{(5,10)}\}$,
$\Pi_6 = \{e_{(1,14)}e_{(14,11)}e_{(11,4)}\}$,
$\Pi_7 = \{e_{(1,13)}e_{(13,2)}e_{(2,9)}\}$,
$\Pi_8 = \{e_{(1,7)}e_{(7,15)}e_{(15,8)}e_{(8,3)}\}$.

Notice that items 5, 6, 7 and 8 have no choice possibilities. Thus, there are 4 items each of which has 2 strategies for 16 different combinations of choices. We show in Figure 5 that none of them represents a Nash equilibrium. ◇

It may be pointed out here that these negative results should suggest us to turn our attention towards the use of mixed Nash equilibria in order to characterize our game in its extended version. This can clearly be a possible research direction, however we must stress that the notion of randomized choices cannot be introduced as easily as one may think. In fact, if from one side one can use randomization in order to chose one of its available connection paths on the underlying graph, on the other side it is not so natural which rule to apply in order to choose the particular bin where to pack its request.

5. Open Problems

Two natural open problems are the determination of a lower bound on the number of steps needed to converge to an equilibrium and the bridging of the gap between upper and lower bound on the price of anarchy.

As an extension, we have also considered the case in which an item may require to be simultaneously packed in more than just one bin and provided results of non existence of pure equilibria even under special restrictions on the input instances. It would be good to improve these results by allowing each item to choose any possible path in the underlying communication graph.

Acknowledgements: the author thanks Michele Flammini for many helpful comments and discussions.

References

[1] A. Czumaj. *Handbook of Scheduling*, chapter Selfish Routing on the Internet. CRC Press, 2004.

[2] A. Czumaj and B. Vocking. Tight bound for the worst-case equilibria. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Mathematics (SODA)*, pages 413–420. SIAM, 2002.

[3] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibria. In *Proc. of the 30th Annual International Colloquium on Automata, Languages and Programming (ICALP)*, volume 2719 of *LNCS*, pages 502–513.

[4] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure equilibria. In *Proc. of the 36th*

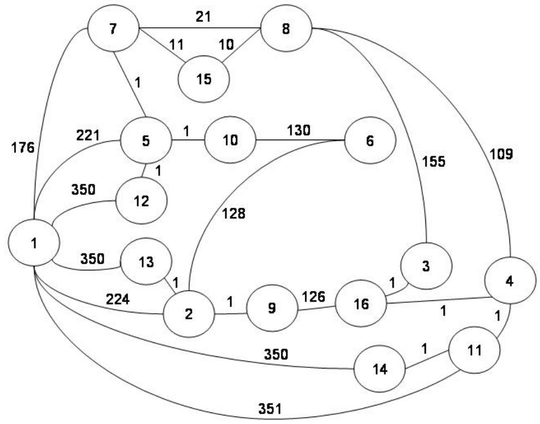

Figure 3. A communication graph inducing a game not admitting a NE even when all the allowed paths are of the same length.

ACM Symposium on Theory of Computing (STOC), pages 604–612. ACM Press, 2004.

[5] R. Feldmann, M. Gairing, T. Lucking, B. Monien, and M. Rode. Nashification and the coordination ratio for a selfish routing game. In *Proc. of the 30th International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS, pages 514–526. Springer, 2003.

[6] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of nash equilibria for a selfish routing game. In *Proc. of the 29th International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS, pages 123–134. Springer, 2002.

[7] M. Gairing, T. Lucking, M. Mavronicolas, and B. Monien. Computing nash equilibria for scheduling on restricted parallel links. In *Proc. of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 613–622. ACM Press, 2004.

[8] M. Gairing, T. Lucking, M. Mavronicolas, and B. Monien. The price of anarchy for polynomial social cost. In *Proc. of the 29th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 3153 of *LNCS*, pages 574–585. Springer, 2004.

[9] M. Gairing, T. Lucking, M. Mavronicolas, B. Monien, and M. Rode. Nash equilibria in discrete routing games with convex latency functions. In *Proc. of the 31st International Colloquium on Automata, Languages and Programming (ICALP)*, volume 3142 of *LNCS*, pages 645–657. Springer, 2004.

[10] M. Gairing, T. Lucking, M. Mavronicolas, B. Monien, and P. Spirakis. Extreme nash equilibria. In *Proc. of the 8th Italian Conference on Theoretical Computer Science, (ICTCS)*, volume 2841 of *LNCS*, pages 1–20. Springer, 2003.

[11] D. Hochbaum and D. Shmoys. A polynomial approximation scheme for scheduling on uniform processors: Using the dual approximation approach. *SIAM Journal of Computing*, 17(3):539–551, 1988.

[12] E. Koutsoupias, M. Mavronicolas, and P. Spirakis. Approximate equilibria and ball fusion. In *Proc. of the 9th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*. Carleton Scientific, 2002.

[13] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *LNCS*, pages 387–396, 1999.

[14] L. Libman and A. Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, 2001.

[15] T. Lucking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. In *Proc. of the 21st Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2996 of *LNCS*, pages 547–558. Springer, 2004.

[16] T. Lucking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis, and I. Vrto. Which is the worst-case nash equilibrium? In *Proc. of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 2747 of *LNCS*, pages 551–561. Springer, 2003.

[17] M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proc. of the 33rd Annual ACM Symposium on the Theory of Computing (STOC)*, pages 510–519, 2001.

[18] J. F. Nash. Equilibrium points in n-person games. In *Proc. of the National Academy of Sciences*, volume 36, pages 48–49, 1950.

[19] C. H. Papadimitriou. Networks and games. *Banquet of the 18th International Parallel and Distributed Processing Symposium (IPDPS)*, 2004.

[20] R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[21] T. Roughgarden. Designing networks for selfish users is hard. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 472–481. IEEE Computer Society Press, 2001.

[22] T. Roughgarden. Stackelberg scheduling strategies. In *Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 104–113. ACM Press, 2001.

[23] T. Roughgarden. How unfair is optimal routing? In *Proc. of the 13rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 203–204. SIAM, 2002.

[24] T. Roughgarden. The price of anarchy is independent of the network topology. In *Proc. of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 428–437. ACM Press, 2002.

[25] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of ACM*, 49(2):236–259, 2002.

p_1	p_2	p_3	p_4	i	$cost(i, p_i)$	$cost(i, p'_i)$
p_1^1	p_2^1	p_3^1	p_4^1	2	$\frac{193}{4} + \frac{128}{2} = 112.25$	$\frac{322}{3} = 107.333...$
p_1^1	p_2^1	p_3^1	p_4^2	4	354	$\frac{322 \cdot 2}{3} + 30 = 244.666...$
p_1^1	p_2^1	p_3^2	p_4^1	2	$\frac{321}{2} = 160.5$	$\frac{322}{3} = 107.333...$
p_1^1	p_2^1	p_3^2	p_4^2	3	$\frac{190}{2} + 130 = 225$	$\frac{193}{2} + 128 = 224.5$
p_1^1	p_2^2	p_3^1	p_4^1	1	$\frac{193}{3} + 128 = 192.333...$	190
p_1^1	p_2^2	p_3^1	p_4^2	2	322	$\frac{193}{4} + \frac{128}{2} = 112.25$
p_1^1	p_2^2	p_3^2	p_4^1	1	321	190
p_1^1	p_2^2	p_3^2	p_4^2	1	321	$\frac{44}{3} + 146 = 160.666...$
p_1^2	p_2^1	p_3^1	p_4^1	1	190	$\frac{193}{4} + \frac{128}{2} = 112.25$
p_1^2	p_2^1	p_3^1	p_4^2	1	$\frac{44}{3} + 146 = 160.666...$	$\frac{193}{4} + \frac{128}{2} = 112.25$
p_1^2	p_2^1	p_3^2	p_4^1	1	190	$\frac{321}{2} = 160.5$
p_1^2	p_2^1	p_3^2	p_4^2	1	$\frac{44}{3} + 146 = 160.666...$	$\frac{321}{2} = 160.5$
p_1^2	p_2^2	p_3^1	p_4^1	3	321	320
p_1^2	p_2^2	p_3^1	p_4^2	3	321	$\frac{190}{2} + 130 = 225$
p_1^2	p_2^2	p_3^2	p_4^1	4	$\frac{322 \cdot 2}{3} + 30 = 244.666...$	$\frac{190}{2} + \frac{44 \cdot 2}{3} + 120 = 244.333...$
p_1^2	p_2^2	p_3^2	p_4^2	2	322	321

Figure 4. The first four columns of the table contain the path chosen by each item thus representing the configuration of choices. The fifth column contains an item i which has an improving migration. Finally, in the two last columns it is reported, respectively, the current cost $cost(i, p_i)$ charged to item i and the new cost $cost'(i, p'_i)$ charged to i after migrating to p'_i.

p_1	p_2	p_3	p_4	i	$cost(i, p_i)$	$cost(i, p'_i)$
p_1^1	p_2^1	p_3^1	p_4^1	2	$\frac{224}{4} + \frac{9}{22} + \frac{126}{2} + 1 = 120.409...$	$\frac{351}{3} + \frac{5}{17} = 117.294...$
p_1^1	p_2^1	p_3^1	p_4^2	4	352	$351 + \frac{10}{17} = 351.588...$
p_1^1	p_2^1	p_3^2	p_4^1	2	$\frac{224}{2} + \frac{9}{22} + \frac{126}{2} + 1 = 176.409...$	$\frac{351}{3} + \frac{5}{17} = 117.294...$
p_1^1	p_2^1	p_3^2	p_4^2	3	$\frac{221}{2} + \frac{5}{6} + 130 = 241.333...$	$\frac{224}{2} + 128 = 240$
p_1^1	p_2^2	p_3^1	p_4^1	1	$\frac{224}{3} + \frac{9}{13} + 127 = 202.358...$	$\frac{331 \cdot 0.25}{0.457} + 21 = 202.072...$
p_1^1	p_2^2	p_3^1	p_4^2	2	$351 + \frac{5}{7} = 351.714...$	$\frac{224}{4} + \frac{9}{22} + \frac{126}{2} + 1 = 120.409...$
p_1^1	p_2^2	p_3^2	p_4^1	1	$351 + \frac{9}{13} = 351.692...$	$\frac{331 \cdot 0.25}{0.457} + 21 = 202.072...$
p_1^1	p_2^2	p_3^2	p_4^2	1	$351 + \frac{9}{13} = 351.692...$	$\frac{331 \cdot 0.25}{0.457} + \frac{21}{3} = 188.072...$
p_1^2	p_2^1	p_3^1	p_4^1	1	$\frac{331 \cdot 0.25}{0.457} + 21 = 202.072...$	$\frac{224}{4} + \frac{9}{22} + \frac{126}{2} + 1 = 120.409...$
p_1^2	p_2^1	p_3^1	p_4^2	1	$\frac{331 \cdot 0.25}{0.457} + \frac{21}{3} = 188.072...$	$\frac{224}{4} + \frac{9}{22} + \frac{126}{2} + 1 = 120.409...$
p_1^2	p_2^1	p_3^2	p_4^1	3	$351 + \frac{5}{6} = 351.833...$	$\frac{224 \cdot 2}{3} + 128 = 277.333...$
p_1^2	p_2^1	p_3^2	p_4^2	1	$\frac{331 \cdot 0.25}{0.457} + \frac{21}{3} = 188.072...$	$1 + \frac{350}{2} + \frac{9}{22} = 176.409...$
p_1^2	p_2^2	p_3^1	p_4^1	3	352	$351 + \frac{5}{6} = 351.833...$
p_1^2	p_2^2	p_3^1	p_4^2	3	352	$\frac{221}{2} + \frac{5}{6} + 130 = 241.333...$
p_1^2	p_2^2	p_3^2	p_4^1	4	$\frac{351 \cdot 2}{3} + \frac{10}{17} = 234.588...$	$\frac{221}{2} + \frac{21 \cdot 2}{3} + 110 = 234.5$
p_1^2	p_2^2	p_3^2	p_4^2	2	$351 + \frac{5}{7} = 351.714$	$351 + \frac{9}{13} = 351.692$

Figure 5. The first four columns of the table contain the path chosen by each item thus representing the configuration of choices. The fifth column contains an item i which has an improving migration. Finally, in the two last columns it is reported, respectively, the current cost $cost(i, p_i)$ charged to item i and the new cost $cost'(i, p'_i)$ charged to i after migrating to p'_i.