

# Scheduling selfish tasks: about the performance of truthful algorithms

George Christodoulou<sup>1</sup>, Laurent Gourvès<sup>2</sup>, and Fanny Pascual<sup>3</sup>

<sup>1</sup> Max-Planck-Institut für Informatik, Saarbrücken, Germany  
gchristo@mpi-inf.mpg.de

<sup>2</sup> LAMSADE, CNRS UMR 7024, Université de Paris-Dauphine, Paris, France  
laurent.gourves@lamsade.dauphine.fr

<sup>3</sup> Equipe MOAIS (CNRS-INRIA-INPG-UJF), Grenoble, France  
fanny.pascual@imag.fr

**Abstract.** This paper deals with problems which fall into the domain of selfish scheduling: a protocol is in charge of building a schedule for a set of tasks without directly knowing their length. The protocol gets these informations from agents who control the tasks. The aim of each agent is to minimize the completion time of her task while the protocol tries to minimize the maximal completion time. When an agent reports the length of her task, she is aware of what the others bid and also of the protocol's algorithm. Then, an agent can bid a false value in order to optimize her individual objective function. With erroneous information, even the most efficient algorithm may produce unreasonable solutions. An algorithm is truthful if it prevents the selfish agents from lying about the length of their task. The central question in this paper is: *“How efficient a truthful algorithm can be?”* We study the problem of scheduling selfish tasks on parallel identical machines. This question has been raised by Christodoulou et al [8] in a distributed system, but it is also relevant in centrally controlled systems. Without considering side payments, our goal is to give a picture of the performance under the condition of truthfulness.

**Key words:** scheduling, algorithmic game theory, truthful algorithms

## 1 Introduction

The Internet is a complex distributed system involving many autonomous entities (*agents*). Protocols organize this network, using the data held by these agents and trying to maximize the social welfare. Agents are often supposed to be trustworthy but this assumption is unrealistic in some settings as they might try to manipulate the protocol by reporting false information in order to maximize their own profit. With false information, even the most efficient protocol may lead to unreasonable solutions if it is not designed to cope with the selfish behavior of the single entities. Then, it is natural to ask the following question: *How efficient a protocol can be if it guarantees that no agent has incentive to lie?*

In this paper, we deal with the problem of scheduling  $n$  selfish tasks on  $m$  identical parallel machines. We consider two distinct settings in which the aim is to minimize the *makespan*, i.e. the maximum completion time. The first setting is centralized, while the second one is distributed. Both problems share the following characteristics. Each task is owned by an agent<sup>4</sup>. The length  $l_i$  of a task  $i$  is known to its owner only. The agents, considered as players of a non-cooperative game, want to minimize the completion time of their tasks. The protocol builds the schedule with rules known to all players and fixed in advance. In particular, mixing the execution of two jobs (like round-robin) is not allowed. Before the execution begins, the agents report a value representing the length of their tasks. We assume that every agent behaves rationally and selfishly. Each one is aware of the situation the others face and tries to optimize her own objective function. Thus an agent can report a value which is not equal to her real length. Practically, an agent can add “fake” data to artificially increase the length of her task if it decreases her completion time. This selfish behavior can prevent the protocol to produce a reasonable (i.e. close to the social welfare) schedule. Without considering side payments, which are often used with the aim of inciting the agents to report their real value, some algorithmic tools can simultaneously offer a guarantee on the quality of the schedule (its makespan is not arbitrarily far from the optimum) and guarantee that the solution is *truthful* (no agent can lie and improve her own completion time). For both centralized and distributed settings, our goal is to give lower and upper bounds on the performance under the condition of truthfulness. It is important to mention that we do not strictly restrict the study to polynomial time algorithms.

Since the length of a task is private, each agent bids a value which represents the length of her task. We assume that an agent cannot shrink the length of her task (otherwise she will not get her result), but if she can decrease her completion time by bidding a value larger than the real one, then she will do so. We also assume that an agent does not report a distribution on different lengths. A player may play according to a distribution, but she just announces the outcome, so the protocol does not know if she lies.

In the *centralized setting*, the strategy of agent  $i$  is a value  $b_i$  representing the length of her task. The protocol, called an algorithm, is in charge of indicating when and on which machine a task will be scheduled. An algorithm is *truthful* when no agent has incentive to report a false value. We focus on the performance of truthful algorithms with respect to the makespan of the schedule. In particular, we are interested in giving lower and upper bounds on the *approximation ratio* that a (deterministic or randomized) truthful algorithm can achieve. For example, a truthful algorithm can be obtained by greedily scheduling the tasks following the increasing order of their lengths. This algorithm, known as SPT, produces a  $(2 - 1/m)$ -approximate schedule [11]. *Are there truthful algorithms with better approximation guarantee for the considered scheduling problem?*

---

<sup>4</sup> We equally refer to a task and its owner since we assume that two tasks cannot be held by the same agent.

In the *distributed setting*, the strategy of agent  $i$  is a couple  $(M_i, b_i)$ , where  $M_i$  is the machine which will execute the task and  $b_i$  is the length bidden. As opposed to the centralized setting, the agents choose their machine and  $M_i$  can be a probability distribution on different machines. The protocol, called a *coordination mechanism* in this context [8], consists in selecting a *scheduling policy* for each machine (e.g. scheduling the tasks in order of decreasing lengths). An important and natural condition is due to the decentralized nature of the problem: the scheduling on a machine should depend only on the tasks assigned to it, and should be independent of the tasks assigned to the other machines. A coordination mechanism is *truthful* when no agent has incentive to lie on the length of her task. Using the *price of anarchy* [14], we study the performance of truthful coordination mechanisms with respect to the makespan. The price of anarchy of a coordination mechanism is, in the context, equal to the largest ratio between the makespan of a schedule where agent's strategies form a *Nash equilibrium*<sup>5</sup> and the optimal makespan.

Interestingly, it is possible to slightly transform the SPT algorithm in a truthful coordination mechanism, as suggested in [8]: each machine  $P_j$  schedules its tasks in order of increasing lengths, and adds at the very beginning of the schedule a small delay equal to  $(j - 1)\varepsilon$  times the length of the first task. By this way, and if  $\varepsilon$  is small enough, the schedule obtained in a Nash equilibrium is similar to the one returned by the SPT algorithm (excepted the small delays at the beginning of the schedule). When  $\varepsilon$  is negligible, the price of anarchy of this coordination mechanism is  $2 - 1/m$ . *Are there truthful coordination mechanisms with better price of anarchy for the considered scheduling problem?*

For both centralized algorithms and coordination mechanisms, we consider the two following execution models:

- **Strong model of execution:** If the owner of task  $i$  bids  $b_i \geq l_i$ , then the execution time will still be  $l_i$  (i.e. the task will be completed  $l_i$  time units after its start).
- **Weak model of execution:** If the owner of task  $i$  bids  $b_i \geq l_i$ , then the execution time will be  $b_i$  (i.e. the task will be completed  $b_i$  time units after its start).

The strong execution model corresponds to the case where tasks have to be linearly executed – from their beginning to their end –, whereas the weak execution model corresponds to the case where a task can be executed in any order<sup>6</sup> (and the “fake” part of the task is not anymore necessarily executed at the end), or when the machine returns the result of the task only at the end of its execution. Depending on the applications of the scheduling problem, either the strong or the weak model of execution will be used.

<sup>5</sup> Situation in which no agent can unilaterally change her strategy and improve her own completion time. A Nash equilibrium is *pure* if each agent has a pure strategy : each agent chooses only one machine. A Nash equilibrium is *mixed* if the agents give a probability distribution on the machines on which they will go.

<sup>6</sup> Nevertheless, the execution of two jobs is never interlaced.

## Related work

The field of *Mechanism Design* can be useful to deal with the selfishness of the agents. Its main idea is to pay the agents to convince them to perform strategies that help the system to optimize a global objective function. The most famous technique for designing truthful mechanisms is perhaps the Vickrey-Clarke-Groves (VCG) mechanism [20, 7, 12]. However, when applied to combinatorial optimization problems, this mechanism guarantees the truthfulness under the hypothesis that the objective function is *utilitarian* (i.e. the value of the objective function is equal to the sum of the agents individual objective functions) and that the mechanism is able to compute the optimum. Archer and Tardos introduce in [4] a method which allows to design truthful mechanisms for several combinatorial optimization problems to which the VCG mechanism does not apply. However, both approaches cannot be applied to our problem.

Scheduling selfish agents has been intensively studied these last years, started with the seminal work of Nisan and Ronen [17], and followed by a series of papers [1, 2, 4, 6, 9, 15, 16]. However, all these works differ from ours since in their case, the selfish agents are the machines while here we consider that the agents are the tasks. Furthermore, they use side payments whereas we focus on truthful algorithms without side payments.

A more closely related work is the one of Christodoulou et al [8] who considered the same model but only in the distributed context of coordination mechanisms. They proposed different coordination mechanisms with a price of anarchy better than the one of the SPT coordination mechanism. Nevertheless, these mechanisms are not truthful. In [13], the authors gave coordination mechanisms for the same model for related machines (i.e. machines can have different speeds), but their mechanisms are also not truthful.

In [3], the authors gave a truthful randomized algorithm for the strong model of execution defined before, and they gave, for the weak model of execution, a coordination mechanism which is truthful if there are two machines and if the lengths of the tasks are powers of a certain constant. An optimal (but exponential time) truthful randomized algorithm and a truthful randomized PTAS for the weak model of execution appear in [18, 19]. The technique consists in computing an optimal (resp. a  $(1 + \epsilon)$ -approximate) schedule and each machine executes its tasks in a random order (the truthfulness is due to the introduction of fictitious tasks which guarantee that all the machines have the same load).

Another related work is the one of Auletta et al. who considered in [5] the problem of scheduling selfish tasks in a centralized case. Their work differs from ours since they considered that each machine uses a round and robin policy and thus that the completion of each task is the completion time of the machine on which the task is (this model is known as the KP model). They considered that the tasks can lie in both directions, and that there are some payments.

## Contribution and organization of the article

Sections 3 and 4 are devoted to the centralized setting. In particular, we study the strong (resp. weak) model of execution in Section 3 (resp. Section 4).

Results on the distributed setting are presented in Section 5 for both execution models.

Table 1 gives a summary of the bounds that we are aware of (those with a † are presented in this article). LB stands for “Lower bound”, UB for “Upper bound” and NE for “Nash equilibria”. Due to the lack of space, some proofs are deferred to the Appendix.

Strong model of execution:

	Deterministic		Randomized	
	LB	UB	LB	UB
centralized setting	$2 - \frac{1}{m}$ †	$2 - \frac{1}{m}$ [8]	$\frac{3}{2} - \frac{1}{2m}$ †	$2 - \frac{1}{m+1} \left( \frac{5}{3} + \frac{1}{3m} \right)$ [3]
distributed setting	$2 - \frac{1}{m}$ (pure NE) † $\frac{3}{2} - \frac{1}{2m}$ (mixed NE) †	$2 - \frac{1}{m}$ [8]	$\frac{3}{2} - \frac{1}{2m}$ †	$2 - \frac{1}{m}$

Weak model of execution:

	Deterministic		Randomized	
	LB	UB	LB	UB
centralized setting	$m = 2 : 1 + \frac{\sqrt{105-9}}{12} > 1.1$ † $m \geq 3 : \frac{7}{6} > 1.16$ †	$\frac{4}{3} - \frac{1}{3m}$ †	1 [18, 19]	1 [18, 19]
distributed setting	$\frac{1+\sqrt{17}}{4} > 1.28$ (pure NE) †	$2 - \frac{1}{m}$	$1 + \frac{\sqrt{13-3}}{4} > 1.15$ † (pure NE)	$2 - \frac{1}{m}$

**Table 1.** Bounds for  $m$  identical machines.

## 2 Notations

We are given  $m$  machines (or processors)  $\{P_1, \dots, P_m\}$ , and  $n$  tasks  $\{1, \dots, n\}$ . Let  $l_i$  denote the real execution time (or length) of task  $i$ . We use the identification numbers to compare tasks of the same (bidding) lengths: we will say that task  $i$ , which bids  $b_i$ , is larger than task  $j$ , which bids  $b_j$ , if and only if  $b_i > b_j$  or ( $b_i = b_j$  and  $i > j$ ). It is important to mention that an agent cannot lie on her (unique) identification number.

A randomized algorithm can be seen as a probability distribution over deterministic algorithms. We say that a (randomized) algorithm is truthful if for every task the expected completion time when she declares her true length is smaller than or equal to her expected completion time in the case where she declares a larger value. More formally, we say that an algorithm is *truthful* if  $E_i[l_i] \leq E_i[b_i]$ , for every  $i$  and  $b_i \geq l_i$ , where  $E_i[b_i]$  is the expected completion time of task  $T_i$  if she declares  $b_i$ . In order to evaluate the quality of a randomized algorithm, we use the notion of expected approximation ratio.

We will refer in the sequel to the list scheduling algorithms LPT and SPT, where LPT (resp. SPT) [11] is the algorithm which greedily schedules the tasks,

sorted in order of decreasing (resp. increasing) lengths: this algorithm schedules, as soon as a machine is available, the largest (resp. smallest) task which has not yet been scheduled. An LPT (resp. SPT) schedule is a schedule returned by the LPT (resp. SPT) algorithm.

### 3 About truthful algorithms for the strong model of execution

#### 3.1 Deterministic algorithms

We saw that the deterministic algorithm SPT, which is  $(2 - \frac{1}{m})$ -approximate, is truthful. Let us now show that there is no truthful deterministic algorithm with a better approximation ratio.

**Theorem 1.** *Let us consider that we have  $m$  identical machines. There is no truthful deterministic algorithm with an approximation ratio smaller than  $2 - \frac{1}{m}$ .*

*Proof.* Let us suppose that we have  $n = m(m - 1) + 1$  tasks of length 1. Let us suppose that we have a truthful deterministic algorithm  $\mathcal{A}$  which has an approximation ratio smaller than  $(2 - 1/m)$ . Let  $t$  be the task which has the maximum completion time,  $C_t$ , in the schedule returned by  $\mathcal{A}$ . We know that  $C_t \geq m$ .

Let us now suppose that task  $t$  bids  $m$  instead of 1. We will show that the completion time of  $t$  is then smaller than  $m$ . Let  $OPT$  be the makespan of an optimal solution where there are  $n - 1 = m(m - 1)$  tasks of length 1 and a task of length  $m$ . We have:  $OPT = m$ . Since the approximation ratio of algorithm  $\mathcal{A}$  is smaller than  $(2 - 1/m)$ , the makespan of the schedule it builds with this instance is smaller than  $(2 - 1/m)m = 2m - 1$ . Thus, the task of length  $m$  starts before time  $(m - 1)$ . Thus, if task  $t$  bids  $m$  instead of 1, it will start before time  $m - 1$  and be completed one time unit after, that is before time  $m$ . Thus task  $t$  will decrease its completion time by bidding  $m$  instead of 1, and algorithm  $\mathcal{A}$  is not truthful.

Note that we can generalize this result to the case of related machines: we have  $m$  machines  $P_1, \dots, P_m$ , such that machine  $P_i$  has a speed  $v_i$ ,  $v_1 = 1$ , and  $v_1 \leq \dots \leq v_m$ . By this way, the bound becomes  $2 - \frac{v_m}{\sum_{i=1}^m v_i}$ .

Concerning the strong model of execution, no deterministic algorithm can outperform SPT in the centralized setting. Then, it is interesting to consider randomized algorithms to achieve a better approximation ratio.

#### 3.2 Randomized algorithms

In [3], the authors present a randomized algorithm which consists in returning a LPT schedule with a probability  $1/(m + 1)$  and a slightly modified SPT schedule with a probability  $m/(m + 1)$ . They obtain a truthful algorithm whose expected approximation ratio improves  $2 - \frac{1}{m}$  but no instance showing the tightness of

their analysis is provided. A good candidate should be simultaneously a tight example for both LPT and SPT schedules. We are not aware of the existence of such an instance and we believe in a future improvement of this upper bound. The following Theorem provides a lower bound.

**Theorem 2.** *Let us consider that we have  $m$  identical machines. There is no truthful randomized algorithm with an approximation ratio smaller than  $\frac{3}{2} - \frac{1}{2m}$ .*

Generalizing this result to the case of related machines, the bound becomes  $\frac{3}{2} - \frac{v_m}{2 \sum_{i=1}^m v_i}$ .

## 4 About truthful algorithms for the weak model of execution

### 4.1 A truthful deterministic algorithm

We saw in the Section 3 that SPT is a truthful and  $(2 - 1/m)$ -approximate algorithm for the strong model of execution, and that no truthful deterministic algorithm can have a better approximation ratio. If we consider the weak model of execution, we can design a truthful deterministic algorithm, called  $LPT_{mirror}$ , with a better performance guarantee. We are given  $n$  tasks  $\{1, \dots, n\}$  which bid lengths  $b_1, \dots, b_n$ . Make a schedule  $\sigma_{LPT}$  with the LPT list algorithm. Let  $C_{max}^{OPT}$  be the optimal makespan. Let  $p(i)$  be the machine on which the task  $i$  is executed in  $\sigma_{LPT}$ . Let  $C_i$  be date at which the task  $i$  ends in  $\sigma_{LPT}$ .  $LPT_{mirror}$  returns the schedule in which task  $i$  is executed on machine  $p(i)$  and starts at time  $(4/3 - 1/(3m))C_{max}^{OPT} - C_i$ .

**Theorem 3.**  *$LPT_{mirror}$  is a deterministic, truthful and  $(\frac{4}{3} - \frac{1}{3m})$ -approximate algorithm.*

*Proof.* We are given  $n$  tasks with true lengths  $l_1, \dots, l_n$ . Let us suppose than each task has bidden a value, and that task  $i$  bids  $b_i > l_i$ . This can make the task  $i$  start earlier in  $\sigma_{LPT}$  but never later. In addition, the optimal makespan when  $i$  bids  $b_i > l_i$  is necessarily larger than or equal to the optimal makespan when task  $i$  reports its true length. Let  $S_i$  be the date at which task  $i$  starts to be executed in  $\sigma_{LPT}$ . The completion time of task  $i$  in  $LPT_{mirror}$  is  $(4/3 - 1/(3m))OPT - C_i + b_i = (4/3 - 1/(3m))OPT - S_i$  because  $S_i = C_i - b_i$ . By bidding  $b_i > l_i$ , task  $i$  can only increase its completion time in the schedule returned by  $LPT_{mirror}$  because  $OPT$  does not decrease and  $S_i$  does not increase. Thus task  $i$  does not have incentive to lie.

Since the approximation ratio of the schedule obtained with the LPT list algorithm is at most  $(4/3 - 1/(3m))$  [11], the schedule returned by  $LPT_{mirror}$  is clearly feasible and its makespan is, by construction,  $(4/3 - 1/(3m))$ -approximate. Thus  $LPT_{mirror}$  is a truthful and  $(\frac{4}{3} - \frac{1}{3m})$ -approximate algorithm.

Note that  $LPT_{mirror}$  is not a polynomial time algorithm, since we need to know the value of the makespan in an optimal solution, which is an NP-hard problem [10]. However, it is possible to have a polynomial time algorithm which is  $(4/3 - 1/(3m))$ -approximate, even if some tasks do not bid their true values. Consider the following simple algorithm: we first compute a schedule  $\sigma_{LPT}$  with the LPT algorithm. Let  $p(i)$  be the machine on which the task  $i$  is executed in  $\sigma_{LPT}$ , let  $C_i$  be the completion time of task  $i$  in  $\sigma_{LPT}$ , and let  $C_{max}$  be the makespan of  $\sigma_{LPT}$ . We then compute the final schedule  $\sigma'$  in which task  $i$  is scheduled on  $p(i)$  and starts at time  $C_{max} - C_i$ .

We can show that this algorithm is  $(4/3 - 1/(3m))$ -approximate (i.e. the schedule returned by this algorithm is at most  $(4/3 - 1/(3m))$  times larger than the optimal schedule in which all the tasks bid their true values). We can show this by the following way. We suppose that all the tasks except  $i$  have bidden some values. Let  $\sigma_{LPT}(b_i)$  be the schedule  $\sigma_{LPT}$  obtained when  $i$  bids  $b_i$ , let  $S_i(b_i)$  be the date at which task  $i$  starts to be executed in  $\sigma_{LPT}(b_i)$ , and let  $C_{max}(\sigma_{LPT}(b_i))$  be the makespan of  $\sigma_{LPT}(b_i)$ . The completion time of task  $i$  (which bids  $b_i$ ) in  $\sigma'$  is equal to  $C_{max}(\sigma_{LPT}(b_i)) - S_i(b_i)$ . Since with the LPT algorithm, tasks are scheduled in decreasing order of lengths, if  $b_i > l_i$  then  $S_i(b_i) \leq S_i(l_i)$ . Thus, whatever the values bidden by the other tasks are,  $i$  has incentive to lie and bid  $b_i > l_i$  only if  $C_{max}(\sigma_{LPT}(b_i)) < C_{max}(\sigma_{LPT}(l_i))$ . Since this is true for each task, no task will unilaterally lie unless this decreases the makespan of the schedule. The makespan of the schedule  $\sigma'$  in which all the tasks bid their true values is  $(4/3 - 1/(3m))$ -approximate, and then the solution returned by this algorithm will also be  $(4/3 - 1/(3m))$ -approximate.

## 4.2 Deterministic algorithms: lower bounds

We suppose that the solution returned by an algorithm depends on the length and the identification number of each task, even those which can be identified with their unique length.

**Theorem 4.** *Let us consider that we have two identical machines. There is no truthful deterministic algorithm with an approximation ratio smaller than  $1 + (\sqrt{105} - 9)/12 \approx 1.1039$ .*

**Theorem 5.** *Let us consider that we have  $m \geq 3$  identical machines. There is no truthful deterministic algorithm with an approximation ratio smaller than  $7/6$ .*

The assumption made to derive Theorems 4 and 5 is, in a sense, stronger than the usual one since we suppose that the solution returned by an algorithm for two similar instances (same number of tasks, same lengths but different identification numbers) can be completely different. If we relax this assumption, i.e. if identification numbers are only required for the tasks which have the same length, the bound presented in Theorem 4 can be improved to  $7/6$ .

**Theorem 6.** *Let us consider that we have two identical machines. No truthful deterministic algorithm can be better than 7/6-approximate if it does not take into account the identification number of tasks whose length is unique.*

## 5 About truthful coordination mechanisms

Let  $\rho \geq 1$ . If there is no truthful deterministic algorithm which has an approximation ratio of  $\rho$ , then there is no truthful deterministic coordination mechanism which always induce pure Nash equilibria and which has a price of anarchy smaller than or equal to  $\rho$ . Indeed, if this was not the case, then the deterministic algorithm which consists in building the schedule obtained in a pure Nash equilibrium with this  $\rho$ -approximate coordination mechanism would be a  $\rho$ -approximate truthful deterministic algorithm.

Likewise, if there is no truthful (randomized) algorithm which has an approximation ratio of  $\rho$ , then there is no truthful coordination mechanism which has a price of anarchy smaller than or equal to  $\rho$ . Indeed, if this was not the case, the algorithm which consists in building the schedule obtained in a Nash equilibrium with this  $\rho$ -approximate coordination mechanism would be a  $\rho$ -approximate truthful algorithm.

This observation leads us to the following results for the strong model of execution. We deduce from Theorem 1 that there is no truthful deterministic coordination mechanism which always induce pure Nash equilibria and which has a price of anarchy smaller than  $2 - 1/m$ . Thus there is no truthful coordination mechanism which performs better than the truthful SPT coordination mechanism, whose price of anarchy tends towards  $2 - 1/m$ . We deduce from Theorem 2 that there is no truthful coordination mechanism which has a price of anarchy smaller than  $\frac{3}{2} - \frac{1}{2m}$ . We now consider the weak model of execution.

**Theorem 7.** *If we consider the weak model of execution, there is no truthful deterministic coordination mechanism which induces pure Nash equilibria, and which has a price of anarchy smaller than  $\frac{1+\sqrt{17}}{4} \approx 1.28$ .*

*Proof.* Let us first prove this result in the case where there are two machines,  $P_1$ , and  $P_2$ . Let  $\varepsilon > 0$ . Let us suppose that there exists a truthful coordination mechanism  $\mathcal{M}$  with a price of anarchy of  $\frac{1+\sqrt{17}}{4} - \varepsilon$ . Let us consider the following instance  $I_1$ : three tasks of length 1. Since  $\mathcal{M}$  is a deterministic coordination mechanism which induces pure Nash equilibria, there is at least a task in  $I_1$  which has a completion time larger than or equal to 2. Let  $t$  be such a task.

Let us first consider this instance  $I_2$ : we have two tasks of length  $\frac{-1+\sqrt{17}}{2} \approx 1.56$ . Since  $\mathcal{M}$  is  $(\frac{1+\sqrt{17}}{4} - \varepsilon)$ -approximate, there is one task on each machine, and each task is completed before time  $\frac{-1+\sqrt{17}}{2} \times \frac{1+\sqrt{17}}{4} = 2$ . Thus, when it has a task of length  $\frac{-1+\sqrt{17}}{2}$ , each machine must end it before time 2.

Let us now consider the following instance  $I_3$ : two tasks of length 1, and a task of length  $\frac{-1+\sqrt{17}}{2}$ . Since  $\mathcal{M}$  is  $(\frac{1+\sqrt{17}}{4} - \varepsilon)$ -approximate, the task of length

$\frac{-1+\sqrt{17}}{2}$  is necessarily alone on its machine (without loss of generality, on  $P_2$ ). As we have seen it,  $P_2$  must schedule this task before time 2. Thus, task  $t$  of instance  $I_1$ , has incentive to bid  $\frac{-1+\sqrt{17}}{2}$  instead of 1: by this way it will end before time 2, instead of a time larger than or equal to 2.

We can easily extend this proof in the case where there are more than 2 machines, by having  $m + 1$  tasks of length 1 in  $I_1$ ;  $m$  tasks of length  $\frac{-1+\sqrt{17}}{2}$  in  $I_2$ ; and  $m$  tasks of length 1 and a task of length  $\frac{-1+\sqrt{17}}{2}$  in  $I_3$ .

**Theorem 8.** *If we consider the weak model of execution, there is no truthful coordination mechanism which induces pure Nash equilibria, and which has a price of anarchy smaller than  $1 + \frac{\sqrt{13}-3}{4} \approx 1.15$ .*

## 6 Conclusion

We showed that, in the strong model of execution, the list algorithm SPT, which has an approximation ratio of  $2 - 1/m$  is the best truthful deterministic algorithm, and that there is no truthful randomized algorithm which has an approximation ratio smaller than  $3/2 - 1/(2m)$ . On the contrary, if we relax the constraints on the execution model, i.e. if the result of a task which bid  $b$  is given to this task only  $b$  time units after its start, then we can obtain better results. In this model of execution, there is a truthful  $4/3 - 1/(3m)$ -approximate deterministic algorithm and a truthful optimal randomized algorithm. For both execution models, we also gave lower bounds on the approximation ratios that a truthful coordination mechanism can have.

As a future work, it would be interesting to improve the results for which a gap between the lower and the upper bound exists. For example, we believe that the lower bound  $\frac{1+\sqrt{17}}{4}$  (lower bound on the performance of a truthful deterministic coordination mechanism for the weak model of execution) can be improved to  $3/2$  for two machines.

Another direction would be to restrict the study to truthful algorithms (or coordination mechanisms) which run in polynomial time. Giving improved lower bounds which rely on a computational complexity argument would be very interesting.

**Acknowledgments.** We thank Elias Koutsoupias for helpful suggestions and discussions on the problem.

## References

1. P. Ambrosio, V. Auletta. *Deterministic Monotone Algorithms for Scheduling on related Machines*. In Proc. of WAOA 2004, 267-280.
2. N. Andelman, Y. Azar, M. Sorani. *Truthful Approximation Mechanisms for Scheduling Selfish Related Machines*. In Proc. of STACS 2005, 69-82.

3. E. Angel, E. Bampis and F. Pascual. *Truthful Algorithms for Scheduling Selfish Tasks on Parallel Machines*. In Proc. of WINE 2005, LNCS 3828, pp. 698-707, 2005.
4. A. Archer, E. Tardos. *Truthful Mechanisms for One-Parameter Agents*. In Proc. of FOCS 2001, 482-491, 2001.
5. V. Auletta, P. Penna, R. De Prisco, P. Persiano. *How to Route and Tax Selfish Unsplittable Traffic*. In Proc. of SPAA 2004, 196-204, 2004.
6. V. Auletta, R. De Prisco, P. Penna, P. Persiano. *Deterministic Truthful Approximation Mechanisms for Scheduling Related Machines*. In Proc. of STACS 2004, 608-619, 2004.
7. E. Clarke. *Multipart pricing of public goods*. Public Choices, pages 17-33, 1971.
8. G. Christodoulou, E. Koutsoupias and A. Nanavati. *Coordination Mechanisms*. In Proc. of ICALP 2004, LNCS 3142, pp. 345-357, 2004.
9. G. Christodoulou, E. Koutsoupias and A. Vidali. *A lower bound for scheduling mechanisms*. In Proc. of SODA 2007, 2007.
10. M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co, 1979.
11. R. Graham. *Bounds on multiprocessor timing anomalies*. In SIAM Jr. on Appl. Math. 17(2), pp. 416-429, 1969.
12. T. Groves. *Incentive in teams*. Econometrica, 41(4):617-631, 1973.
13. N. Immorlica, L. Li and V.S. Mirrokni and A. Schulz. *Coordination Mechanisms for Selfish Scheduling*. In Proc. of WINE 2005, LNCS 3828, pp. 55-69, 2005.
14. E. Koutsoupias and C. Papadimitriou. *Worst Case Equilibria*. In Proc. of STACS 1999, LNCS 1563, pp. 404-413, 1999.
15. A. Kovács. *Fast monotone 3-approximation algorithm for scheduling related machines*. In Proc. of ESA 2005, LNCS 3669, pp. 616-627, 2005.
16. A. Mu'alem and M. Schapira. *Setting lower bounds on truthfulness*. In Proc. of SODA 2007, 2007.
17. N. Nisan, A. Ronen. *Algorithmic mechanism design*. In Proc. STOC 1999, 129-140, 1999.
18. F. Pascual. *Optimisation dans les réseaux : de l'approximation polynomiale à la théorie des jeux*, Ph. D Thesis, University of Evry, France, 2006 (in french).
19. A-A. Tchegnia. *Truthful algorithms for some scheduling problems*, Master Thesis MPRI, École Polytechnique, France, 2006.
20. W. Vickrey. *Counterspeculation, auctions and competitive sealed tenders*. J. Finance, 16:8-37, 1961.

## Appendix

### Proof of Theorem 2:

*Proof.* Let  $\varepsilon > 0$ , and let us suppose that we have a truthful algorithm  $\mathcal{A}$  whose expected approximation ratio is  $(3/2 - 1/(2m) - \varepsilon)$ .

Let us consider that we have  $xm(m-1) + m$  tasks of length 1, where  $x$  is a positive integer such that  $x > 1/(2\varepsilon m) - 1/(2\varepsilon m^2) - 1/m$ . Since  $m$  machines are available, the optimal makespan is  $x(m-1) + 1$ . With any randomized algorithm (including  $\mathcal{A}$ ), there is a task  $t$  whose expected completion time is at

least  $(x(m-1))/2 + 1$ . Since algorithm  $\mathcal{A}$  is truthful,  $t$  should not improve its completion time by bidding  $xm + 1$  instead of 1. Suppose that  $t$  unilaterally lies and bids  $xm + 1$  instead of 1. For algorithm  $\mathcal{A}$ , there is then  $xm(m-1) + m - 1$  tasks of length 1 and a task of length  $xm + 1$ . The optimal makespan is then  $xm + 1$  and the expected makespan of the schedule returned by  $\mathcal{A}$  is smaller than or equal to  $(3/2 - 1/(2m) - \varepsilon)(xm + 1)$ . As a consequence, the expected completion time of the task of length  $xm + 1$  is also bounded above by  $(3/2 - 1/(2m) - \varepsilon)(xm + 1)$ . Since  $t$  increased its length with  $xm$  time units, its real expected completion time is smaller than or equal to  $(3/2 - 1/(2m) - \varepsilon)(xm + 1) - xm$ . With  $x > 1/(2\varepsilon m) - 1/(2\varepsilon m^2) - 1/m$ , the expected completion time of  $t$  is strictly smaller than  $(x(m-1))/2 + 1$  when it bids  $xm + 1$  and larger than or equal to  $(x(m-1) + 2)/2$  when it bids its true length 1. This contradicts the fact that  $\mathcal{A}$  is truthful.

### • Results of Section 3 for related machines

**Theorem 9.** *Let us consider that we have a fixed number  $m$  of machines  $P_1, \dots, P_m$ , such that machine  $P_i$  has a speed  $v_i$ ,  $v_1 = 1$ , and  $v_1 \leq \dots \leq v_m$ . There is no truthful deterministic algorithm with an approximation ratio smaller than  $2 - \frac{v_m}{\sum_{i=1}^m v_i}$ .*

*Proof.* Let us suppose that we have  $n$  tasks of length 1, and that  $n \gg \sum_{i=1}^m v_i$ . Let us suppose that we have a truthful algorithm  $\mathcal{A}$  which has an approximation ratio equal to  $\rho$  (we will then show that  $\rho$  has to be larger than or equal to  $2 - \frac{v_m}{\sum_{i=1}^m v_i}$ ). Let  $t$  be the task which has the maximum completion time,  $C_t$ , in the schedule returned by  $\mathcal{A}$ . We know that  $C_t \geq \frac{n}{\sum_{i=1}^m v_i}$ . Indeed, if it is possible to assign  $\frac{nv_j}{\sum_{i=1}^m v_i}$  tasks to each machine  $P_j$ , with  $j \in \{1, \dots, m\}$ , then every machines will end at time  $\frac{n}{\sum_{i=1}^m v_i}$ . If, for some  $j \in \{1, \dots, m\}$ , it is not possible to assign  $\frac{nv_j}{\sum_{i=1}^m v_i}$  tasks to each machine  $P_j$  (i.e.  $\frac{nv_j}{\sum_{i=1}^m v_i}$  is not an integer number), then there is a machine  $P_j$  on which there are more than  $\frac{nv_j}{\sum_{i=1}^m v_i}$  tasks, and then on which the last task is completed after  $\frac{n}{\sum_{i=1}^m v_i}$  time units.

Let us now suppose that task  $t$  bids  $l'_t = \frac{(n-1)v_m}{\sum_{i=1}^{m-1} v_i}$  instead of 1. Let  $OPT$  be the makespan of an optimal solution where there are  $n-1$  tasks of length 1 and a task of length  $l'_t$ . We have:  $\frac{(n-1)}{\sum_{i=1}^{m-1} v_i} \leq OPT < \frac{(n-1)}{\sum_{i=1}^{m-1} v_i} + 1$ . Indeed,  $\frac{(n-1)}{\sum_{i=1}^{m-1} v_i}$  is the date at which all the machines end at the same time, if this is possible: in this case the task of length  $l'_t$  is on  $P_m$  and there are  $\frac{(n-1)v_j}{\sum_{i=1}^{m-1} v_i}$  tasks of length 1 on machine  $P_j$  (for every  $j \in \{1, \dots, m-1\}$ ). If it is not possible that all the machines end at the same time, then we schedule (at most)  $\lceil \frac{(n-1)v_j}{\sum_{i=1}^{m-1} v_i} \rceil$  tasks of length 1 on machine  $P_j$  ( $j \in \{1, \dots, m-1\}$ ): all the tasks are scheduled and the completion time on each machine  $P_i$  is smaller than  $\frac{(n-1)}{\sum_{i=1}^{m-1} v_i} + \frac{1}{v_i} \leq \frac{(n-1)}{\sum_{i=1}^{m-1} v_i} + 1$ .

Algorithm  $\mathcal{A}$  is not truthful if the completion time of  $t$  which bids  $l'_t$  is smaller than  $\frac{n}{\sum_{i=1}^m v_i}$ , since the completion time of  $t$  when it bids 1 is at least  $\frac{n}{\sum_{i=1}^m v_i}$ . Let  $S_A$  be the schedule returned by algorithm  $\mathcal{A}$  when it has  $(n-1)$  tasks of

length 1 and a task of length  $l'_t$ . Algorithm  $\mathcal{A}$  is not truthful if at least one unit of the task of length  $l'_t$  is completed before  $\frac{n}{\sum_{i=1}^m v_i}$  time units in  $S_A$ . Thus, if we want  $\mathcal{A}$  to be truthful, then less than one unit of the task of length  $l'_t$  has to be completed before  $\frac{n}{\sum_{i=1}^m v_i}$  time units in  $S_A$ . In this case, the makespan of  $S_A$  is at least  $\frac{n}{\sum_{i=1}^m v_i} + (\frac{n-1}{\sum_{i=1}^{m-1} v_i} - \frac{1}{v_m})$ , where  $\frac{n-1}{\sum_{i=1}^{m-1} v_i} - \frac{1}{v_m}$  is the smallest execution time to execute a task of length  $l'_t - 1$ . Thus, if  $\mathcal{A}$  is truthful, the makespan of  $S_A$  is at least  $\frac{n}{\sum_{i=1}^m v_i} + (\frac{n-1}{\sum_{i=1}^{m-1} v_i} - \frac{1}{v_m}) \geq (\frac{\sum_{i=1}^{m-1} v_i}{n-1+\sum_{i=1}^{m-1} v_i}) (\frac{n}{\sum_{i=1}^m v_i} + \frac{n-1}{\sum_{i=1}^{m-1} v_i} - \frac{1}{v_m}) OPT$ , because we have seen that  $OPT \leq \frac{(n-1)}{\sum_{i=1}^{m-1} v_i} + 1$ . This last expression is equal to  $\frac{n-1}{n-1+\sum_{i=1}^{m-1} v_i} + (\frac{n}{n-1+\sum_{i=1}^{m-1} v_i}) (\frac{\sum_{i=1}^{m-1} v_i}{\sum_{i=1}^m v_i} - \frac{\sum_{i=1}^{m-1} v_i}{v_m(n-1+\sum_{i=1}^{m-1} v_i)})$ , and tend towards  $1 + \frac{\sum_{i=1}^{m-1} v_i}{\sum_{i=1}^m v_i} = 2 - \frac{v_m}{\sum_{i=1}^m v_i}$ , when  $n \gg \sum_{i=1}^m v_i$ .

Thus, if  $\rho < 2 - \frac{v_m}{\sum_{i=1}^m v_i}$ , there is an instance of  $n$  tasks of length 1 in which one of the tasks of length 1 will have incentive to bid  $l'_t$  instead of its true value. Thus, there is no truthful deterministic algorithm with an approximation ratio smaller than  $2 - \frac{v_m}{\sum_{i=1}^m v_i}$ .

Theorem 1 is thus a corollary of this theorem when all the machines have the same speed.

We can adapt in the same way the proof of Theorem 2, to show the following Theorem for related machines:

**Theorem 10.** *Let us consider that we have a fixed number  $m$  of machines  $P_1, \dots, P_m$ , such that machine  $P_i$  has a speed  $v_i$ ,  $v_1 = 1$ , and  $v_1 \leq \dots \leq v_m$ . There is no truthful randomized algorithm with an approximation ratio smaller than  $\frac{3}{2} - \frac{v_m}{2 \sum_{i=1}^m v_i}$ .*

*Proof.* Let us consider an instance  $I$  with  $n$  tasks of length 1 ( $n \gg \sum_{i=1}^m v_i$ ). Let  $O = \frac{n}{\sum_{i=1}^m v_i}$ . Let  $OPT$  be the makespan of an optimal solution of this instance. We have  $O - 1 < OPT \leq O$ . In any schedule of these tasks, there is at least one task, denoted by  $t$ , whose expected completion time is larger than or equal to  $\frac{O-1}{2}$ . Indeed, in any schedule, the number of tasks which are completed at time  $(O-1)/2$  is smaller than or equal to the total number of tasks, otherwise  $OPT > O - 1$  would not be the makespan of an optimal solution.

Let us now consider the instance  $I'$  with  $n-1$  tasks of length 1 and one task of length  $\frac{(n-1)v_m}{\sum_{i=1}^{m-1} v_i}$ . Let  $O' = \frac{n-1}{\sum_{i=1}^{m-1} v_i}$ . Let  $OPT'$  be the optimal makespan of  $I'$ . We have  $O' \leq OPT' < O' + 1$ . Indeed, if the large task is on machine  $P_m$  of speed  $v_m$ , and at most  $\lceil \frac{(n-1)v_j}{\sum_{i=1}^{m-1} v_i} \rceil$  tasks of length 1 are on machine  $P_j$  of speed  $v_j$  (for all  $j \in \{1, \dots, m-1\}$ ), then all the tasks are scheduled and the makespan is smaller than or equal to  $\lceil \frac{n-1}{\sum_{i=1}^{m-1} v_i} \rceil \leq O' + 1$ . Moreover the minimum completion time of the large task is equal to  $O'$ .

Let us now consider that, in  $I$ , task  $t$  bids  $\frac{(n-1)v_m}{\sum_{i=1}^{m-1} v_i}$  instead of 1. Let us consider that we have an algorithm  $\mathcal{A}$  which has an approximation ratio  $\rho$ . We

will see that, if  $\rho < \frac{3}{2} - \frac{v_m}{\sum_{i=1}^m v_i}$ , then  $t$  decreases its expected completion time by lying, and thus that  $\mathcal{A}$  is not truthful.

The expected completion time of  $t$  which bids  $\frac{(n-1)v_m}{\sum_{i=1}^m v_i}$  instead of 1 is smaller than or equal to

$$\rho OPT' - \left( \frac{(n-1)v_m}{\sum_{i=1}^{m-1} v_i} - 1 \right) / v_j$$

where  $P_j$  is the machine on which  $t$  is assigned by  $\mathcal{A}$ . Since  $OPT' < O' + 1$  and  $v_j \leq v_m$ , this is smaller than:

$$\rho \left( \frac{n-1}{\sum_{i=1}^{m-1} v_i} + 1 \right) - \frac{n-1}{\sum_{i=1}^{m-1} v_i} + \frac{1}{v_m}.$$

If  $t$  does not lie, we have seen that its expected completion time is larger than or equal to  $(\frac{n}{\sum_{i=1}^m v_i} - 1)/2$ . Then,  $\mathcal{A}$  is truthful if

$$\rho \left( \frac{n-1}{\sum_{i=1}^{m-1} v_i} + 1 \right) - \frac{n-1}{\sum_{i=1}^{m-1} v_i} + \frac{1}{v_m} < \left( \frac{n}{\sum_{i=1}^m v_i} - 1 \right) / 2$$

which is equivalent to

$$\rho < \frac{\sum_{i=1}^{m-1} v_i}{n-1 + \sum_{i=1}^{m-1} v_i} \left( \frac{n}{2 \sum_{i=1}^m v_i} - \frac{1}{2} + \frac{n-1}{\sum_{i=1}^{m-1} v_i} - \frac{1}{v_m} \right).$$

When  $n$  tends towards the infinity, the right hand part of this inequality tends towards:

$$\frac{\sum_{i=1}^{m-1} v_i}{2 \sum_{i=1}^m v_i} + 1 = \frac{1}{2} \left( \frac{\sum_{i=1}^m v_i - v_m}{\sum_{i=1}^m v_i} \right) + 1 = \frac{3}{2} - \frac{v_m}{\sum_{i=1}^m v_i}.$$

Thus, if  $\rho$  is smaller than  $\frac{3}{2} - \frac{v_m}{\sum_{i=1}^m v_i}$ , then  $\mathcal{A}$  is not truthful.

Theorem 2 is thus a corollary of this theorem when all the machines have the same speed.

## • Results of Section 4.2

### Proof of Theorem 4 :

*Proof.* For the sake of simplicity, we first prove this theorem with a ratio 1.1. At the end of the proof, we simply replace some numerical values to get the bound  $1 + (\sqrt{105} - 9)/12$ .

Let us suppose that we have a truthful algorithm  $\mathcal{A}$  with an approximation ratio  $\rho < 1.1$ . Let  $I$  be the following instance: one task of length 5, one task of length 4 and three tasks of length 3. The optimal makespan is 9. When  $I$  is the input,  $\mathcal{A}$  returns a schedule  $\sigma$  whose makespan is (strictly) smaller than 9.9.

Since  $\mathcal{A}$  is a deterministic algorithm, it must execute the three tasks of length 3 on the same machine. Then, two of them have a completion time larger than or equal to 6. Let  $I'$  be the following instance: two tasks of length 5, one task of length 4 and two tasks of length 3. The optimal makespan is 10. When  $I'$  is the input,  $\mathcal{A}$  returns a schedule  $\sigma'$  whose makespan is (strictly) smaller than 11. Since  $\mathcal{A}$  is a deterministic algorithm, it must execute the two tasks of length 5 on the same machine. Then, one task of length 5 ends strictly before 6 time units (the first to be scheduled).

Since  $\mathcal{A}$  is truthful, no task can bid a larger length and improve its completion time. Then, among the two tasks of length 3 which are completed after 6 time units in  $\sigma$ , none can bid 5 and be the task which ends strictly before 6 time units in  $\sigma'$ . We now show that  $\mathcal{A}$  cannot avoid this since its approximation ratio is strictly smaller than 1.1.

Let  $ID = \{a, b, c, d\}$  be a set of 4 distinct identification numbers (*id* stands for *identification number* in the sequel). For each  $x \in ID$ , we define  $I_x$  as an instance similar to  $I$  and for which the ids are assigned as follows: the task of length 5 gets the id  $x$  and the 3 tasks of length 3 get an id in  $ID - \{x\}$  (the task of length 4 is always given id  $e \notin ID$ ). Let  $\sigma_x$  be the schedule returned by  $\mathcal{A}$  when  $I_x$  is the input. Let  $S = \{I_x \mid x \in ID\}$ . We have  $|S| = 4$ .

For each couple  $\{x, y\} \subset ID$ , we define  $I'_{x,y}$  as an instance similar to  $I'$  and for which the ids are assigned as follows: the two tasks of length 5 get the ids  $x$  and  $y$ , while the two tasks of length 3 get their ids from  $ID - \{x, y\}$  (the task of length 4 is still given id  $e$ ). Let  $\sigma'_{x,y}$  be the schedule returned by  $\mathcal{A}$  when  $I'_{x,y}$  is the input. Let  $S' = \{I'_{x,y} \mid \{x, y\} \subset ID\}$ . We have  $|S'| = 6$ .

Let  $f : S' \rightarrow S$  be a function such that  $f(I'_{x,y}) = I_x$  if the task with id  $x$  is scheduled after the one with id  $y$  in  $\sigma'_{x,y}$ , otherwise  $f(I'_{x,y}) = I_y$ . Since  $|S'| > |S|$ , there exists a couple of instances in  $S'$  which both have the same image in  $S$  by  $f$ . Without loss of generality, we suppose that these instances are  $I'_{a,b}$  and  $I'_{a,c}$ . Moreover, we suppose that  $f(I'_{a,b}) = f(I'_{a,c}) = I_a$ . Since  $I_a$  contains three tasks of length 3 (with ids  $b, c$  and  $d$ ) which are scheduled on the same machine in  $\sigma_a$ , two of them have a completion time larger than or equal to 6. This set of two tasks must have a nonempty intersection with  $\{b, c\}$ . Without loss of generality, we suppose that task  $b$  with length 3 has a completion time larger than or equal to 6 in  $\sigma_a$ . Due to the definition of  $f$ ,  $b$  can bid 5 instead of 3 and be executed before  $a$  in  $\sigma'_{a,b}$ . As previously observed,  $b$  will then end strictly before 6 time units. As a consequence,  $\mathcal{A}$  cannot be a truthful and  $\rho$ -approximate with  $\rho < 1.1$ .

Using the same technique with slightly modified instances, we can get an improved lower bound. Indeed  $\mathcal{A}$  cannot be a truthful and  $\rho$ -approximate with  $\rho < 1 + (\sqrt{105} - 9)/12 \approx 1.1039$ . Instance  $I$  contains three tasks of length 2, one task of length  $2 + 2\varepsilon$  and one task of length  $3 + \varepsilon$ . Instance  $I'$  contains two tasks of length 2, one task of length  $2 + 2\varepsilon$  and two tasks of length  $3 + \varepsilon$ . The bound is obtained when  $\varepsilon = (\sqrt{105} - 9)/4$ .

**Proof of Theorem 5 :**

Let us suppose that we have a truthful algorithm  $\mathcal{A}$  with an approximation ratio  $\rho < 7/6$ . Let  $I$  be the following instance: one task of length 3 and  $3m - 2$  tasks of length 2. Let  $\sigma$  be the schedule returned by  $\mathcal{A}$  when  $I$  is the input. Since there are only  $m$  machines available, at most  $m$  tasks of length 2 are completed before four time units, and thus at least  $2m - 2$  tasks with the same length have a completion time which is larger than or equal to 4.

Let us now consider the following instance  $I'$ : two tasks of length 3 and  $3m - 3$  tasks of length 2. The optimal makespan is 6. When  $I'$  is the input,  $\mathcal{A}$  returns a schedule  $\sigma'$  whose makespan is (strictly) smaller than 7. One can remark that  $\mathcal{A}$  must execute the two tasks of length 3 on the same machine. Then, the task of length 3 which is scheduled first ends strictly before 4 time units in  $\sigma'$ .

Since  $\mathcal{A}$  is truthful, no task can bid a larger length and improve its completion time. In particular, among the tasks of length 2 which are completed after 4 time units in  $\sigma$ , none can unilaterally bid 3 and be the task which ends strictly before 4 time units in  $\sigma'$ . We now show that  $\mathcal{A}$  cannot avoid this since its approximation ratio is strictly smaller than  $7/6$ .

Let  $ID$  be a set of  $3m - 1$  distinct identification numbers (ids). For each subset  $\{a, b\} \subset ID$ , the instance  $I'_{a,b}$  is equal to the instance  $I'$  in which the ids are as follows: the two tasks of length 3 get the ids  $a$  and  $b$ , while the  $3m - 3$  tasks of length 2 get an id in  $ID - \{a, b\}$ . Let  $\sigma'_{a,b}$  be the schedule returned by  $\mathcal{A}$  when  $I'_{a,b}$  is the input.

$ID$  must contain  $m + 2$  distinct ids denoted by  $i_1, \dots, i_{m+2}$  and such that:  $\forall x \in \{1, \dots, m + 1\}$ , the task with id  $i_x$  is scheduled before the one with id  $i_{m+2}$  in  $\sigma'_{i_x, i_{m+2}}$ .

Now, we build an instance  $I_{i_{m+2}}$  similar to  $I$  as follows: the task of length 3 get the id  $i_{m+2}$  and the  $3m - 2$  tasks of length 2 get an id in  $ID - \{i_{m+2}\}$ . Let  $\sigma_{i_{m+2}}$  be the schedule returned by  $\mathcal{A}$  when  $I_{i_{m+2}}$  is the input. There exists at least one task with id  $i_y$  such that  $y \in \{1, \dots, m + 1\}$  and its completion time in  $\sigma_{i_{m+2}}$  is larger than or equal to 4. Indeed, at most  $m$  tasks of length 2 can end strictly before 4 time units since  $m$  machines are available. As consequence, the task with id  $i_y$  can bid 3 instead of 2 and improve its completion time.

As a consequence,  $\mathcal{A}$  cannot be truthful and this shows that there is no truthful algorithm which has an approximation ratio  $\rho < 7/6$ .  $\square$

### Proof of Theorem 6:

*Proof.* Let us suppose that we have a truthful and  $(7/6 - \epsilon)$ -approximate algorithm. We consider an instance  $I_1$  with four tasks of length 2 (with identification numbers  $a, b, c$  and  $d$ ) and one task of length  $3 - 5\epsilon$  (with id  $e$ ). In the solution built by the algorithm, at least two tasks of length 2 have a completion time larger than or equal to 4. We can observe that if one of them bids  $3 - 4\epsilon$  or  $3 - 3\epsilon$  then it will necessarily be executed on the same machine as task  $e$  of length  $3 - 5\epsilon$  (this is due to the fact that the algorithm is  $(7/6 - \epsilon)$ -approximate). Since the algorithm is truthful, the task which lied must be executed after task  $e$  (otherwise, its completion decreased).

Consider the instance  $I_2$  with four tasks of length 2 (with identification numbers  $a, b, c$  and  $d$ ) and one task of length  $3 - 4\epsilon$  (with id  $e$ ). In the solution built by the algorithm, at least two tasks of length 2 have a completion time larger than or equal to 4. If one of them bids  $3 - 3\epsilon$  or  $3 - 5\epsilon$  then it will necessarily be executed on the same machine as task  $e$  of length  $3 - 4\epsilon$ . Since the algorithm is truthful, the task which lied must be executed after task  $e$ .

Consider the instance  $I_3$  with four tasks of length 2 (with identification numbers  $a, b, c$  and  $d$ ) and one task of length  $3 - 3\epsilon$  (with id  $e$ ). In the solution built by the algorithm, at least two tasks of length 2 have a completion time larger than or equal to 4. If one of them bids  $3 - 4\epsilon$  or  $3 - 5\epsilon$  then it will necessarily be executed on the same machine as task  $e$  of length  $3 - 3\epsilon$ . Since the algorithm is truthful, the task which lied must be executed after task  $e$ .

Let  $T_1$  (resp.  $T_2, T_3$ ) be the set of tasks of length 2 whose completion time is equal or larger than or equal to 4 in the solution returned by the algorithm when  $I_1$  (resp.  $I_2, I_3$ ) is the input. One can find a couple of sets in  $\{T_1, T_2, T_3\}$  such that their intersection is non empty. W.l.o.g., we suppose that task  $d$  is in  $T_1 \cap T_3$ . Now consider the following instances:

1.  $(2, a), (2, b), (2, c), (3 - 3\epsilon, d), (3 - 5\epsilon, e)$
2.  $(2, a), (2, b), (2, c), (3 - 5\epsilon, d), (3 - 3\epsilon, e)$

For the first instance, we observed that the algorithm executes the task of length  $3 - 5\epsilon$  before the one of length  $3 - 3\epsilon$ . For the second instance, we observed the opposite. Though some tasks have a unique length, the algorithm must take into account their identification number to be truthful and  $(7/6 - \epsilon)$ -approximate.

### Proof of Theorem 8:

*Proof.* Let us first prove this result in the case where there are two machines,  $P_1$ , and  $P_2$ . Let  $\epsilon > 0$ . Let us suppose that there exists a truthful coordination mechanism  $\mathcal{M}$  with a price of anarchy of  $\rho < \frac{2+\alpha}{2}$  (with  $\alpha < 1$ ). Let us consider the following instance  $I_1$ : three tasks of length 1. Since  $\mathcal{M}$  is a coordination mechanism which induces pure Nash equilibria, there is at least a task in  $I_1$  which has a completion time larger than or equal to 1.5 (because there is a machine where there is at least 2 tasks). Let  $t$  be such a task. Let us suppose that  $t$  bids  $1 + \alpha$  (with  $\alpha \leq 1$ ). Since  $\rho < \frac{2+\alpha}{2}$ , then task  $t$  is alone on its machine (without loss of generality, on  $P_1$ ), otherwise the makespan of the schedule would be greater than  $\rho$  times the makespan of an optimal schedule which is 2. Since  $\mathcal{M}$  is truthful, the expected completion time of  $t$  when it bids  $1 + \alpha$  should be larger than or equal to 1.5 (otherwise  $t$  would have incentive to lie). Thus, when machine  $P_1$  has a task of length  $1 + \alpha$ , the expected completion time of this task is at least 1.5. Let us consider the instance  $I_2$  where we have two tasks of length  $1 + \alpha$ . Since  $\rho < \frac{2+\alpha}{2} < 2$ , there is one task on each machine. Since, when there is one task of length  $1 + \alpha$  on machine  $P_1$ , the expected completion time of this task is larger than or equal to 1.5, and since the optimal makespan of a schedule of these tasks is  $1 + \alpha$ , then, the price of anarchy of  $\mathcal{M}$  should be larger

than or equal to  $\frac{1.5}{1+\alpha}$ . By setting  $\frac{2+\alpha}{2} = \frac{1.5}{1+\alpha}$ , we obtain  $\alpha = \frac{\sqrt{13}-3}{2}$ , and thus  $\rho > 1 + \frac{\sqrt{13}-3}{2} > 1.15$ .

We can easily extend this proof in the case where there are more than 2 machines, by having  $m+1$  tasks of length 1 in  $I_1$ ; and  $m$  tasks of length  $1+\alpha$  in  $I_2$ .