

# An Optimal Lower Bound for Anonymous Scheduling Mechanisms

Itai Ashlagi  
Harvard Business School  
Harvard, MA, USA  
iashlagi@hbs.edu

Shahar Dobzinski\*  
The School of Computer  
Science and Engineering  
Hebrew University of  
Jerusalem, Israel  
shahard@cs.huji.ac.il

Ron Lavi†  
IE&M  
Technion, Israel  
ronlavi@ie.technion.ac.il

## ABSTRACT

We consider the problem of designing truthful mechanisms to minimize the makespan on  $m$  unrelated machines. In their seminal paper, Nisan and Ronen [14] showed a lower bound of 2, and an upper bound of  $m$ , thus leaving a large gap. They conjectured that their upper bound is tight, but were unable to prove it. Despite many attempts that yield positive results for several special cases, the conjecture is far from being solved: the lower bound was only recently slightly increased to 2.61 [5, 10], while the best upper bound remained unchanged.

In this paper we show the optimal lower bound on truthful anonymous mechanisms: no such mechanism can guarantee an approximation ratio better than  $m$ . This is the first concrete evidence to the correctness of the Nisan-Ronen conjecture, especially given that the classic scheduling algorithms are anonymous, and all state-of-the-art mechanisms for special cases of the problem are anonymous as well.

## Categories and Subject Descriptors

J.4 [SOCIAL AND BEHAVIORAL SCIENCES]: Economics; K.4.4 [Computers and Society]: Electronic Commerce—*Payment schemes*

## General Terms

Algorithms, Economics, Theory

## Keywords

Mechanism design, truthful, scheduling

\*Supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities, and by a grant from the Israeli Academy of Sciences.

†Supported by grants from the Israeli Science Foundation and the Binational Science Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'09, July 6–10, 2009, Stanford, California, USA.

Copyright 2009 ACM 978-1-60558-458-4/09/07 ...\$5.00.

## 1. INTRODUCTION

### 1.1 Background

In their STOC'99 paper, Nisan and Ronen [14] introduced the field of Algorithmic Mechanism Design. In their paper Nisan and Ronen study a problem that takes a central place in the field of algorithmic mechanism design ever since: truthful scheduling on unrelated machines. In this problem,  $n$  jobs are to be allocated to  $m$  machines, where it takes machine  $i$   $t_i^j$  time units to process job  $j$  (and this is also the cost machine  $i$  incurs from executing job  $j$ ). The processing times are private information of the machines. The load of machine  $i$  is the sum of processing times of the jobs  $i$  is assigned. The designer aims to allocate the jobs to minimize the maximal load over all machines (the “makespan”).

Nisan and Ronen design a simple mechanism that achieves an approximation ratio of  $m$  (the VCG mechanism). They also show that *no* mechanism, even one with unlimited computational power, can guarantee an approximation ratio better than 2. Nisan and Ronen were unable to close this large gap. However, they conjecture that their upper bound is tight:

**Conjecture (Nisan-Ronen):** No mechanism for the truthful scheduling problem on unrelated machines can achieve an approximation ratio better than  $m$ .

A decade later, although widely believed and extensively studied, this conjecture is far from being solved. In fact, despite many efforts by the community, no additional evidence, to either support the conjecture or to weaken the belief in it, has been provided since. Only recently, Christodoulou, Koutsoupias, and Vidali [5] were able to slightly improve the lower bound from 2 to 2.41, further improving it to 2.61 later on [10]. Mu’alem and Schapira [13] and Christodoulou et al. [4] prove a lower bound of 2 for randomized and fractional mechanisms, respectively. Dobzinski and Sundararajan [9] and Christodoulou, Koutsoupias, and Vidali [3] attempted to characterize truthful mechanisms for this problem, but succeeded in doing so only for the very limited case of 2 machines. All these papers involve many non-trivial observations and technicalities, further emphasizing the basic difficulties that this problem entails.

A different line of research attempted to *relax* the problem, by finding tractable special cases. Archer and Tardos [2] observe that if we consider related machines, then

the problem becomes single-parameter<sup>1</sup>. Mechanisms for single-parameter problems are easier to design as only monotonicity has to be ensured, a relatively easy task. Indeed, a truthful PTAS that essentially concludes all efforts in this direction was recently constructed [8]. Back in the multi-parameter setting, Lavi and Swamy [12] consider a special case where processing times can take only two possible values, “low” and “high”, and give several truthful algorithms that guarantee constant-factor approximation ratios.

Interestingly, all the known lower bounds are just small constants, and there are no super-constant lower bounds. Furthermore, as mentioned before, several truthful mechanisms that provide good approximation ratios for non-trivial special cases have been presented. At this point, the skeptical reader may start doubting the correctness of the Nisan-Ronen conjecture. Perhaps one should interpret the low values of the upper bounds for the special cases and the previous lower bounds as a signal that a truthful mechanism with a good approximation ratio do exist?

## 1.2 Our Result

In this paper we give the first strong, concrete evidence to the correctness of the Nisan-Ronen conjecture.

**Theorem:** No *anonymous* mechanism for the truthful scheduling problem on unrelated machines can achieve an approximation ratio better than  $m$ .

An algorithm is anonymous, roughly speaking, if whenever two machines switch costs, the job assignments of the two machines also switch<sup>2</sup>. Note that this is the best lower bound possible as the Nisan-Ronen algorithm is anonymous. Let us explicitly spell out why the class of anonymous mechanisms is of interest:

- **Algorithmic Perspective:** The classic algorithms for scheduling unrelated machines are indeed anonymous. There does not seem to be an *algorithmic* reason that explains why a specific *naming* of the machines can help.
- **Mechanism Design Perspective:** Indeed, it is very easy to come up with non-anonymous mechanisms. However, are non-anonymous mechanisms more powerful than anonymous ones? All state-of-the-art mechanisms for the special cases in the literature are anonymous [12, 8]. This might suggest that anonymous mechanisms for this problem are as powerful as non-anonymous mechanisms<sup>3</sup>. In fact, a separation between the power of these two classes will be remarkable.

<sup>1</sup>A problem is *single parameter* when, informally, the private information of each player consists of only one number.

<sup>2</sup>For the mechanism to “notice” that the machines have switched costs, we of course require that the cost vectors of the machines are unique. See the preliminaries section for a formal definition.

<sup>3</sup>There is a single example in a different context (“digital goods”) where it is proved that anonymous mechanisms are less powerful [1]. However, this setting is a single-parameter one, comparing to our much complicated multi-parameter setting. Thus, one may doubt whether a complicated derandomization process, similar to the one used in [1], might be applicable in our multi-parameter setting. Furthermore, we have no evidence that randomized mechanisms for scheduling are significantly more powerful than deterministic ones.

- **Game-Theoretic Importance:** Not only are anonymous games interesting from a mechanism-design perspective, they are well-studied also from a wider game-theoretic point of view. In particular, anonymous mechanisms are desired from a pure game theoretic point of view as there is no discrimination between the bidders. For further motivation and recent examples of CS papers studying anonymous games see [6, 7]<sup>4</sup>.

To the very least, our result shows that if the Nisan-Ronen conjecture is false and there are mechanisms that provide a reasonable approximation ratio, then they must be very “strange”.

## 1.3 Tools and Techniques

Our proof shows that the “difficult” instance is actually the most simple instance from an algorithmic point of view: it is the instance in which all costs are between 1 and  $1 + \epsilon$ , and the cost vectors are ordered so that one cost vector completely dominates the former one, coordinate-wise. We prove that every truthful anonymous mechanism with a finite approximation ratio to the makespan must allocate *all jobs* to the machine with the lowest cost vector. This immediately yields a makespan which is  $m$  times the optimum. It is the same difficulty that the VCG mechanism encounters, and we show that all anonymous mechanisms suffer from the same drawback.

To show that we have to develop significant amount of new machinery. Specifically, the proof works inductively, starting from the observation that if we have only a single job, then every anonymous mechanism must allocate it to the lowest-cost machine. Unfortunately, this simple fact is not true when more jobs are added (for mechanisms that do not provide a good approximation ratio). The proof proceeds by following a subtle inductive process that requires substantial amount of technical work, which allows us to consider instances with increasing number of jobs. The approximation property is used in few crucially important places throughout the induction.

As the previous proofs do, we bootstrap the basic difficulty that truthfulness casts: given a specific assignment for one instance, truthfulness implies some restrictions on the possible assignments of all other instances that differ from the original instance by exactly one machine. Nisan and Ronen prove their lower bound by simply considering two such neighboring instances (i.e., a single transition from one instance to another). The other lower bounds that were mentioned above consider slightly longer paths of instances, thus improving the lower bound only slightly. The inductive techniques we develop let us tackle much longer paths of instances<sup>5</sup>. This is the key point that enables us to obtain the optimal lower bound. We believe that this is the main novelty of our proof.

Another important technical reason for our success in proving the lower bound is the way we exploit the *weak monotonicity* property. It is known that every truthful mechanism is also weakly monotone (see the preliminaries for a definition). However, the existing lower bounds prior to our work

<sup>4</sup>We note that the definition of anonymity in [6, 7] is slightly different than ours, but very related.

<sup>5</sup>An Exception to this is [10] which also uses induction. The induction used in the current paper is very different and more involved.

mainly used a limited and straightforward corollary of weak monotonicity: a machine that declares a vector of costs  $t$  and is allocated a bundle  $S$ , will be allocated the same bundle  $S$  when raising the costs of the jobs not in  $S$  and lowering the costs of the jobs in  $S$ . We use the weak monotonicity property in more sophisticated ways, by taking into account also the *amount* of which the costs of the jobs changes, whether allocated to the machine or not. This is another novel technical tool that enables us to prove the optimal lower bound. We believe that this tool will turn out to be even more helpful, as it might be used in solving other intriguing problems in algorithmic mechanism design.

## 1.4 Future Directions

Our proof gives strong evidence that deterministic mechanisms for this central problem do not have much power. Can the anonymity assumption be dropped? Our novel inductive method enables us to reach what seems to be the “correct” difficulty of truthful mechanisms. Thus, we believe that the inductive method and the new technical machinery we introduce might be the basis for further enhancements, to construct lower bounds without the anonymity assumption (though we have not managed to do so yet).

To this end, an interesting observation is that the anonymity property can easily be dropped in the *fractional* case, since, given any truthful approximation mechanism, one can construct an anonymous mechanism that averages over all permutations of players. This keeps the truthfulness and the approximation properties, and inserts anonymity. It might be possible, thus, to extend our proof to the fractional case, assuming anonymity without loss of generality, and obtain a general lower bound using this route.

## 2. PRELIMINARIES

### 2.1 Problem Definition and Notation

Let  $m$  be the number of machines, and let  $N$ ,  $|N| = n$ , be the set of jobs. We denote by  $t_i^j \geq 0$  the cost for machine  $i$  to process job  $j$  (this is also the time it takes machine  $i$  to process job  $j$ ). A cost vector for machine  $i$  is a vector  $t_i = (t_i^j)_{j=1}^n$ . For every bundle  $S \subseteq N$  we denote by  $t_i(S)$  the total cost for machine  $i$  to process the jobs in  $S$ . That is,  $t_i(S) = \sum_{j \in S} t_i^j$ .

Note that an instance of this problem can be viewed as a matrix in which the  $i^{\text{th}}$  row is the cost vector of machine  $i$ ,  $\vec{t}_i$ . We often use matrices to present instances, and use stars, “\*”, to indicate jobs assignments; a star next to the entry  $t_i^j$  indicates that machine  $i$  is assigned job  $j$ . For example, in the following matrix machine  $i$ ’s cost vector is  $(t_i^1, \dots, t_i^n)$  and all the jobs are assigned to machine 1:

$$\begin{pmatrix} t_1^1* & \dots & t_1^n* \\ \vdots & & \vdots \\ t_{m-1}^1 & \dots & t_{m-1}^n \\ t_m^1 & \dots & t_m^n \end{pmatrix}$$

### 2.2 Truthfulness and Anonymity

A mechanism consists of an allocation function  $f$ , and a payment function  $p_i$  for every machine  $i$ . I.e., for every valuation profile  $t = (t_1, \dots, t_m)$ ,  $f(t)$  is the allocation ( $f_i(t)$  is the bundle assigned to machine  $i$ ), and  $p_i(t)$  denotes the payment for machine  $i$  when the valuation profile is  $t$ . The

mechanism is *truthful* if, for every machine  $i$ , declaring its true cost vector is at least as good as declaring any other cost vector. Formally, for every machine  $i$ , every  $t_i$  and every  $t_{-i}$ ,

$$p_i(t_i, t_{-i}) - t_i(f_i(t)) \geq p_i(t'_i, t_{-i}) - t_i(f_i(t'_i, t_{-i}))$$

where  $f_i(t)$  denotes the bundle assigned to machine  $i$  in the cost vector  $t = (t_1, \dots, t_m)$ . We follow the standard notation:  $t_{-i}$  denotes the costs of all machines but  $i$ . In particular if  $t = (t_1, \dots, t_m)$  then  $t = (t_i, t_{-i})$ .

The following is a well-known necessary condition for an allocation function to be *implementable* (i.e., there exist payment functions that make the resulting mechanism truthful):

**DEFINITION 2.1** (WEAK MONOTONICITY [11]). *An allocation function  $f$  is called weakly monotone if for every machine  $i$ , and every  $t_i, t'_i, t_{-i}$  the following property is satisfied: suppose that  $f_i(t_i, t_{-i}) = S_i$  and that  $f_i(t'_i, t_{-i}) = S'_i$ , then  $t_i(S_i) - t_i(S'_i) \leq t'_i(S_i) - t'_i(S'_i)$ .*

A direct consequence of weak monotonicity, which we often use in our proof, is the following: if a machine is allocated a bundle  $S$  (including the empty set) in some given instance, then it will be allocated the same bundle if it lowers the costs for the jobs in  $S$  and raises the costs for all other jobs, while all other machines costs are fixed. Formally,

**CLAIM 2.2.** *Let  $f$  be a truthful mechanism for the makespan minimization problem. Let  $t = (t_1, \dots, t_m)$  be some set of valuations, and suppose that machine  $i$  is allocated some bundle  $S$  in  $f(t)$ . I.e.,  $f_i(t) = S$ . Let  $t'_i$  be some other cost vector of machine  $i$  where  $t'_i(\{j\}) < t_i(\{j\})$  for all  $j \in S$ , and  $t'_i(\{j\}) > t_i(\{j\})$  for all  $j \notin S$ . Then,  $f_i(t'_i, t_{-i}) = S$ .*

This paper deals with anonymous mechanisms. Roughly speaking, one expects a mechanism to be anonymous if for every two machines that switch their costs their assigned bundles switch too. There are some technical issues with this definition; for example, if all machines have the same costs, “switching the costs” makes no sense. To handle this and similar problems we only require the following: if machine  $k$  has a unique cost vector (a cost vector is said to be *unique* if it differs from every other cost vector in at least one coordinate) and is assigned bundle  $S$ , then whenever  $k$  switches costs with some other machine  $l$ , machine  $l$  is allocated the bundle  $S$  (we stress that no constraints are imposed on what all other machines are assigned after the switch, in particular machine  $k$ , nor on instances where there is no machine with a unique cost vector).

We note that this is only one possible definition of anonymity. This definition was chosen only because it is simple to present. Our proof seems to be general enough to hold using essentially any other reasonable definition of anonymity. In particular, the proof uses the anonymity property only in instances where all machines have unique cost vectors.

**DEFINITION 2.3** (AN ANONYMOUS MECHANISM). *A mechanism  $f$  is called anonymous if for every  $m$  cost vectors,  $\vec{t}_1, \dots, \vec{t}_m$ , and for each  $k$  such that  $\vec{t}_k$  is unique, it holds that for every  $l \neq k$ :*

$$f_l(\vec{t}_1, \dots, \vec{t}_{k-1}, \vec{t}_l, \vec{t}_{k+1}, \dots, \vec{t}_{l-1}, \vec{t}_k, \vec{t}_{l+1}, \dots, \vec{t}_m) = f_k(\vec{t}_1, \dots, \vec{t}_m).$$

### 3. THE LOWER BOUND

In this section we prove our main result:

**THEOREM 3.1.** *Let  $f$  be a truthful anonymous mechanism for  $m$  machines and  $n$  jobs that provides a  $c$ -approximation to the makespan. Consider an instance in which for every job  $j$ ,  $t_m^j > t_{m-1}^j > \dots > t_1^j$ .*

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{n*} \\ t_2^1 & t_2^2 & \dots & t_2^n \\ & & \vdots & \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^n \\ t_m^1 & t_m^2 & \dots & t_m^n \end{pmatrix}$$

*Then,  $f$  allocates all jobs to machine 1 (the allocation is indicated by stars). As a corollary, if for each  $j$ ,  $1+\epsilon > t_m^j > t_{m-1}^j > \dots > t_2^j > t_1^j = 1$  and  $m = n$  the approximation ratio of  $f$  approaches  $m$  as  $\epsilon \rightarrow 0$ .*

We prove this theorem by induction on the number of jobs, first considering mechanisms for one job, then gradually increasing the number of jobs. Next we formally define the induction hypothesis.

#### 3.1 The Induction Hypothesis and the Base Case

We will actually prove a stronger version of the main theorem. The stronger version also serves as our induction hypothesis:

**The Induction Hypothesis:** Fix  $\beta > 0$ . Let  $f$  be a truthful anonymous mechanism for  $m$  machines and  $n$  jobs that is guaranteed to provide a  $c$ -approximation to the makespan, but only to instances with an optimal makespan of at least  $\beta$ . Consider an instance in which for every job  $j$ ,  $t_m^j > t_{m-1}^j > \dots > t_2^j > t_1^j$ .

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{n*} \\ t_2^1 & t_2^2 & \dots & t_2^n \\ & & \vdots & \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^n \\ t_m^1 & t_m^2 & \dots & t_m^n \end{pmatrix}$$

Then,  $f$  allocates all jobs to machine 1 (the allocation is indicated by stars).

All the instances we consider during the proof will have an optimal makespan larger than some constant  $\beta$ . Since the proof holds for any  $\beta > 0$ , it implies Theorem 3.1.

We first prove the induction hypothesis for the base case where there is only one job. In this case the proof is very simple.

**LEMMA 3.2.** *Let  $f$  be an anonymous mechanism for one job and  $m$  machines. Suppose that  $t_1 < t_2 < \dots < t_m$ . Then, machine 1 must be allocated the job in the following instance:*

$$\begin{pmatrix} t_1^* \\ \vdots \\ t_{m-1} \\ t_m \end{pmatrix} \quad (1)$$

**PROOF.** Assume towards a contradiction that the job is allocated to machine  $i \neq 1$ . Now lower machine  $i$ 's cost to  $t_1$ . By weak monotonicity (see Claim 2.2) machine  $i$  keeps

the job. In particular, machine 1 is not allocated the job. Now raise machine 1's cost to  $t_i$ . The job is still not allocated to machine 1, by weak monotonicity. Notice however that machine 1 and  $i$  have switched their costs comparing to the instance (1), so by anonymity machine 1 should have received the job. A contradiction.  $\square$

#### 3.2 The Central Argument

We assume the induction hypothesis is true for any  $n' < n$ . The following Lemma will almost prove the induction hypothesis.

**LEMMA 3.3.** *Let  $f$  be a truthful and anonymous mechanism that provides a  $c$ -approximation ratio (for some real number  $c$ ) for  $m$  machines and  $n$  jobs. Fix any  $k \leq n-1$ . Suppose an instance  $t$  such that for every job  $1 \leq j \leq k$ ,  $t_m^j > t_{m-1}^j > \dots > t_2^j > t_1^j$ , and for every job  $k < j \leq n$ ,  $t_m^j \geq t_{m-1}^j \geq \dots \geq t_2^j$ . Fix any constant  $a > 0$  such that  $200 \cdot c \cdot n \cdot a < \min_{1 \leq j \leq k} t_1^j$ ,  $200 \cdot c \cdot n \cdot a < \min_{1 \leq j \leq n} t_2^j$ , and  $2na < \min_{1 \leq j \leq k} (t_2^j - t_1^j)$ . Then, machine 1 is allocated all jobs in the following instance:*

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{k*} & a^* & \dots & a^* \\ t_2^1 & t_2^2 & \dots & t_2^k & t_2^{k+1} & \dots & t_2^n \\ & & & \vdots & & & \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^k & t_{m-1}^{k+1} & \dots & t_{m-1}^n \\ t_m^1 & t_m^2 & \dots & t_m^k & t_m^{k+1} & \dots & t_m^n \end{pmatrix} \quad (2)$$

Note that, since  $c \cdot n \cdot a < \min_j t_1^j$  we also have  $n \cdot a < \beta$ .

Notice that it is not clear that machine 1 will receive even the jobs  $k+1, \dots, n$  for which its cost is very small,  $a$ , although allocating these jobs to machine 1 seems like the "right" thing to do. For example, allocating jobs  $1, \dots, k$  to machine 1 and the rest of the jobs to machines  $2, \dots, m$  might also provide a good approximation. In particular, recall that we cannot assume that these jobs are allocated to machine 1 "without loss of generality": a priori it might be the case that such weird allocations are necessary to construct a truthful mechanism that provides a good approximation ratio.

The proof of the lemma is by induction on  $k$ . The case of  $k=0$  follows immediately since  $f$  is a  $c$ -approximation. We now fix some  $0 < k \leq n-1$ , assume the lemma is correct for every  $k' < k$ , and prove it for  $k$ . Before proving the lemma, we provide an informal overview of the proof.

#### An Overview of the Proof of Lemma 3.3

The proof of Lemma 3.3 is, yet again, inductive. Below we give an overview of the induction step for an arbitrary  $k$ .

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{k*} & a & \dots & a \\ t_2^1 & t_2^2 & \dots & t_2^k & a & \dots & a \\ & & & \vdots & & & \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^k & a & \dots & a \\ t_m^1 & t_m^2 & \dots & t_m^k & a & \dots & a \end{pmatrix}$$

*Step 1:* We show that if the mechanism was "restricted" only to the first  $k$  jobs, then machine 1 was allocated all of them. Using this we show that even in the unrestricted mechanism machine 1 is still allocated the first  $k$  jobs (but not necessarily jobs  $k+1, \dots, n$ ).

$$\begin{pmatrix} t_1^1 - \epsilon^* & t_1^2 - \epsilon^* & \dots & t_1^k - \epsilon^* & \delta^* & \dots & \delta^* \\ t_2^1 & t_2^2 & \dots & t_2^k & a & \dots & a \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^k & a & \dots & a \\ t_m^1 & t_m^2 & \dots & t_m^k & a & \dots & a \end{pmatrix}$$

*Step 2:* Here we verify that machine 1 is allocated all jobs. We do that by reducing the costs of the first  $k$  jobs by more than  $a$ , and reducing the costs of jobs  $k+1, \dots, n$  from  $a$  to a much smaller  $\delta$ . Weak monotonicity and approximation arguments guarantee that machine 1 is allocated all jobs.

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{k*} & \delta^* & \dots & \delta^* \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ t_{i-1}^1 & t_{i-1}^2 & \dots & t_{i-1}^k & a & \dots & a \\ t_i^1 & t_i^2 & \dots & t_i^k & t_i^{k+1} & \dots & t_i^n \\ t_{i+1}^1 & t_{i+1}^2 & \dots & t_{i+1}^k & t_{i+1}^{k+1} & \dots & t_{i+1}^n \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ t_m^1 & t_m^2 & \dots & t_m^k & t_m^{k+1} & \dots & t_m^n \end{pmatrix}$$

*Step 3:* This is the most technically involved step. We raise the costs of machines 2 to  $m$  one by one, starting from machine  $m$ , and show that machine 1 is still allocated all jobs during the process.

### Step 1 – Adding Small Jobs

This step shows that by “adding” small jobs that all machines value with the same very low cost does not change the allocation of the first  $k$  jobs. The statement of the lemma is almost obvious given the induction hypothesis, but the proof is a bit subtle.

LEMMA 3.4. *Let  $f$  be a truthful and anonymous mechanism that provides a finite approximation ratio for  $m$  machines and  $n$  jobs, and suppose an instance  $t$  and a constant  $a$  as in Lemma 3.3 above. Then, jobs  $1, \dots, k$  are allocated to machine 1 in the instance below (no claim is made regarding the allocation of jobs  $k+1, \dots, n$ ):*

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{k*} & a & \dots & a \\ t_2^1 & t_2^2 & \dots & t_2^k & a & \dots & a \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^k & a & \dots & a \\ t_m^1 & t_m^2 & \dots & t_m^k & a & \dots & a \end{pmatrix} \quad (3)$$

PROOF. Consider the following allocation function  $f^k$  on  $m$  machines and  $k$  jobs, which is defined as follows: given a cost vector  $\vec{t}_i$  of  $k$  jobs for machine  $i$  define the cost vector  $t'_i: t'_i(\{j\}) = t_i(\{j\})$  for  $j \leq k$ , and  $t'_i(\{j\}) = a$  otherwise. Let  $(S_1, \dots, S_m) = f(t'_1, \dots, t'_m)$ . Set  $f^k(t_1, \dots, t_n) = (S_1 \setminus \{k+1, \dots, n\}, \dots, S_m \setminus \{k+1, \dots, n\})$ . In other words,  $f^k$  “runs”  $f$  when the cost of each of the last  $n-k$  jobs is  $a$ , and allocates its  $k$  jobs the same way  $f$  allocates its first  $k$  jobs.

PROPOSITION 3.5. *The allocation function  $f^k$  is implementable.*

PROOF. Recall the following characterization of truthful mechanisms:

THEOREM 3.6 (SAKS AND YU [15]). *Let  $f$  be an allocation function defined on a convex domain of valuations.  $f$  is truthfully implementable if and only if it is weakly monotone.*

By this theorem (which is applicable since the scheduling domain is convex), all we have to prove is that  $f^k$  satisfies weak monotonicity. Fix some  $n$ -jobs cost vectors  $t_1, \dots, t_{i-1}, t_i, \dots, t_m$  of all machines but machine  $i$ , where the cost of each machine for jobs  $k+1, \dots, n$  is  $a$ . Suppose machine  $i$  is assigned  $S_i$  when declaring  $t_i$  and  $S'_i$  when declaring  $t'_i$  in  $f$  (the cost of machine  $i$  in  $t'_i$  for jobs  $k+1, \dots, n$  is still  $a$ ). The function  $f$  is implementable and thus satisfies weak monotonicity. Hence,

$$t_i(S_i) - t_i(S'_i) \leq t'_i(S_i) - t'_i(S'_i).$$

Therefore, since machine  $i$  did not change its costs of jobs  $k+1, \dots, n$ , we also have that

$$t_i(S_i \setminus \{k+1, \dots, n\}) - t_i(S'_i \setminus \{k+1, \dots, n\}) \leq t'_i(S_i \setminus \{k+1, \dots, n\}) - t'_i(S'_i \setminus \{k+1, \dots, n\})$$

which implies that  $f^k$  also satisfies weak monotonicity.  $\square$

By construction  $f$  and  $f^k$  allocate jobs  $j = 1, \dots, k$  identically. Therefore, it is enough to show that  $f^k$  assigns jobs  $j = 1, \dots, k$  to machine 1. Towards this end, observe that  $f^k$  is anonymous, since the costs of all machines for jobs  $k+1, \dots, n$  is  $a$  and  $f$  is anonymous. Observe also that  $f^k$  has a finite approximation ratio  $\tilde{c} = 2c$  on all instances of  $k$  jobs that has optimal makespan of at least  $\beta$ , by the following argument: let  $t^k$  be such an instance on  $k$  jobs, and let  $t$  be an instance on  $n$  jobs which is created from  $t^k$  by adding  $n-k$  “a-columns”. Let  $OPT(t)$  and  $OPT(t^k)$  denote the optimal makespans of these two instances. Since  $OPT(t) \geq OPT(t^k) \geq \beta$  we know by assumption that the makespan of the allocation  $f(t)$  is a  $c$ -approximation to  $OPT(t)$ . We get: the makespan of  $f^k(t^k)$  is at most the makespan of  $f(t)$  which is at most  $c \cdot OPT(t)$ . Now  $OPT(t) \leq OPT(t^k) + n \cdot a$  and  $n \cdot a < \beta < OPT(t^k)$ . Therefore the makespan of  $f^k(t^k)$  is at most  $2c \cdot OPT(t^k)$ , as we need. Hence, by the induction hypothesis machine 1 is allocated jobs  $1, \dots, k$  in  $f^k(t^k)$  and hence also by  $f(t)$ .

### Step 2 – Lowering the Cost of Machine 1

The massaging process of our instance will be much easier by guaranteeing that all jobs, not just jobs  $1, \dots, k$ , are allocated to machine 1.

LEMMA 3.7. *Let  $f$  be a truthful and anonymous mechanism that provides a  $c$ -approximation for  $m$  machines and  $n$  jobs, and suppose cost vectors  $\vec{t}_1, \dots, \vec{t}_n$  and a constant  $a$ , all satisfy the conditions of Lemma 3.3 above. Fix  $\epsilon > 0$  and  $\delta > 0$  such that  $\min_j(t_j^1) > \epsilon > n \cdot a$ ,  $\delta < a/(200nc)$ , and  $(3c)^{2m+1}2n^2c\delta < a - \delta$ . Then all jobs are allocated to machine 1 in the following instance:*

$$\begin{pmatrix} t_1^1 - \epsilon^* & t_1^2 - \epsilon^* & \dots & t_1^k - \epsilon^* & \delta^* & \dots & \delta^* \\ t_2^1 & t_2^2 & \dots & t_2^k & a & \dots & a \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^k & a & \dots & a \\ t_m^1 & t_m^2 & \dots & t_m^k & a & \dots & a \end{pmatrix} \quad (4)$$

PROOF. Notice that only machine 1 changed its costs comparing to instance (3). To see that in the instance (4) machine 1 gets all jobs we use two claims. Both weak monotonicity arguments and the assumption that  $f$  provides a finite approximation ratio are used to prove this claim.

CLAIM 3.8. *Machine 1 is allocated at least jobs  $1, \dots, k$  in the instance (4).*

PROOF. We have shown that in the instance (3) machine 1 gets some bundle  $S$  where  $S$  contains all jobs  $1, \dots, k$ , and possibly some more jobs. By our choice of  $\epsilon$ , the cost of every job in  $S$  decreased more than the cost of all jobs  $k+1, \dots, n$  together. Thus, by weak monotonicity machine 1 is allocated in the instance (4) some bundle that includes all jobs  $1, \dots, k$ .  $\square$

CLAIM 3.9. *Machine 1 is allocated all jobs  $1, \dots, n$  in the instance (4).*

PROOF. Suppose not. Therefore, by the previous claim, machine 1 is allocated all jobs  $1, \dots, k$  but is not allocated some jobs from  $k+1, \dots, n$ :

$$\begin{pmatrix} t_1^1 - \epsilon^* & t_1^2 - \epsilon^* & \dots & t_1^k - \epsilon^* & \delta & \dots & \delta^* \\ t_2^1 & t_2^2 & \dots & t_2^k & a & \dots & a \\ & & & & \vdots & & \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^k & a & \dots & a \\ t_m^1 & t_m^2 & \dots & t_m^k & a & \dots & a \end{pmatrix} \quad (5)$$

Consider the following change in machine 1's costs: the costs of all jobs that it receives in the instance (5) are lowered to  $\frac{\delta}{2}$ , and the costs of the jobs that it did not receive are raised to  $2\delta$ :

$$\begin{pmatrix} \frac{\delta}{2}^* & \frac{\delta}{2}^* & \dots & \frac{\delta}{2}^* & 2\delta & \dots & \frac{\delta}{2}^* \\ t_2^1 & t_2^2 & \dots & t_2^k & a & \dots & a \\ & & & & \vdots & & \\ t_{m-1}^1 & t_{m-1}^2 & \dots & t_{m-1}^k & a & \dots & a \\ t_m^1 & t_m^2 & \dots & t_m^k & a & \dots & a \end{pmatrix}$$

By weak monotonicity the allocation of machine 1 stays the same. Thus, the mechanism does not provide a good approximation ratio: the optimal makespan is less than  $2n\delta$  (allocate all jobs to machine 1), while the makespan provided here is at least  $a > 200nc\delta$ . A contradiction.  $\square$

### Step 3 – Raising the Costs of Machines 2 to $m$

In this step we raise machine  $i$ 's costs of every job  $j$ ,  $k+1 \leq j \leq n$  from  $a$  to  $t_i^j$  (starting with the instance (4)), and show that machine 1 still receives all jobs. We first need the following two useful claims.

CLAIM 3.10. *In every instance that can be represented as in the matrix (2) that satisfies the conditions detailed in Lemma 3.3, machine 1 is either allocated all jobs  $1, \dots, k$  or none of them.*

PROOF. Suppose that machine 1 is allocated bundle  $T$  that includes some of the jobs  $1, \dots, k$  (but not all of them). Raise the cost of each job  $j \in \{1, \dots, k\} \setminus T$  to some  $t_1^j$ , so that for each such  $j$  we still have that  $2na < \min_j(t_2^j - t_1^j)$ . Lower the cost of all jobs in  $T$  to  $a$ , and keep the costs of the other jobs the same. By weak monotonicity machine 1 still

does not get the jobs  $\{1, \dots, k\} \setminus T$ , and machine 1 receives  $n - k' \delta$  jobs for some  $k' \leq k - 1$ . However, by induction Lemma 3.3 is already proved for smaller values of  $k$  (note that all requirements of the Lemma regarding  $t'$  and  $a$  are satisfied), and this lemma says that if the cost of more than  $n - k$  jobs to machine 1 is  $a$ , then machine 1 is allocated all jobs<sup>6</sup>. This is a contradiction and the claim follows.  $\square$

CLAIM 3.11. *In every instance that can be represented as in the matrix (2) that satisfies the conditions detailed in Lemma 3.3, if machine 1 is allocated jobs  $1, \dots, k$  then it is allocated all jobs.*

PROOF. The proof is similar to the proof of Claim 3.9. Suppose towards a contradiction that machine 1 is not allocated some set  $T$  of jobs  $k+1, \dots, n$  (and is allocated the rest of the jobs). Lower the costs of all jobs that machine 1 received to  $\frac{a}{2}$ , and raise the cost of all jobs in  $T$  to  $2a$ . Notice that the allocation to machine 1 stays the same, in particular machine 1 does not get the jobs in  $T$ . However, the approximation ratio provided is poor: the optimal solution allocates all jobs to machine 1 and has a makespan of less than  $2an$ , while the allocation provided by the mechanism has a makespan of at least  $\min_j t_2^j > c \cdot 2an$ .  $\square$

We now start to raise the costs of the machines on jobs  $k+1 \dots n$ . We first raise the cost of machine  $m$ , then machine  $m-1$  and so on. The definition of the family  $I_i$  aims to capture our state after raising the cost of machine  $i$ .

DEFINITION 3.12. *For each  $i \geq 2$ , an instance  $U$  belongs to the family  $I_i$  if it can be represented as the matrix below, where the  $t_i^j$ 's and  $a$  satisfy the conditions of Lemma 3.3, and  $\delta$  satisfies the conditions of Lemma 3.7.*

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{k*} & \delta^* & \dots & \delta^* \\ t_2^1 & t_2^2 & \dots & t_2^k & a & \dots & a \\ & & & & \vdots & & \\ t_{i-1}^1 & t_{i-1}^2 & \dots & t_{i-1}^k & a & \dots & a \\ t_i^1 & t_i^2 & \dots & t_i^k & t_i^{k+1} & \dots & t_i^n \\ t_{i+1}^1 & t_{i+1}^2 & \dots & t_{i+1}^k & t_{i+1}^{k+1} & \dots & t_{i+1}^n \\ & & & & \vdots & & \\ t_m^1 & t_m^2 & \dots & t_m^k & t_m^{k+1} & \dots & t_m^n \end{pmatrix} \quad (6)$$

LEMMA 3.13. *In every instance that belongs to the family  $I_i$ , for  $i \geq 2$ , machine 1 is allocated all jobs.*

Notice that proving that machine 1 is allocated all jobs in instances that belong to  $I_2$  will immediately give us Lemma 3.3 for this value of  $k$ .

PROOF. The proof of the lemma will make use of the following claim:

CLAIM 3.14. *If machines  $i, \dots, m$  do not receive jobs in every instance that belongs to  $I_i$  then machine 1 is allocated all jobs (for  $i \geq 2$ ).*

<sup>6</sup>Formally speaking, Lemma 3.3 is proved for smaller values of  $k$  only for the case where the value of the right-most jobs is  $a$ , and the value of the rest of the jobs is more than  $\delta$ , while here we use the lemma for instances where this is not the case. Notice, however, that the numbering (or positioning) of the jobs was not used during the proof, hence the Lemma holds for every possible numbering of the jobs.

PROOF. Notice that every instance in  $I_i$  satisfies the conditions of claims 3.10 and 3.11. <sup>7</sup> Therefore we need to examine only the case where machine 1 is allocated none of the jobs  $1, \dots, k$ . Suppose by contradiction that machine 1 does not get any of the jobs  $1, \dots, k$ . For any  $\gamma > 0$ , let  $t_1(\gamma)$  be a cost vector where  $t_1^j(\gamma) = t_1^j + \gamma$  for  $j \leq k$  and  $t_1^j(\gamma) = \delta + \gamma/n$  for  $j > k$ . If machine 1 increases its cost vector to  $t_1(\gamma)$  where then by weak monotonicity it is still not allocated any of the jobs  $1, \dots, k$ . We will show that there exist  $\gamma_2 > \gamma_1 > 0$  such that: (1) in the instances where machine 1 declares either  $t_1(\gamma_1)$  or  $t_1(\gamma_2)$  and the rest of the machines declare as before, there exists a machine  $l$  (the same machine for both declarations  $t_1(\gamma_1)$  and  $t_1(\gamma_2)$ ) that receives some of the jobs  $1, \dots, k$ , (2) the instance where machine 1 declares  $t_1(\gamma_1)$ , machine  $l$  declares  $t_1(\gamma_2)$ , and the rest of the machines declare as before, satisfies the conditions of claims 3.10 and 3.11.

We will now explicitly show the existence of  $\gamma_1, \gamma_2$ . Divide the interval  $[n\delta, a - \delta]$  to  $2m$  subintervals where the  $d$ 'th interval is of length  $(3c)^d 2n^2 c \delta$ . Choose an arbitrary point  $\gamma$  in each interval and label the interval as  $r$ , where machine  $r$  receives some of the jobs  $1 \dots k$  when machine 1 declares  $t_1(\gamma)$  (and the rest of the machines declare as before). Since there are more intervals than machines there exist two non-adjacent subintervals that are labeled by the same  $r$ , and we choose  $\gamma_1, \gamma_2$  as the points taken from these intervals. To verify that  $\gamma_1, \gamma_2$  satisfy the conditions of above, we need to verify (1)  $\gamma_2 - \gamma_1 > n(\delta + \gamma_1/n)$ , (2)  $\delta + \gamma_2/n < a$ , (3) For every  $1 \leq j \leq k$ ,  $t_1^j + \gamma_2 < t_2^j$ , and (4)  $2nc(\delta + \gamma_1/n) < \delta + \gamma_2/n$ , and this follows by basic algebra.

Now consider the following sequence of instances: (a) machine 1 declares  $t_1(\gamma_2)$  and the other machines declare as before. By construction machine 1 does not get any job and machine  $l$  gets some jobs from  $1 \dots k$ . (b) machine 1 declares  $t_1(\gamma_2)$ , machine  $l$  declares  $t_1(\gamma_1)$ , and the other machines declare as before. By weak monotonicity machine  $l$  continues to receive some of the jobs  $1, \dots, k$  since for every job  $1 \leq j \leq k$  we have  $t_l^j - (t_1^j + \gamma_1) > na$  (since  $t_2^j - t_1^j - na > a > \gamma_1$ ). Since this instance satisfies the conditions of claims 3.10 and 3.11 we have that machine  $l$  receives all jobs. In particular machine 1 does not receive any job. (c) machine 1 declares  $t_l$ , machine  $l$  declares  $t_1(\gamma_1)$ , and the other machines declare as before. By weak monotonicity machine  $l$  does not receive any job. However this contradicts the anonymity assumption since when machine 1 declares  $t_1(\gamma_1)$ , machine  $l$  declares  $t_l$ , and the other machines declare as before machine  $l$  receives some jobs.  $\square$

Now we are ready to prove Lemma 3.13. The proof is by induction on  $i$ , from  $m$  to 2. We start by considering  $I_m$ . Notice that the only difference between the instance (4) and an instance that belongs to  $I_m$  is the bigger costs of machine  $m$ . Hence, by weak monotonicity we have that machine  $m$  is allocated no jobs at all in every instance that belongs to  $I_m$ . This proves that the conditions of Claim 3.14 hold for the family of instances  $I_m$ , so machine 1 is allocated all jobs in every instance that belongs to  $I_m$ .

We now prove the induction step, i.e. we fix  $i < m$ , assume that the lemma is correct for  $I_{i+1}, \dots, I_m$ , and show correctness for  $I_i$ . We first argue that machines  $i, \dots, m$  do not receive any jobs in every instance that belongs to  $I_i$ . Suppose by contradiction that there exists an instance  $U \in I_i$

<sup>7</sup> $\delta$  plays the role of  $a$  and  $a$  belongs to the  $t_i^j$ 's.

(with corresponding cost vectors  $t_1, \dots, t_m$  and corresponding constants  $a, \delta$ ) and a machine  $l \geq i$  such that machine  $l$  gets a non-empty set of jobs in  $f(U)$ . But now decrease every  $t_l^j$  to some  $\tilde{t}_l^j$  such that  $t_{i-1}^j < \tilde{t}_l^j < t_l^j$  for all  $1 \leq j \leq k$  (note that indeed we have  $t_{i-1}^j < t_l^j$ ) and  $\tilde{t}_l^j = a$  for all  $k < j \leq n$ . By weak monotonicity since machine  $l$  got some jobs previously it still gets a non-empty set of jobs after the change. However notice that this new instance belongs to the family  $I_{i+1}$ , and machine  $l$  is not the lowest machine, contradicting the inductive hypothesis that for all instances in  $I_{i+1}$  the lowest machine gets all jobs.

Therefore we have that machines  $i, \dots, m$  do not receive jobs in every instance that belongs to  $I_i$ . Claim 3.14 now tells us that the lowest machine gets all jobs, and the inductive argument follows. This finishes the proof of Lemma 3.13.  $\square$

Lemma 3.13 immediately finishes the inductive argument of Lemma 3.3.

### 3.3 Concluding the Main Inductive Step

As a result of Lemma 3.3 (for  $k = n - 1$ ) we are now in the following situation:

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{n-1*} & \delta^* \\ t_2^1 & t_2^2 & \dots & t_2^{n-1} & t_2^n \\ & & & & \vdots \\ t_m^1 & t_m^2 & \dots & t_m^{n-1} & t_m^n \end{pmatrix} \quad (7)$$

Note that this holds for every instance  $t$  such that  $t_i^j > t_{i-1}^j$  for all  $j$  and  $i > 1$  and for all appropriately low  $\delta$  (as detailed above in Lemma 3.7.). Therefore it only remains to show is that when machine 1 increases the  $n$ 'th cost to  $t_1^n$ , it still gets all jobs, i.e. that we have

$$\begin{pmatrix} t_1^{1*} & t_1^{2*} & \dots & t_1^{n-1*} & t_1^n \\ t_2^1 & t_2^2 & \dots & t_2^{n-1} & t_2^n \\ & & & & \vdots \\ t_m^1 & t_m^2 & \dots & t_m^{n-1} & t_m^n \end{pmatrix} \quad (8)$$

The proof is very similar to the proof of Lemma 3.13. In particular, the first claim is almost identical to claim 3.10, and the second claim is almost identical to claim 3.14. We explicitly state and prove both claims for completeness.

CLAIM 3.15. *In instance (8), if machine 1 gets a non-empty set of jobs then it must get all jobs.*

PROOF. Suppose by contradiction that machine 1 gets some non-empty set of jobs  $S$  that does not contain all jobs. Decrease the costs of machine 1 for all  $j \in S$  to  $\delta$  and slightly increase the costs of machine 1 for all  $j \notin S$  to  $\tilde{t}_1^j > t_1^j$  such that  $\tilde{t}_1^j < t_2^j$ . By weak monotonicity machine 1 now gets the same set of jobs  $S$ . However this new instance fits exactly the requirement of Lemma 3.3 (for  $k = n - |S| < n$ ) and since the Lemma states that machine 1 should get all jobs we have a contradiction.  $\square$

CLAIM 3.16. *In instance (8) machine 1 must get a non-empty set of jobs.*

PROOF. Suppose by contradiction that machine 1 does not get any job. There exist two cost vectors for machines

1,  $t_1'$  and  $t_1''$ , such that for every job  $j$ ,  $t_1^j < t_1^{j'} < t_1^{j''} < t_2^j$ , and such that (1) machine 1 does not get any job in both declarations, and (2) there exists a machine  $i > 1$  such that machine  $i$  gets a non-empty set of jobs in both declarations.

Now consider the following sequence of instances: (a) Machine 1 declares  $t_1^{j''}$  and the rest of the machines declare as in  $t$ . By construction we have that machine 1 gets no jobs and machine  $i$  gets some jobs. (b) Machine 1 declares  $t_1^{j'}$ , machine  $i$  declares  $t_1^{j'}$ , and the rest of the machines declare as in  $t$ . Since in the previous instance machine  $i$  got some jobs then by weak monotonicity machine  $i$  continues to get some jobs (not necessarily the same set of jobs). However this instance is equivalent to the instance of (8), where machine  $i$  is the lowest machine, and therefore by claim 3.15 we get that machine  $i$  gets all jobs. In particular machine 1 gets no jobs. (c) Machine 1 declares  $t_i$ , machine  $i$  declares  $t_1$ , and the rest of the machines declare as in  $t$ . Since in the previous instance machine 1 did not get any job then by weak monotonicity machine 1 does not get any job here as well. However notice that this last instance (c) is the same as the instance  $t' = (t_1, t_{-1})$  defined above, except that machine 1 and machine  $i$  have switched costs. By anonymity since in  $t'$  machine  $i$  gets some jobs then in the last instance (c) machine 1 should get some jobs, however earlier we have argued that it does not get any jobs in instance (c), a contradiction.  $\square$

As a result of the two last claims we get that in instance (8) machine 1 gets all jobs, which completes the main inductive argument.

### Acknowledgments

We thank Michal Feldman, Amos Fiat, Haim Kaplan, and Svetlana Olonetsky for helpful discussions. We are also grateful to the anonymous referees of EC for their many helpful comments.

## 4. REFERENCES

- [1] Gagan Aggarwal, Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, Nicole Immorlica, and Madhu Sudan. Derandomization of auctions. In *STOC'05*.
- [2] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *FOCS'01*.
- [3] G. Christodoulou, E. Koutsoupias, and A. Vidali. A characterization of 2-player mechanisms for scheduling. In *ESA*, September 2008.
- [4] George Christodoulou, Elias Koutsoupias, and Annamária Kovács. Mechanism design for fractional scheduling on unrelated machines. In *ICALP'07*.
- [5] George Christodoulou, Elias Koutsoupias, and Angelina Vidali. A lower bound for scheduling mechanisms. In *SODA '07*.
- [6] Constantinos Daskalakis and Christos H. Papadimitriou. Computing equilibria in anonymous games. In *FOCS*, pages 83–93, 2007.
- [7] Constantinos Daskalakis and Christos H. Papadimitriou. Discretized multinomial distributions and nash equilibria in anonymous games. In *FOCS*, 2008.
- [8] Peerapong Dhangwatnotai, Shahar Dobzinski, Shaddin Dughmi, and Tim Roughgarden. Truthful approximation schemes for single-parameter agents. In *FOCS'08*.
- [9] Shahar Dobzinski and Mukund Sundararajan. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *EC'08*.
- [10] Elias Koutsoupias and Angelina Vidali. A lower bound of  $1+\phi$  for truthful scheduling mechanisms. In *MFCS'07*.
- [11] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *FOCS'03*.
- [12] Ron Lavi and Chaitanya Swamy. Truthful mechanism design for multi-dimensional scheduling via cycle monotonicity. In *EC'07*.
- [13] Ahuva Mu'alem and Michael Schapira. Setting lower bounds on truthfulness: extended abstract. In *SODA*, 2007.
- [14] Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *STOC*, 1999.
- [15] M. E. Saks and L. Yu. Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 286–293, 2005.