The Price of Envy-Freeness in Machine Scheduling

Vittorio Bilò [a] Angelo Fanelli [b] Michele Flammini [c] [d] Gianpiero Monaco [c] and Luca Moscardelli [e]

[a]Dipartimento di Matematica e Fisica "Ennio De Giorgi", Università del Salento, Italy

[b]CNRS, (UMR-6211), France

[c]Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica, Università de L'Aquila, Italy

[d]Gran Sasso Science Institute, L'Aquila, Italy

[e]Dipartimento di Economia, Università di Chieti-Pescara, Italy

The evolution of scheduling closely tracked the development of computers. Given m machines that have to process n jobs, minimizing the makespan of an assignment of the jobs to the machines is one of the most well-studied problem in the Theory of Algorithms. In more details, assuming that the processing of job i on machine j requires time $p_{ij} > 0$, the completion time of machine j (under a certain assignment) is given by the sum of the processing times of all the jobs allocated to j. The makespan of an assignment is the maximum completion time among all the machines (we stress that an assignment is not forced to use all the available machines) and the objective of the scheduling problem is to find an assignment of minimum makespan.

In the literature, three different models of machines have been adopted. The general setting illustrated above is called scheduling problem with *unrelated machines*. An interesting particular scenario is the case with *related machines*, where each job i has a load $l_i > 0$ and each machine j has a speed of processing $s_j > 0$, and thus the processing time of job i on machine j is given by $p_{ij} = l_i/s_j$. Finally, the even more specific setting in which the speed of each machine is 1 is referred to as the scheduling problem with *identical machines*. Even this latter problem is NP-hard.

The approximability of the scheduling problem has been well understood for all the three models described above. However, all the proposed solutions do not envisage fair allocations in which no machine prefers (or envies) the set of the tasks assigned to another machine, i.e., for which her completion time would be strictly smaller. In the literature, such fairness property is referred to as "envy-freeness" [3,4]. Specifically, consider a scenario in which a set of tasks (jobs) has to be allocated among employees (machines) in such a way that the last task finishes as soon as possible. It is natural to consider fair allocations, that is allocations where no employee prefers (or envies) the set of tasks assigned to some other employee, i.e., a set of tasks for which her completion time would be strictly smaller than her actual one.

It is possible to consider two different variants of this model, depending on the fact that an employee (i) can envy the set of tasks assigned to any other employee or (ii) can only envy the set of tasks of other employees getting at least one job: in the latter case, employees not getting any job do not create envy. In the following, we provide some scenarios motivating both variants.

For the first variant, consider a company that receives an order of tasks that must be assigned among its m employees. For equity reasons, in order to make the workers satisfied with their task assignment so that they are as productive as they can, the tasks should be assigned in such a way that no envy is induced among the employees.

For the second variant, consider a scenario in which a company, in order to fulfill a complex job composed by several tasks, has to engage a set of employees that, for law or trade union reasons have to be all paid out the same wage. Again, for making the workers as productive as they can, it is required that no envy is induced, but in this case we are interested only in the envy among the *engaged* employees, i.e. the ones receiving at least a task to perform.

We notice that the existence of envy-free schedules is not guaranteed in the first variant of the model. For instance, consider a scenario where the number of machines is strictly greater than the number of jobs. Clearly at least one machine would not get any job and all the machines getting at least one job would be envious. Therefore, in the following of this paper we focus on the second varant of the model, in which envy-freeness is required only among machines getting at least one job.

We adopt a more general definition of envy-free allocations, namely the *k-envy-freeness* (for any $k \geq 1$): Given an assignment and two machines j, j' (where both j and j' get jobs), we say that j k-envies j' if the completion time of j is at least k times the completion time she would have when getting the set of jobs assigned to j'. In other words, an assignment is k-envy-free if no machine would decrease her completion time by a factor at least k by being assigned all the jobs allocated to another machine. Notice that a k-envy-free assignment always exists: a trivial one can be obtained by allocating all the jobs to a single machine, even if it might have a dramatically high makespan.

We are interested in analyzing the loss of performance due to the adoption of envy-free allocations. Our study has an optimistic nature and, then, aims at quantifying the efficiency loss in the best k-envy-free assignment. Therefore, we introduce the **price of k-envy-freeness**, defined as the ratio between the makespan of the best k-envy-free assignment and that of an optimal assignment. In the literature, other papers performed similar optimistic studies, see, for instance, [1,2]. The price of k-envy-freeness represents an ideal limitation to the efficiency achievable by any k-envy-free assignment. In our work, we also show how to efficiently compute k-envy-free assignments which nicely compare with the performance of the best possible ones. We point out that the computation of non-trivial k-envy-free assignments is necessary to achieve good quality solutions, since the ratio between the makespan of the worst k-envy-free assignment and that of an optimal assignment can be very high. In particular, it is unbounded for unrelated machines, $n\frac{s_{max}}{s_{min}}$ for related ones, where s_{max} (resp. s_{min}) is the maximum (resp. minimum) speed among all the machines, and n for identical machines.

We consider the price of k-envy-freeness in the scheduling problem, that is, the ratio between the makespan of the best k-envy-free assignment and that of an optimal assignment. We investigate the cases of unrelated, related and identical machines and provide exact or asymptotically tight bounds on the price of k-envy-freeness. We stress that low values of k implies a greater attitude to envy, which tremendously reduces the set of k-envy-free assignments. A natural threshold that arose in our analysis of the cases with related and identical machines is the value $k = 2$, as it can be appreciated in the following table where we summarize our main results.

		Identical	Related	Unrelated
$k = 1$	UB and LB	$\min\{n, m\}$	$\min\{n, m\}$	$2^{\min\{n,m\}-1}$
$k \in (1, 2)$	UB	$\frac{2k}{k-1}$	$2k\sqrt{\frac{m}{k-1}}$	$\left(1 + \frac{1}{k}\right)^{\min\{n,m\}-1}$
	LB	$\Omega\left(\frac{2k}{k-1}\right)$	$\Omega\left(\sqrt{\frac{m}{k-1}}\right)$	$\left(1 + \frac{1}{k}\right)^{\min\{n,m\}-1}$
$k \geq 2$	UB	$1 + \frac{1}{k}$	$2 + \max\left\{1, \sqrt{\frac{m}{k}}\right\}$	$\left(1 + \frac{1}{k}\right)^{\min\{n,m\}-1}$
	LB	$1 + \frac{1}{k}$	$\max\left\{1, \sqrt{\frac{m}{k}}\right\}$	$\left(1 + \frac{1}{k}\right)^{\min\{n,m\}-1}$

A further result derives from the fact that our upper bound proofs are constructive and, therefore, they de facto provide polynomial time algorithms able to calculate good k-envy-free assignments.

REFERENCES

1. E. Anshelevich, A. Dasgupta, J. M. Kleinberg, E. Tardos, T. Wexler, T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM Journal on Computing*, 38(4): 1602–1623, 2008.
2. I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou. The Efficiency of Fair Division. *Theory of Computing Systems* 50(4): 589–610, 2012.
3. L.E. Dubins, E.H. Spanier. How to cut a cake fairly. *American Mathematical Monthly*, 68:1–17, 1961.
4. D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.