

On the Nisan-Ronen conjecture

George Christodoulou
University of Liverpool
gchristo@liverpool.ac.uk

Elias Koutsoupias
University of Oxford
elias@cs.ox.ac.uk

Annamária Kovács
Goethe University Frankfurt
panni@cs.uni-frankfurt.de

Abstract—The Nisan-Ronen conjecture states that no truthful mechanism for makespan-minimization when allocating m tasks to n unrelated machines can have approximation ratio less than n . Over more than two decades since its formulation, little progress has been made in resolving it and the best known lower bound is still a small constant. This work makes progress towards validating the conjecture by showing a lower bound of $1 + \sqrt{n-1}$.

Keywords—mechanism design; algorithmic game theory;

I. INTRODUCTION

This work makes progress on one of the most important open problems of algorithmic mechanism design [26], the Nisan-Ronen conjecture. Mechanism design, a celebrated branch of Game Theory and Microeconomics, studies a special class of algorithms, called mechanisms, which are robust under selfish behavior and produce a social outcome with a certain guaranteed quality. Unlike traditional algorithms that get their input from a single user, mechanisms solicit the input from different participants (called agents, players, bidders), in the form of preferences over the possible outputs (outcomes). The challenge stems from the fact that the actual preferences of the participants are private information, unknown to the algorithm. The participants are assumed to be utility maximisers that will provide some input that suits their objective and may differ from their true preferences. A *truthful* mechanism provides incentives such that a truthful input is the best action for each participant.

The question is what kind of problems can be solved within this framework. In their seminal paper that launched the field of algorithmic mechanism design, Nisan and Ronen [25] proposed the scheduling problem on unrelated machines as a central problem to capture the algorithmic aspects of mechanism design. In the classical form of the problem, which has been extensively studied from the algorithmic perspective, there are n machines that process a set of m tasks; each machine i takes time t_{ij} to process task j . The objective of the algorithm is to allocate each task to a machine in order to minimize the makespan, i.e., the maximum completion time of all machines. In the mechanism design setting, each machine provides as input its processing times for each task. The selection by Nisan and Ronen of this version of the scheduling problem to study the limitations that truthfulness imposes on algorithm design

was a masterstroke, because it turned out to be an extremely rich and challenging setting.

Nisan and Ronen applied the VCG mechanism [28], [12], [18], the most successful generic machinery in mechanism design, which truthfully implements the outcome that maximizes the social welfare. In the case of scheduling, the allocation of the VCG is the greedy allocation: each task is independently assigned to the machine with minimum processing time. This mechanism is truthful, but has poor approximation ratio, n . They boldly conjectured that this is the best guarantee that can be achieved by any deterministic (polynomial-time or not) truthful mechanism.

Conjecture 1 (Nisan-Ronen). *There is no deterministic truthful mechanism with approximation ratio better than n for the problem of scheduling n unrelated machines.*

The conjectured bound refers only to the limitation that truthfulness imposes and does not preclude any computational limitations. In that sense, this is a very strong information-theoretic bound which should hold for all deterministic mechanisms, regardless of their running time.

The Nisan-Ronen conjecture is perhaps the most famous problem and, arguably, one of the most important open ones in algorithmic mechanism design. Despite intensive efforts, very sparse progress has been made towards its resolution. All the known lower bounds are small constants; Nisan and Ronen originally showed that no truthful deterministic mechanism can achieve an approximation ratio better than 2. This was improved to 2.41 [11], and later to 2.61 [19], which was the current best for over a decade. Very recently the bound was improved to 2.755 by Giannakopoulos, Hamerl and Poças [16] and then to 3 by Dobzinski and Shaulker [14]. Still the gap from the best known upper bound of n remains huge.

In this work we take a solid step towards the resolution of the conjecture by providing a drastic improvement on the lower bound.

Theorem 2. *There is no deterministic truthful mechanism with approximation ratio better than $1 + \sqrt{n-1}$ for the problem of scheduling n unrelated machines.*

The proof of the theorem¹ is presented in Section III. At its core, it is based on the use of a known characterization of 2×2 instances, with only two machines and 2 tasks [8]. Roughly speaking, this characterization says essentially that the only available algorithms for 2×2 instances are either task independent (algorithms that allocate each task independently of the values of the other tasks) or affine minimizers (algorithms that compare simple linear expressions of the input values to allocate the tasks). This is an extremely limited class of algorithms, a fact that provides some weak evidence that the Nisan-Ronen conjecture may be true. No such characterization is known for 3 or more machines, and a widely-held belief has been that attacking the Nisan-Ronen conjecture requires further progress in this direction of characterizing truthful mechanisms for n machines. It is therefore quite surprising that one can obtain a lower bound of $1 + \sqrt{n-1}$ based mainly on the characterization of 2×2 instances.

Another cornerstone of the proof of the main theorem is the extensive use of weak monotonicity, the characterizing property of truthful allocation functions. It is well known [27], [1] that a mechanism is truthful if its allocation function is monotone in the values of each machine. Weak monotonicity in one dimension (i.e., a single task) is the usual notion of monotonicity of the allocation function, and for two or more dimensions, it takes a particular very natural form. Thus one can restate the main theorem as “no monotone algorithm, polynomial-time or not, has approximation ratio less than $1 + \sqrt{n-1}$ for the problem of scheduling n unrelated machines”. In contrast, the approximation ratio for the usual (non-monotone) class of algorithms is trivially 1, for exponential-time algorithms, and 2 for polynomial-time ones [21]. To use weak monotonicity, we consider perturbations of an input that lie in a small open hyperrectangle. The precise definition of these perturbations needs to balance different requirements of the proof, but we believe that the general idea can be useful in other similar situations.

Previous attempts for resolving the Nisan-Ronen conjecture employed a characterization of 2×2 instances and, of course, weak monotonicity. What distinguishes this work from these previous attempts, with the exception of [8], is the use of particular inputs with some very high values that force any reasonable algorithm to allocate each task to only two machines. These inputs avoid the problems that weak monotonicity presents for multiple players: with such inputs, when one machine does not take a task, we know precisely which machine takes it.

Some ideas of this work, such as the use of the 2×2 characterization on particular inputs, first appeared in a recent publication [8]. The main result in [8] is a lower bound of $\sqrt{n-1}$ for all deterministic truthful mechanisms,

when the cost of processing a subset of tasks is given by a submodular (or supermodular) set function, instead of an additive function of the standard scheduling setting.

The instances employed in [8] for submodular costs use only $n-1$ pairs of tasks. But here, where costs are additive, we use exponentially many tasks. We don’t know whether this large number of tasks is needed in general, but it is essential in our proof. It definitely needs to be $\Omega(n^{3/2})$ for the type of instances that we use, since for fewer tasks we know truthful mechanisms with approximation ratio less than $\sqrt{n-1}$ [9].

The real challenge, not present in [8], is that the characterization of 2×2 truthful mechanisms for additive costs includes mechanisms which are difficult to handle (for example, relaxed task independent mechanisms and relaxed affine minimizers). This obstacle adds many complications and requires a different proof structure.

The major open problem left open is to settle the Nisan-Ronen conjecture. The techniques of this work may be helpful in this direction. The case of randomized or fractional mechanisms appears to be more challenging; the best known lower bound of the approximation ratio is 2 [24], [7], embarrassingly lower than the best known upper bound $(n+1)/2$ [7]. The bottleneck of applying the techniques of the current work to these variants appears to be the lack of a good characterization of 2×2 mechanisms. Finally, although the result of this work indicates that mechanisms constitute a limited subclass of allocation algorithms, a more direct demonstration would be to find other useful properties of mechanisms, or even obtain a more global characterization of mechanisms for the domain of scheduling and its generalizations.

A. Further related work

The lack of progress in the original unrelated machine problem led to the study of variants and special cases, for which significant results have been obtained. Ashlagi et al. [2], resolved a restricted version of the Nisan-Ronen conjecture, for the special but natural class of *anonymous* mechanisms. Lavi and Swamy [20] studied a restricted input domain which however retains the multi-dimensional flavour of the setting. They considered inputs with only two possible values “low” and “high”, that are publicly known to the designer of the algorithm. For this case they showed an elegant deterministic mechanism with an approximation factor of 2. They also showed that even for this setting achieving the optimal makespan is not possible under truthfulness, and provided a lower bound of 11/10. Yu [30] extended the results for a range of values, and Auletta et al. [3] studied multi-dimensional domains where the private information of the machines is a single bit.

Randomized mechanisms have also been studied and have slightly improved guarantees. There are two notions of truthfulness for randomized mechanisms; a mechanism is *uni-*

¹Due to space limitations most proofs are omitted. We refer the reader to the full version <https://arxiv.org/abs/2011.14434>.

versally truthful if it is defined as a probability distribution over deterministic truthful mechanisms, and it is *truthful-in-expectation*, if in expectation no player can benefit by lying. In [25], a universally truthful mechanism was proposed for the case of two machines, and was later extended to the case of n machines by Mu'alem and Schapira [24] with an approximation guarantee of $0.875n$, which was later improved to $0.837n$ by [22]. Lu and Yu [23] showed a truthful-in-expectation mechanism with an approximation guarantee of $(n+5)/2$. Mu'alem and Schapira [24], showed a lower bound of $2 - 1/n$, for both types of randomized truthful mechanisms. Christodoulou, Koutsoupias and Kovács [7] extended the lower bound for fractional mechanisms, where each task can be fractionally allocated to multiple machines. They also showed a fractional mechanism with a guarantee of $(n+1)/2$. Even for the special case of two machines the upper and lower bounds are close but not tight [23], [6].

In the Bayesian setting, Daskalakis and Weinberg [13] showed a mechanism that is at most a factor of 2 from the *optimal truthful mechanism*, but not with respect to optimal makespan. Chawla et al. [5] provided bounds of prior-independent mechanisms (where the input comes from a probability distribution unknown to the mechanism). Giannakopoulos and Kyropoulou [17] showed that the VCG mechanism achieves an approximation ratio of $O(\log n / \log \log n)$ under some distributional and symmetry assumptions. Finally, in a recent work, Christodoulou, Koutsoupias and Kovács [9] showed positive results for settings where each task can only be allocated to at most two machines, a property shared by the instances of the lower bound of this paper.

II. PRELIMINARIES

In the unrelated machines scheduling problem, we are given a set M of m tasks that need to be scheduled on a set N of n machines. The processing time or cost that each machine i needs to process task j is t_{ij} , and the completion time of machine i for a subset S of tasks is equal to the sum of the individual task costs $t_i(S) = \sum_{j \in S} t_{ij}$. The objective is to find an allocation of tasks to machines that minimize the *makespan*, that is, the maximum completion time of a machine. We sometimes refer to the cost of a set of tasks S , as the optimal makespan with respect to S .

A. Mechanism design setting

We assume that each machine $i \in N$ is controlled by a selfish agent (player) that is reluctant to process tasks and the costs t_{ij} is private information known only to them (also called the *type* of agent i). The set \mathcal{T}_i of possible types of agent i consists of all vectors $b_i = (b_{i1}, \dots, b_{im}) \in \mathbb{R}_+^m$. Let also $\mathcal{T} = \times_{i \in N} \mathcal{T}_i$ be the space of type profiles.

A mechanism defines for each player i a set \mathcal{B}_i of available strategies the player can choose from. We consider *direct revelation* mechanisms, i.e., $\mathcal{B}_i = \mathcal{T}_i$ for all i , meaning

that the players strategies are to simply report their types to the mechanism. Each player i provides a *bid* $b_i \in \mathcal{T}_i$, which not necessarily matches the true type t_i , if this serves their interests. A mechanism (A, \mathcal{P}) consists of two parts:

An allocation algorithm: The allocation algorithm A allocates the tasks to machines based on the players' inputs (bid vector) $b = (b_1, \dots, b_n)$. Let \mathcal{A} be the set of all possible partitions of m tasks to n machines. The allocation function $A : \mathcal{T} \rightarrow \mathcal{A}$ partitions the tasks into the n machines; we denote by $A_i(b)$ the subset of tasks assigned to machine i for bid vector $b = (b_1, \dots, b_n)$.

A payment scheme: The payment scheme $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_n)$ determines the payments, which also depends on the bid vector b . The functions $\mathcal{P}_1, \dots, \mathcal{P}_n$ stand for the payments that the mechanism hands to each agent, i.e., $\mathcal{P}_i : \mathcal{T} \rightarrow \mathbb{R}$.

The *utility* u_i of a player i is the payment that they get minus the *actual* time that they need to process the set of tasks assigned to them, $u_i(b) = \mathcal{P}_i(b) - t_i(A_i(b))$. We are interested in *truthful* mechanisms. A mechanism is truthful, if for every player, reporting his true type is a *dominant strategy*. Formally,

$$u_i(t_i, b_{-i}) \geq u_i(t'_i, b_{-i}), \quad \forall i \in [n], t_i, t'_i \in \mathcal{T}_i, b_{-i} \in \mathcal{T}_{-i},$$

where notation \mathcal{T}_{-i} denotes all parts of \mathcal{T} except its i -th part.

The quality of a mechanism for a given type t is measured by the makespan $\text{MECH}(t)$ achieved by its allocation algorithm A , $\text{MECH}(t) = \max_{i \in N} t_i(A_i(t))$, which is compared to the optimal makespan $\text{OPT}(t) = \min_{A \in \mathcal{A}} \max_{i \in N} t_i(A_i)$.

It is well known that only a subset of algorithms can be allocation algorithms of truthful mechanisms. In particular, no mechanism's allocation algorithm is optimal for every t , so it is natural to focus on the approximation ratio of the mechanism's allocation algorithm. A mechanism is *c*-approximate, if its allocation algorithm is *c*-approximate, that is, if $c \geq \frac{\text{MECH}(t)}{\text{OPT}(t)}$ for all possible inputs t .

In this work, we do not place any requirements on the time to compute the mechanism's allocation A and payments \mathcal{P} . In other words, the lower bound does not make use of any computational assumptions.

B. Weak monotonicity

Mechanisms consist of two algorithms, the allocation algorithm and the payment algorithm. However, we are only interested in the performance of the allocation algorithm. Therefore it is natural to ask whether it is possible to characterize the class of allocation algorithms that are part of a truthful mechanism with no reference to the payment algorithm. Indeed the following definition provides such a characterization.

Definition 3. An allocation algorithm A is called *weakly monotone (WMON)* if it satisfies the following property:

for every two inputs $t = (t_i, t_{-i})$ and $t' = (t'_i, t_{-i})$, the associated allocations A and A' satisfy

$$t_i(A_i) - t_i(A'_i) \leq t'_i(A_i) - t'_i(A'_i).$$

An equivalent condition, using $a_{ij}(t) \in \{0, 1\}$ indicator variables of whether task j is allocated to player i , is

$$\sum_{j \in M} (a_{ij}(t') - a_{ij}(t))(t'_{ij} - t_{ij}) \leq 0.$$

It is well known that the allocation function of every truthful mechanism is weakly monotone [4]. Although it is not needed in establishing a lower bound on the approximation ratio, it turns out that this is also a sufficient condition for truthfulness in convex domains [27] (which is the case for the scheduling domain).

A useful tool in our proof relies on the following immediate consequence of weak monotonicity (see [8] for a simple proof). Intuitively, it states that when you fix the costs of all players for a subset of tasks, then the *restriction* of the allocation to the rest of the tasks is still weakly monotone.

Lemma 4. *Let A be a weakly monotone allocation, and let (S, T) a partition of M . When we fix the costs of the tasks of T , the restriction of the allocation A on S is also weakly monotone.*

The following implication of weak monotonicity, that was first used in [25], is a standard tool for showing lower bounds for truthful mechanisms (see for example [25], [11], [24], [2], [16], [14]).

Lemma 5. *Consider a truthful mechanism (A, \mathcal{P}) and its allocation for a bid vector t . Let S be a subset of the tasks allocated to player i and let S' be a subset of the tasks allocated to the other players. Consider any bid profile $t' = (t'_i, t_{-i})$ that is obtained from t by decreasing the values of player i in S , i.e., $t'_{ij} < t_{ij}, j \in S$, and increasing the values of player i in S' , i.e., $t'_{ij} > t_{ij}, j \in S'$. Then the allocation of player i for t and t' agree for all tasks in $S \cup S'$.*

Notice that this lemma guarantees that only the set S of tasks allocated to player i remains the same. This does not preclude changing the allocation of the other players for the tasks in S' , unless there are only two players. This is a major obstacle in multi-player settings, that we avoid in this work by focusing on only two players for each task.

C. Characterization of truthful mechanisms for 2 players and 2 tasks.

In this section we discuss the class of truthful 2×2 mechanisms (i.e., for 2 machines and 2 tasks). We present here a full characterization of these mechanisms as it appears in the full version of [8].

There are other such characterizations in the literature (for example [15], [10]), but they are somewhat incomplete, in the sense that they either impose additional restrictions

to mechanisms (for example, bounded approximation ratio) or consider a bigger domain with negative values. These characterizations do not fit the purpose of this work. For example, we deal with domains with multiple players and tasks and conditions on the approximation ratio of restrictions to only 2 tasks are not applicable.

An important aspect of our approach is to restrict the allocation of each task to only two specific machines, which we call player (machine) t and player s . To achieve this, we fix two large constants Θ and B with $\Theta \gg B \gg 1$, and consider instances in which the value for the t -player can be arbitrary, for the s player in $[0, B)$, and for the other players equal to Θ . Since the values of the s -player are much smaller than Θ , if the mechanism allocates even one task to a player with value Θ , its approximation ratio is really high.

A convenient aspect of the characterization of [8], is that it characterizes the 2×2 mechanisms with values exactly in this domain. We first present the theorem and then define and discuss the different types of mechanisms. The characterization assumes that for sufficiently high values of t -player, both tasks will be allocated to the s -player, otherwise the approximation ratio is unbounded, even when there are more players and tasks.

Theorem 6 (Characterization of 2×2 mechanisms [8]). *For every $B \in \mathbb{R}_{>0} \cup \{\infty\}$, every weakly monotone allocation for two tasks and two additive players with bids $t \in [0, \infty) \times [0, \infty)$ and $s \in [0, B) \times [0, B)$ —such that for every s there exists t for which both tasks are allocated to the second player— belongs to the following classes: (1) relaxed affine minimizers (including the special case of affine minimizers), (2) relaxed task independent mechanisms (including the special case of task independent mechanisms), (3) 1-dimensional mechanisms, (4) constant mechanisms.²*

Note that the theorem applies also to unbounded domain, i.e., the case of $B = \infty$. Some mechanisms can be of two types; for example the VCG mechanism is both an affine minimizer and task independent.

To discuss the different types of mechanisms of the theorem, let (A, \mathcal{P}) be a truthful 2×2 mechanism, where A is a weakly monotone allocation function, and \mathcal{P} denotes the payment function. For given fixed $s \in [0, B) \times [0, B)$ let us consider the allocation for the t -player depending on his own bids (t_1, t_2) (see Figure 1). The allocation regions R_{12} (resp. R_1, R_2, R_\emptyset) $\subseteq \mathbb{R}_{\geq 0}^2$ are defined to be the interior (wrt. $\mathbb{R}_{\geq 0}^2$) of the set of all t values such that the t -player gets the set of tasks $\{1, 2\}$ (resp. $\{1\}, \{2\}, \emptyset$). The allocation of the points on the boundaries can take any allocation of the adjacent regions.

It is known that in the case of two tasks, for a fixed s , these

²To be precise, the restriction of every weakly monotone allocation to strictly positive t and s is of type (1) to (4); on the border 0 there might be singularities, see [8] for details. For this reason, in the main argument we use strictly positive t and s .

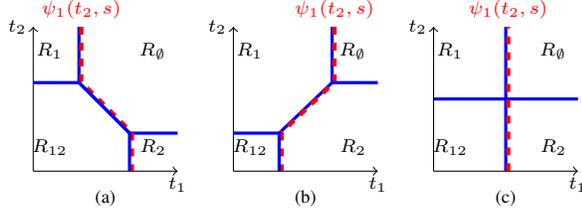


Figure 1. The allocation to the t -player depending on their own bid vector (t_1, t_2) for fixed values $s = (s_1, s_2)$ of the other player: (a) quasi-bundling allocation; (b) quasi-flipping allocation; (c) crossing allocation. The boundary $\psi_1(t_2, s)$ for task 1 is shown by a dashed line.

regions in a weakly monotone allocation subdivide $\mathbb{R}_{\geq 0}^2$ basically in three possible ways (as in Figure 1), which are characteristic for the type of the whole allocation-function A . Similar subdivisions for multiple tasks exist in higher dimensions [10], [29].

Definition 7. The allocation of a mechanism is defined by its boundary functions $\psi_i(t_{-i}, s)$, the infimum of the values of task t_i for which the t -player does not get task i . In other words, the t -player gets task i when $t_i < \psi_i(t_{-i}, s)$ and it does not get it when $t_i > \psi_i(t_{-i}, s)$.

For a given s , we call the allocation for the t -player

- *quasi-bundling*, if there are at least two points $t \neq t'$ on the boundary of R_{12} and R_\emptyset . In this case, we say that the boundary between R_{12} and R_\emptyset is a *bundling boundary*.
- *quasi-flipping*, if there are at least two points $t \neq t'$ on the boundary of R_1 and R_2 .
- *crossing*, otherwise. (See Figure 1 for an illustration.)

Relaxed affine minimizers.: An allocation is an *affine minimizer*, if there exist positive constants λ', λ and constants $\pi_a \in \mathbb{R} \cup \{-\infty, \infty\}$ for each allocation $a \in \{1, 2, \emptyset\}$, so that for every input (t, s) the allocation of the t -player minimizes the corresponding expression among the following:

$$\lambda'(t_1 + t_2) + \pi_{12}, \quad \lambda' t_1 + \lambda s_2 + \pi_1, \quad (1)$$

$$\lambda' t_2 + \lambda s_1 + \pi_2, \quad \lambda(s_1 + s_2) + \pi_\emptyset. \quad (2)$$

An affine minimizer allocation has typically either the form of Figure 1.a or 1.b.

Since we consider nonnegative values, the domain is bounded below. When an affine minimizer has the bundling form (Figure 1.a), it may be the case that for small $s_1 + s_2$ and $t_1 + t_2$, the mechanism locally ‘looks like’ a bundling mechanism, with only the regions R_\emptyset and R_{12} for both players. If this is the case for *all* (s_1, s_2) with $s_1 + s_2 < D_s$ and *all* (t_1, t_2) such that $t_1 + t_2 < D_t$ for some given $D_t, D_s > 0$, then the mechanism becomes locally a bundling mechanism. As such, the boundary functions for the allocation need not remain linear for these small

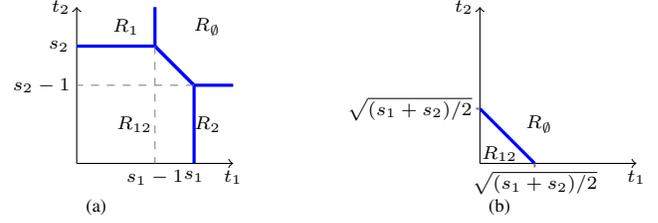


Figure 2. An example of a relaxed affine minimizer, which shows the allocation of the t -player for two distinct values of s . The left figure shows the allocation for some $s_1 \geq 1$ and $s_2 \geq 1$; in this case, the boundaries are linear; for example, when $t_2 \in [s_2 - 1, s_2]$ the allocation for task 1 is determined by comparing t_1 to $\psi_1(t_2, s) = s_1 + s_2 - 1$. The right figure shows the allocation at the bundling tail for some $s_1, s_2 \leq 1$, in which case the boundary does not have to be linear; for example for $t_2 \in [0, \sqrt{(s_1 + s_2)/2}]$, we have $\psi_1(t_2, s) = \sqrt{(s_1 + s_2)/2}$.

s and t and the boundary between R_{12} and R_\emptyset is given by $t_1 + t_2 = \zeta(s_1 + s_2)$, for some non-decreasing function ζ . See Figure 2 for an example. For a formal definition of these mechanisms, which are called *relaxed affine minimizers*, we refer the reader to [8], where they were introduced.

We will refer to the (potentially) non-linear bundling allocation of a relaxed affine minimizer as its ‘bundling tail’. In other words, the bundling tail of a relaxed affine minimizer is determined by the values of s for which the allocation of the t -player contains only the regions R_{12} and R_\emptyset . Note that affine minimizers are simply the special subclass of these mechanisms in which the bundling tail is missing or conforms to (1).

The important aspect of affine minimizers and of relaxed affine minimizers outside their bundling tail is that the boundary $\psi_i(t_{-i}, s)$ is a truncated linear function in s_i , and in particular of the form

$$\psi_i(t_{-i}, s) = \max(0, \lambda s_i - \gamma(t_{-i}, s_{-i})),$$

where λ is a constant and $\gamma(t_{-i}, s_{-i})$ a function that does not depend on the values of task i . Note that when we consider a 2×2 mechanism as a restriction of a mechanism with multiple players and tasks, the coefficient λ is not an absolute constant but it can depend in an arbitrary way on values outside the 2×2 slice (see Definition 9 for the definition of slice); similarly for the function γ .

Relaxed task independent allocations.: An allocation function A is *task independent* if for both tasks the allocation of task i depends only on the input values s_i and t_i . For the t -player, the boundary of t_1 , i.e., the lowest value, above which t_1 does not get task 1, is determined by an *arbitrary non-decreasing function* $\psi_1 : [0, B) \rightarrow [0, \infty)$ of s_1 and analogously for the boundary $\psi_2(s_2)$. Geometrically, in a task independent mechanism, the allocations of both players are always crossing (Figure 1.c). In a *relaxed task independent* mechanism the latter property is fulfilled in all

but countably many s (resp. t) points, in which both ψ_1 and ψ_2 (resp. ψ_1^{-1} and ψ_2^{-1}) have a jump discontinuity. See [8] for a formal definition. *The important property that we use here is that every relaxed task independent allocation is identical with a task independent allocation on $t_i \in [0, \infty) \setminus T_i$, $s_i \in [0, B) \setminus S_i$, where the T_1, T_2, S_1, S_2 are countable sets.*

1-dimensional mechanisms.: In a 1-dimensional mechanism at most two possible allocations are ever realized. If the two occurring allocations are \emptyset and 12, we call the mechanism *bundling mechanism*. One can consider bundling 1-dimensional mechanisms as degenerate *relaxed affine minimizers* with $\pi_1 = \pi_2 = \infty$.

The other cases when the allocations to the t -player are \emptyset and 1 (or \emptyset and 2) are degenerate *task independent* allocations, and for our purposes they can be treated as task independent.

Constant mechanisms.: In a *constant or dictatorial mechanism* the allocation is independent of the bids of at least one of the players. This property can also be interpreted as being an affine minimizer with a multiplicative constant $\lambda = 0$.

III. LOWER BOUND

In this section, we give a proof of our main result. In Sections III-A and III-B, we describe the general setting, provide important definitions, and set the main goal. With these at hand, we outline the proof in Section III-C. Then in Sections III-D, III-E, and III-F we provide the proofs of all the technical lemmas needed in order to establish the main theorem, the proof of which is presented in the last subsection (Section III-G).

A. The construction

We consider instances, with n players and $m = (\ell + 1)(n - 1) + n$, for some large ℓ to be determined later. Player 0 is special and for convenience we use the symbol t for its values; sometimes we refer to it as the t -player or 0-player. We use the symbol s for the values of the remaining players $1, \dots, n - 1$, and sometimes we refer to them as the s -players.

There are n tasks d_0, \dots, d_{n-1} , which are special that are called *dummy tasks*; they play a limited role in the proof and their only purpose is to increase the lower bound from $\sqrt{n-1}$ to $1 + \sqrt{n-1}$.

The remaining tasks are partitioned into $n - 1$ clusters C_1, \dots, C_{n-1} , where each cluster C_i contains $\ell + 1$ tasks and is associated with player $i \in [n - 1]$. We call two tasks j and j' that belong to the same cluster *siblings*.

Definition 8 (Range of input values). The processing times for a task $j \in C_i$, $i \in [n - 1]$, is described by two values t_j and s_j , as follows:

- player 0 has processing time $t_j \in (0, \infty)$; with very few exceptions, the argument uses values $t_j \in (0, 1]$.

- player i has processing time $s_j \in (0, B)$, for some B ; again with very few exceptions, the argument uses values $s_j \in (0, 1]$.
- every other player $k \notin \{0, i\}$ has processing time Θ , for some Θ .

Dummy task d_i has value Θ for all players except for player i for which the value is initially 0.

The values of B and Θ are arbitrarily large functions of n (exponential functions suffice to get a good approximation ratio) with $\Theta/((\ell + 1)B) \gg n$.

Since every task has at least one processing time in $[0, B)$, no algorithm with approximation ratio less than $\Theta/((\ell + 1)B) \gg n$, allocates any task to players with value Θ .

An instance (input) T is described only by two values t_j and s_j per task j . Let's denote by t , and s the respective vectors, hence, $T = (t, s) = (t_j, s_j)_{j \in [m]}$.

B. Definitions

We use the following fixed values throughout this section.

- $\alpha = 1/\sqrt{n-1}$
- β , an arbitrarily small positive value
- δ , a small value; think of this as n^{-2} ; any value $o(n^{-3/2})$ gives lower bound $1 + (1 - o(1))\sqrt{n-1}$. For simplicity, we will assume that δ is selected so that $2n/\delta$ is an integer.
- $\delta' = 2\delta$, upper bound on the cost of a trivial cluster (see definition of a trivial cluster below)
- ρ , the targeted lower bound on the approximation ratio, which is given by

$$\rho = 1 - \delta' + \min \left\{ \frac{1}{\alpha + (n-1)\delta'}, \frac{(n-1)\alpha}{1 + (n-1)\delta'} \right\} \\ = 1 + (1 - o(1))\sqrt{n-1},$$

- $\ell + 1$ number of tasks per cluster; exponential in n , and greater than $1 + 3n^3 \left(\frac{3n}{\delta}\right)^{n-3} = (n/\delta)^{\Theta(n)}$.

The proof is based on the characterization of 2×2 mechanisms. To be able to use it, we fix all other values except for the values of two tasks and then use the characterization. The next definition formalizes this.

Definition 9 (Slice). Fix an instance T and two tasks p and p' . The set of instances that agree with T on all tasks except for the tasks p and p' is called a (p, p') -slice for T or simply slice of tasks p and p' . A slice may involve values of 3 different players, if p, p' belong to different clusters, or values of 2 players if p, p' are siblings. In the latter case, the allocation of these tasks (due to Lemma 4) is defined by a 2×2 mechanism which we call (p, p') -slice mechanism for T .

To take advantage of weak monotonicity, we need to consider perturbations of instances, that is, instances that differ by a small amount from a given one. Here is a precise definition of the perturbations that we use.

Definition 10 (Perturbations of an instance). Fix an instance $T = (t, s)$ and a set of tasks P . Let V be a vector of open intervals V_j , one for each task $j \in P$, such that

- $V_j = (t_j, t_j + \theta_j)$, for every task $j \in P$, for some $\theta_j \in (0, \beta)$

The set of instances $T' = (t', s)$ with $t'_j \in V_j$, when $j \in P$, and $t'_j = t_j$ when $j \notin P$, is called a set of V -perturbations of T for tasks P or simply set of V -perturbations of T , when P is understood from the context.

Note that we consider perturbations only of t -values and only in one direction (towards higher values). The values of instance T itself are *not* in the perturbation for tasks in P .

Perturbations of an instance T satisfy a few important properties:

- for any given allocation, the cost of all instances of a perturbation is almost the same as the cost of T (within $(n-1)|P|\beta$).
- they allow us to select points not on boundaries of the mechanism, thus when we apply weak monotonicity for the 0-player, we can guarantee that certain allocations do not change.

The central part of the argument is an induction on the number k of clusters. The values of the tasks in the remaining $n - k - 1$ clusters, which we call trivial clusters, play a limited role, but it is important that they do not affect substantially the approximation ratio. We allow their values to be arbitrary (within the limits of Definition 8), but we require that one of the values is very small in the following sense.

Definition 11 (Trivial cluster). A cluster is called *trivial* for a given instance T if the optimal allocation for all tasks of the cluster has cost at most δ' .

We usually select a single task from each non-trivial cluster. We now fix the terminology for such sets of tasks:

Definition 12 (Regular set of tasks). A set of tasks is called *regular* if they are from different clusters.

We make repeated use of the following set of instances:

Definition 13 (Standard instance and $\hat{T}(P)$). An instance T is *standard* for a set of clusters \mathcal{C} if the following conditions hold

- the value of every task $j \in \mathcal{C}$ is $[t_j = \beta, s_j = 1]$, and
- the remaining clusters are trivial

We also say that T is *standard for a regular set of tasks* $P = \{p_1, \dots, p_k\}$, if it is standard for the set of clusters that intersect P . We denote by $\hat{T}(P)$ the instance that agrees with T everywhere except for tasks in P for which $[\hat{t}_{p_i} = \alpha, \hat{s}_{p_i} = 1]_{i=1}^k$. See Figure 3 for an illustration.

The following definition is at the heart of the proof. Roughly speaking, the aim of the proof is to show by

induction that there exist good sets of $n - 1$ tasks, otherwise the mechanism has high approximation ratio.

Definition 14 (Good set of tasks). Fix a mechanism, a set of regular tasks $P = \{p_1, \dots, p_k\}$ from a set of clusters \mathcal{C} , and a standard instance T for P . The set of tasks P is called *good for instance* T if there exists a vector V of open intervals for the tasks in P such that the mechanism allocates all tasks in P to the 0-player for every instance in the set of V -perturbations of $\hat{T}(P)$ for tasks in P . We call V the witness of goodness of P . (See Figure 3 for an illustration.)

If no such V exists, we call P a *bad set*. A singleton bad set will be simply called *bad task*. For technical reasons, we will also call a task bad when it has all the above properties of a bad task, but its s -value is $1 - \epsilon$, for some arbitrarily small $\epsilon > 0$.³

$$T = \begin{bmatrix} \beta & \beta & \beta & \beta & \beta & \beta & \delta & \beta & \beta \\ 1 & 1 & 1 & & & & & & \\ & & & 1 & 1 & 1 & & & \\ & & & & & & 1/2 & 1 & 1 \\ \alpha & \beta & \beta & \beta & \alpha & \beta & \delta & \beta & \beta \\ 1 & 1 & 1 & & & & & & \\ & & & 1 & 1 & 1 & & & \\ & & & & & & 1/2 & 1 & 1 \end{bmatrix},$$

$$\hat{T}(\{1, 5\}) = \begin{bmatrix} \alpha & \beta & \beta & \beta & \alpha & \beta & \delta & \beta & \beta \\ 1 & 1 & 1 & & & & & & \\ & & & 1 & 1 & 1 & & & \\ & & & & & & 1/2 & 1 & 1 \end{bmatrix}$$

Figure 3. T and \hat{T} are two instances with 4 players and 3 clusters, where rows correspond to players and columns to tasks. The values Θ and the dummy tasks are omitted. Typical values of the constants are $\beta \approx 0$, $\delta \approx n^{-2}$, $\alpha \approx n^{-1/2}$. T is a standard instance for clusters $\{C_1, C_2\}$. C_3 is a trivial cluster because it has cost $\delta + 2\beta \leq \delta'$. If there exists a V such that tasks 1 $\in C_1$, and 5 $\in C_2$ are assigned to the 0-player for every instance in the set of V -perturbations of $\hat{T}(\{1, 5\})$, then $\{1, 5\}$ is a good set for T .

Now that we have the definition of a set of good tasks, we can state the main lemma of this section.

Lemma 15 (Main Lemma). *At least one of the following three properties hold for every truthful mechanism:*

- the approximation ratio is at least ρ
- there exists a bad task
- there exists a good set of $n - 1$ tasks.

The aim of this section is to prove this lemma by assuming that Properties (i) and (ii) do not hold and then show Property (iii).

Assumption 16. For the rest of this section, we assume that the approximation ratio is less than

$$\rho = 1 - \delta' + \min \left\{ \frac{1}{(\alpha + (n-1)\delta')}, \frac{(n-1)\alpha}{(1 + (n-1)\delta')} \right\}$$

and that there is no bad task.

³In particular, a task will be bad if for every $\epsilon' > 0$, there is a $0 \leq \epsilon < \epsilon'$, so that setting its s -value to $1 - \epsilon$, the task has the above properties. This will serve to exclude singular points of relaxed task independent mechanisms.

Before we give the proof of the Main Lemma, we discuss how it will be used to prove a lower bound of $1 + \sqrt{n-1}$ on the approximation ratio (Section III-G contains a detailed derivation). We provide here a high-level argument. Property (i) immediately implies the lower bound. The existence of a bad task (Property (ii)) means that there is only one non-trivial task j with values $[t_j = \alpha, s_j = 1]$ and the mechanism gives it to the s -player. In this case, the approximation ratio is approximately $1/\alpha = \sqrt{n-1}$, which can be improved to $1 + \sqrt{n-1}$ using the dummy tasks. Finally, a good set of $n-1$ tasks (Property (iii)) has approximation ratio approximately $(n-1)\alpha = \sqrt{n-1}$, which again can be improved to $1 + \sqrt{n-1}$ using the dummy tasks, giving the desired result.

To obtain a proof of the Main Lemma, we will show that there exists a good set of k tasks for every $k \in [n-1]$, by induction on k . To show existence of good sets of k tasks, we start with some set of k tasks, which we call potentially-good set, such that all its subsets of $k-1$ tasks are good. To satisfy all the requirements in the proof, the precise structure of a potentially-good set of tasks is complicated and it is detailed in the following definition (see Figure 4).

Definition 17 (Potentially-good set of tasks). Fix a mechanism. A set of regular tasks $P = (p_1, \dots, p_k)$ from a set of clusters \mathcal{C} is called *potentially-good for an instance T* if T is standard for P and the following conditions hold

- for every $i \in [k]$, $P_{-i} = P \setminus \{p_i\}$ is good for T ; let V^i denote a witness of goodness of P_{-i}
- P_{-k} is good (with witness V^k) for every instance that results from T when we replace the values⁴ of task p_k with $[t_{p_k} = \delta, s_{p_k} = q\delta/(2n)]$, for $q = 0, 1, \dots, 2n/\delta$.

The witness V of potential goodness of P is defined as follows: for task p_i take the intersection of all relevant intervals in V^j , $j \neq i$.

Remark. One can replace $\delta/(2n)$ by $\delta/(2\rho)$ in the above definition to reduce the required number of $\ell + 1$ tasks per cluster, but we opt for simplicity.

C. Outline

We now give a rough outline of the argument that establishes the Main Lemma (Lemma 15). We consider standard instances T for sets of k clusters \mathcal{C} . We will show that for $k = n-1$, there is a good set of tasks. By induction on k , we show the stronger claim that there are *many sets of tasks that are good*. The base case ($k = 1$) is the assumption that all single tasks are good (Assumption 16).

We use a probabilistic argument to try to keep things simple. In particular, let $1 - b_k$ be (a lower bound on)

⁴Note that the cluster of task p_k must be trivial after the change. This is the point where we need the extra parameter $\delta' = 2\delta$. Actually any value of $\delta' \geq \delta + (\ell + 1)\beta$ works.

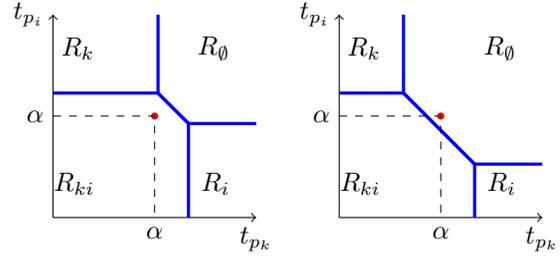


Figure 4. Allocation of the t -player in the (p_i, p_k) -slice for a regular set of tasks $P = \{p_i, p_k\}$, and some given standard instance T for P . Point (α, α) represents $\hat{T}(P)$ when P is good (left figure) and when P is potentially-good but not good (right figure).

the probability that a random regular set of tasks⁵ $P = (p_1, \dots, p_k)$ from \mathcal{C} is good for T . We want to show that the probability b_k of P being a bad set of tasks is small, and in particular that $b_{n-1} < 1$, which establishes the existence of a good set of $n-1$ tasks.

Showing that b_k is small.: We show that b_k is small by considering potentially-good sets $P = (p_1, \dots, p_k)$ of size k that contain tasks from distinct clusters (Q_1, \dots, Q_k) to establish the following two facts:

Fact a. the probability that a randomly selected set $P = (p_1, \dots, p_k)$ is potentially-good is at least $1 - (3n/\delta - 1)b_{k-1}$ (Lemma 30).

Fact b. if $P = (p_1, \dots, p_k)$ is a potentially-good set of tasks, either P is good itself or (P_{-k}, p'_k) is good with probability at least $1 - 2n^3/\ell$, where p'_k is a random sibling of p_k (Lemma 28). Roughly speaking, either P is good or *almost all other sets are good* (with exponentially small probability of the negative event).

By Assumption 16, there is no bad task, hence $b_1 = 0$, and then in Lemma 31 we show how to combine these two facts in order to get that b_k is bounded above by $3n^3/\ell(3n/\delta)^{k-2}$ (Lemma 31). If we allow the number $\ell+1$ of tasks per cluster to be sufficiently large, this shows that $b_{n-1} < 1$.

Showing Fact b.: The difficult part is to establish the second of the above two facts (Fact b). Let's assume that P is potentially-good but not good. We show that (P_{-k}, p'_k) is good for many p'_k 's, as follows

- first, we observe that there is a witness of potential goodness of P for which all tasks in $P = (p_1, \dots, p_k)$ are given to the s -players. This essentially follows from the definition of potentially-good and weak monotonicity (Lemma 20)
- let p'_k be a sibling of p_k and consider the (p_k, p'_k) -slice mechanism. This is exactly the point where we exploit the 2×2 characterization that we provide in Section II-C. The proof proceeds by treating carefully all possible cases

⁵In the argument, we select random tasks and instances. All random selections are from uniform distributions and independent.

affine minimizers: we show that the mechanism is not an affine minimizer almost surely (Lemma 22) and the stronger claim that they do not exist (Lemma ??)

relaxed affine minimizers: we show that the probability that the mechanism is a relaxed affine minimizer is at most $2n^2/\ell$ (Lemma 25)

1-dimensional and constant mechanisms: we show that 1-dimensional and constant mechanisms do not occur, otherwise the approximation ratio is high (Lemma 26)

task independent or relaxed task independent mechanisms: we show that if the mechanism is task independent or relaxed task independent for each of k appropriately selected random instances from the witness, then (P_{-k}, p'_k) is good (Lemma 27)

- we conclude that for a random sibling p'_k , the mechanism must be either task independent or relaxed task independent for all these k instances with probability at least $1 - k2n^2/\ell$; therefore (P_{-k}, p'_k) is good with probability at least $1 - 2n^3/\ell$. (Lemma 28).

The first item, i.e., to show that affine minimizers are sparse, exploits an interesting use of goodness and linearity. This is where the strange second point of Definition 17 is needed. The proof of relaxed affine minimizers uses the same machinery, but it has an extra layer of difficulty, which arises from having to guarantee that the action happens at the linear part and not at the bundling tail of these mechanisms. In fact, we might have a positive probability (at most $2n^3/\ell$) to pick a wrong sibling p'_k solely due to bundling tails of relaxed affine minimizers. The proof of the last item about task independent and relaxed task independent mechanisms has very similar flavor. It is essentially this part that takes away the complications that arise from having to deal with an additive domain.

The core of the argument for (relaxed) affine minimizers uses the characterization of the (p_k, p'_k) -slice mechanism in a way that seems peculiar at first glance. Instead of focusing on the (p_k, p'_k) -slice mechanism, it focuses on the slice (p_i, p_k) , for some $p_i \in P_{-k}$. The reason is simple: the characterization is only used to extract the property that the allocation of p_k has linear boundaries, and then taking advantage of the fact—due to potentially-goodness property—that the (p_i, p_k) -slice mechanism exhibits a tight connection between the allocation boundaries of p_i and p_k , we can argue about the allocation of task p_i . A similar argument is used when the (p_k, p'_k) -slice mechanism is task independent.

D. General lemmas

We start by establishing some useful facts.

Lemma 18. *Let $P = \{p_1, \dots, p_k\}$ be a potentially-good set of tasks for T with witness V and let T' be an instance in the set of V -perturbations of $\hat{T}(P)$. Consider a sibling p'_k of p_k and assume that the (p_k, p'_k) -slice mechanism for T' is an affine minimizer or a relaxed affine minimizer with boundary function $\psi_{p_k}(s_{p_k})$ for task p_k , which is linear when we fix the values of all other tasks i.e., $\psi_{p_k}(s_{p_k}) = \max(0, \lambda s_{p_k} - \gamma)$ for some λ and γ that may depend on all other values but not on t_{p_k} and s_{p_k} . Suppose further that $\psi_{p_k}(1) > 0$. Then the approximation ratio of the mechanism (for all tasks) is at least $1 + \max(\lambda, 1/\lambda) - \delta'$.*

We assume that the mechanism has approximation ratio less than ρ , which means that if a (p, p') -slice mechanism is an affine minimizer or a relaxed affine minimizer, the coefficient λ of its linear boundary is in the interval $[1/(\rho - 1 + \delta'), \rho - 1 + \delta']$. This guarantees that for some task p_k , if the s_{p_k} changes by Δs_{p_k} , then the boundary for t_{p_k} changes by at least $\Delta s_{p_k}/(\rho - 1 + \delta')$ and at most $(\rho - 1 + \delta')\Delta s_{p_k}$. We will take advantage of this fact in Lemma 21, where we will combine this with the second condition in the definition of potentially-good set of tasks. This second condition states that P_{-k} is good for instances in which task p_k has values $[t_{p_k} = \delta, s_{p_k} = q\delta/(2n)]$, for $q = 0, 1, \dots, 2n/\delta$. For successive values of q , s_{p_k} changes by $\delta/(2n)$. The above lemma allows us to conclude that successive values of the boundary of t_{p_k} change by at most $\rho\delta/(2n) \leq \delta/2$ and that one of these values is less than δ (see the proof of Lemma 21 for details). This goal lies behind the complicated definition of potentially-good set of tasks.

Lemma 19. *Let T be a standard instance for a regular set of tasks $P = (p_1, \dots, p_k)$. Let $T' = (t', s)$ be an instance with $t'_{p_i} \in (\alpha, \alpha + \beta)$, $i \in [k]$, which agrees with T on the remaining tasks. If all tasks in P are allocated to the 0-player in the allocation of T' , then there exists V such that P is good for T with witness V .*

The next lemma shows that all tasks of a potentially-good but not good set P of tasks must be allocated to the s -players. Furthermore, P is a minimal or critical set in the sense that if we lower the t -value of one of its tasks, the allocation for all of them changes. The first part of the definition of potentially-good sets of tasks is designed to achieve exactly this goal.

Lemma 20. *Suppose that $P = (p_1, \dots, p_k)$ is a potentially-good set of tasks for T with witness V . If P is not good for T and there exists no bad task, then for every instance of the set of V -perturbations of $\hat{T}(P)$, all tasks in P are allocated to the s -players. Furthermore, for each $i \in [k]$, if we change the t -value of only task p_i to β , all tasks in P are allocated to the 0-player.*

E. Almost all sets are good

We now consider a potentially-good but not good set of tasks $P = (p_1, \dots, p_k)$ from clusters (Q_1, \dots, Q_k) . Let p'_k be a sibling of p_k , i.e., another task of cluster Q_k . The main result of this section establishes that *almost all other sets* (P_{-k}, p'_k) are good; we show via a probabilistic argument that (P_{-k}, p'_k) is good with probability at least $1 - 2n^3/\ell$, where p'_k is a random sibling of p_k (Lemma 28).

In order to show this result, we consider the (p_k, p'_k) -slice mechanism, utilizing the 2×2 characterization that we provide in Section II-C, considering all possible cases. Specifically, in Section III-E1 we exclude (almost surely) affine minimizers (Lemma 22), in Section III-E2 we bound from above the probability that the mechanism is a relaxed affine minimizer by $2n^2/\ell$ (Lemma 25) and in Section III-E3 we exclude 1-dimensional and constant mechanisms (Lemma 26). Finally in Section III-E4 we show that if the mechanism is task independent or relaxed task independent for each of k appropriately selected random instances from the witness, then (P_{-k}, p'_k) is good (Lemma 27).

We conclude that for a random sibling p'_k , the mechanism must be either task independent or relaxed task independent for all these k instances with probability at least $1 - k \cdot 2n^2/\ell$; therefore (P_{-k}, p'_k) is good with probability at least $1 - 2n^3/\ell$. (Lemma 28).

1) *Affine minimizers*: The following lemma (Lemma 21) deals with the case that the (p_k, p'_k) -slice mechanism is an affine minimizer or in the linear part of a relaxed affine minimizer. It is essential in showing that *very few* slice-mechanisms are of this type.

Lemma 21. *Assume that there is no bad task and that the approximation ratio is at most ρ . Suppose that $P = (p_1, \dots, p_k)$ is a potentially-good but not good set of tasks for T with witness V . There are no instances T_1, T_2 such that*

- both T_1 and T_2 belong to the set of V -perturbations of $\hat{T}(P)$
- T_1 and T_2 differ only in the value of t_{p_i} , for some fixed $i \in [k-1]$
- for both T_1 and T_2 , the boundary function ψ_{p_k} is truncated linear in s_{p_k} . That is,

$$\psi_{p_k}(t_{-p_k}, s) = \max(0, \lambda(t_{-p_k}, s_{-p_k})s_{p_k} - \gamma(t_{-p_k}, s_{-p_k}))$$

for both T_1 and T_2 .

The above lemma is crucial and its importance is captured by the next lemma that essentially states that linear 2×2 mechanisms do not exist almost surely.

Lemma 22. *Suppose that $P = (p_1, \dots, p_k)$ is a potentially-good but not good set of tasks for T with witness V . Take a random instance T' from the set of V -perturbations of $\hat{T}(P)$. Fix all other values to T' except for task p_k , and*

consider the event \mathcal{E} that the boundary function $\psi_{p_k}(s_{p_k})$ is linear in s_{p_k} . The probability of event \mathcal{E} is 0.

2) *Relaxed affine minimizers*: Lemma 22 essentially excludes the case of 2×2 affine minimizers. We now want to deal with the case of relaxed affine minimizers. This presents an additional difficulty as we may not be in the linear part of the mechanism but at its tail. Although this seems like a special case, it adds another complication to the proof and we need to handle it carefully. We deal with this problem by showing that for most siblings p'_k , the 2×2 mechanism cannot be a relaxed affine minimizer, since then it will be in its linear part and it will remain in the linear part when we change the s_{p_k} value, and hence we can treat it like in Lemma 21. To achieve this, we first show that not many tasks can be allocated to the k player, otherwise the approximation ratio is high (Lemma 23 below). This is useful, because when a task p'_k is allocated to the 0-player, and the (p_k, p'_k) -slice mechanism is a relaxed affine minimizer, we can guarantee that it is in the linear part of the mechanism and it will remain there when we change s_{p_k} (see Lemma 24 below).

Lemma 23. *Suppose that $P = (p_1, \dots, p_k)$ is a potentially-good set of tasks for T with witness V and that the approximation ratio is less than ρ . For every instance of the set of V -perturbations of $\hat{T}(P)$, fewer than $2n^2$ siblings of p_k are allocated to the k -player.*

In the above lemma, we can easily improve the number of tasks from $2n^2$ to $O(n)$, by a simple tightening of its proof, but we again opt for simplicity.

Lemma 24. *Fix an instance T with $t_{p'} > 0$, two tasks p and p' of the same cluster and assume that the (p, p') -slice mechanism for T is a relaxed affine minimizer. If the mechanism allocates task p' to 0-player but not task p , then this instance is not in the bundling tail of the relaxed affine minimizer, and the same holds for every value of s_p . That is, the boundary of $\psi_p(s_p)$ is linear for any value of s_p , for fixed values for t_{-p} and s_{-p} .*

One of the premises of the previous lemma is that $t_{p'} > 0$. We will apply this lemma when $t_{p'} = \beta$.

We now use a probabilistic argument that combines the above lemmas to show that the probability of relaxed affine minimizers is bounded by $2n^2/\ell$.

Lemma 25. *Suppose that $P = (p_1, \dots, p_k)$ is a potentially-good but not good set of tasks for T with witness V . Take a random instance T' from the set of V -perturbations of $\hat{T}(P)$ and a random sibling p'_k of p_k . We consider the (p_k, p'_k) -slice mechanism of tasks p_k and p'_k when we fix the values according to T' for the other tasks. The probability that this 2×2 mechanism is a relaxed affine minimizer or affine minimizer is at most $2n^2/\ell$.*

3) *1-dimensional and constant mechanisms*: Here, we show that 1-dimensional and constant slice mechanisms can be immediately excluded.

Lemma 26. *Suppose that $P = (p_1, \dots, p_k)$ is a potentially-good set of tasks for T with witness V . Fix a sibling p'_k of p_k and some instance T' in the set of V -perturbations of $\hat{T}(P)$. Suppose that the (p_k, p'_k) -slice mechanism for T is a 1-dimensional or a constant mechanism. Then the mechanism has approximation ratio at least ρ .*

4) *Task independent and relaxed task independent mechanisms*: In this subsection we deal with task independent and relaxed task independent slice mechanisms. The next lemma is analogous to Lemma 21 for task independent or relaxed task independent slice mechanisms. There is another layer of difficulty in this case, since we need to take k instances instead of just two instances that we considered in the linear case. The difference is that the goal of Lemma 21 was to essentially exclude linear mechanisms, but the goal of the next lemma is to show that task independent slice mechanisms imply abundance of good sets of tasks.

Lemma 27. *Suppose that $P = (p_1, \dots, p_k)$ is a potentially-good but not good set of tasks for T with witness V . Fix a sibling p'_k of p_k . Suppose that there exists instances T_0, T_1, \dots, T_{k-1} such that for every $i \in \{0, \dots, k-1\}$*

- T_i belongs to the set of V -perturbations of $\hat{T}(P)$
- for every $i \neq 0$, T_i differs from T_0 only in the value of t_{p_i} ; furthermore, the value of T_i for this task is higher than the value of T_0
- the (p_k, p'_k) -slice mechanism for T_i is a task independent or a relaxed task independent mechanism.

Assuming that there is no bad task and that approximation ratio is less than ρ , the set of tasks (P_{-k}, p'_k) is good for T .

Taking advantage of the above lemma, we now show that either a potentially-good set of tasks is also good or there are many other good sets.

Lemma 28. *Suppose that $P = (p_1, \dots, p_k)$ is a potentially-good set of tasks for T with witness V . Then either P is good for T or the probability that (P_{-k}, p'_k) is good for T , for a random sibling p'_k of p_k , is at least $1 - 2n^3/\ell$.*

F. Existence of a good set of tasks of size $n - 1$

The above analysis culminated in Lemma 28. In this subsection, we use this lemma to show the existence of a good set of tasks of size $n - 1$ (Corollary 32), which immediately concludes the proof of the Main Lemma (Lemma 15). We will use a probabilistic argument again, so we first define the probability b_k that a regular set of k tasks is bad.

Definition 29 (Probability b_k). Fix a mechanism and suppose that T is a standard instance for a set \mathcal{C} of k clusters. Let us denote by $b(T, \mathcal{C})$ the probability that a random regular set P of k tasks from \mathcal{C} is not good and let b_k

denote the maximum possible $b(T, \mathcal{C})$ —or more precisely, its supremum—for all choices of T and \mathcal{C} , for the given mechanism.

First, we find an upper bound on the probability that a random regular set of tasks of size k is not potentially-good, based on the probability that a set of size $k - 1$ is not good.

Lemma 30. *Fix any instance T which is standard for a set \mathcal{C} of $k \geq 2$ clusters, and let P be a random regular set of k tasks from \mathcal{C} . The probability that P is not potentially-good for T is at most $(3n/\delta - 1)b_{k-1}$.*

The following lemma estimates a rough upper bound on the probability that a random regular set is not good. It shows that if ℓ has exponential size in n , the probability that there exists a good regular set of $n - 1$ tasks is positive.

Lemma 31 (Probability of bad sets). *Fix a mechanism with approximation ratio less than ρ and suppose that there is no bad task. Then the probabilities of bad events are bounded above by*

$$b_k \leq \left(\frac{3n}{\delta}\right)^{k-2} \frac{3n^3}{\ell}, \quad (3)$$

for every $k \in [n - 1]$.

Using this, we can now reach the target of this section, showing that there exists a good set of size $n - 1$.

Corollary 32. *If $\ell > 3n^3 \left(\frac{3n}{\delta}\right)^{n-3}$, we get $b_{n-1} < 1$, and there exists at least one good set of $n - 1$ tasks.*

G. Lower bound on the approximation ratio

In this subsection we use the Main Lemma (Lemma 15) to prove a lower bound of $1 + \sqrt{n - 1}$.

Theorem 33. *The approximation ratio of truthful mechanisms is at least $1 + \sqrt{n - 1}$.*

Remark. To achieve a ratio very close to $1 + \sqrt{n - 1}$, we must select a small δ' . By Corollary 32, this has implications on the number of tasks required by the proof of the main lemma (Lemma 15). In particular the number of tasks per cluster must be more than $3n^3 \left(\frac{3n}{\delta}\right)^{n-3} = n^{\Theta(n)}/\delta^{n-3}$. Note that by any choice of $\delta < 1$, this number is at least exponential in n .

IV. DISCUSSION

Although we believe that some of the ideas in this work can be used to fully resolve the Nisan-Ronen conjecture, there are also some limitations. For example, the idea of considering instances where tasks can essentially only be assigned to two machines is quite promising, as it allows to exploit the 2×2 characterization. We refer the reader to our previous work [8] for a systematic treatment of such instances, which can be modeled as multi-graphs. Using the terminology of [8], the instances considered in this work

(Definition 8) are multi-stars. The next theorem shows that the lower bound of this work is tight for multi-stars.

Theorem 34. *The approximation ratio for the set of instances of Definition 8 is at most $1 + \sqrt{n-1}$ and it is achieved by a weighted version of the VCG mechanism.*

REFERENCES

- [1] Aaron Archer and Robert Kleinberg. Truthful germs are contagious: a local to global characterization of truthfulness. In *EC*, pages 21–30. ACM, 2008.
- [2] Itai Ashlagi, Shahar Dobzinski, and Ron Lavi. Optimal lower bounds for anonymous scheduling mechanisms. *Mathematics of Operations Research*, 37(2):244–258, 2012.
- [3] Vincenzo Auletta, George Christodoulou, and Paolo Penna. Mechanisms for scheduling with single-bit private values. *Theory Comput. Syst.*, 57(3):523–548, 2015.
- [4] S. Bikhchandani, S. Chatterji, R. Lavi, A. Mu’alem, N. Nisan, and A. Sen. Weak monotonicity characterizes deterministic dominant strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.
- [5] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Prior-independent mechanisms for scheduling. In *STOC*, pages 51–60. ACM, 2013.
- [6] Xujin Chen, Donglei Du, and Luis Fernando Zuluaga. Copula-based randomized mechanisms for truthful scheduling on two unrelated machines. *Theory Comput. Syst.*, 57(3):753–781, 2015.
- [7] George Christodoulou, Elias Koutsoupias, and Annamária Kovács. Mechanism design for fractional scheduling on unrelated machines. *ACM Transactions on Algorithms*, 6(2), 2010.
- [8] George Christodoulou, Elias Koutsoupias, and Annamária Kovács. On the Nisan-Ronen conjecture for submodular valuations. In *STOC*, pages 1086–1096. ACM, 2020. (Full version arXiv:1907.12733).
- [9] George Christodoulou, Elias Koutsoupias, and Annamária Kovács. Truthful allocation in graphs and hypergraphs. In *ICALP*, pages 56:1–56:20, 2021.
- [10] George Christodoulou, Elias Koutsoupias, and Angelina Vidali. A characterization of 2-player mechanisms for scheduling. In *ESA*, volume 5193 of *Lecture Notes in Computer Science*, pages 297–307. Springer, 2008.
- [11] George Christodoulou, Elias Koutsoupias, and Angelina Vidali. A lower bound for scheduling mechanisms. *Algorithmica*, 55(4):729–740, 2009.
- [12] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 8, 1971.
- [13] Constantinos Daskalakis and S. Matthew Weinberg. Bayesian truthful mechanisms for job scheduling from bi-criterion approximation algorithms. In *SODA*, pages 1934–1952. SIAM, 2015.
- [14] Shahar Dobzinski and Ariel Shaulker. Improved Lower Bounds for Truthful Scheduling, 2020.
- [15] Shahar Dobzinski and Mukund Sundararajan. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *EC*, pages 38–47. ACM, 2008.
- [16] Yiannis Giannakopoulos, Alexander Hammerl, and Diogo Poças. A new lower bound for deterministic truthful scheduling. In *SAGT*, volume 12283 of *Lecture Notes in Computer Science*, pages 226–240. Springer, 2020.
- [17] Yiannis Giannakopoulos and Maria Kyropoulou. The VCG mechanism for Bayesian scheduling. *ACM Trans. Economics and Comput.*, 5(4):19:1–19:16, 2017.
- [18] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
- [19] Elias Koutsoupias and Angelina Vidali. A lower bound of $1+\phi$ for truthful scheduling mechanisms. *Algorithmica*, pages 1–13, 2012.
- [20] Ron Lavi and Chaitanya Swamy. Truthful mechanism design for multidimensional scheduling via cycle monotonicity. *Games and Economic Behavior*, 67(1):99–124, 2009.
- [21] Jan Karel Lenstra, David B Shmoys, and Eva Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical programming*, 46(1-3):259–271, 1990.
- [22] Pinyan Lu and Changyuan Yu. An improved randomized truthful mechanism for scheduling unrelated machines. In *STACS*, volume 1, pages 527–538, 2008.
- [23] Pinyan Lu and Changyuan Yu. Randomized truthful mechanisms for scheduling unrelated machines. In *WINE*, volume 5385 of *Lecture Notes in Computer Science*, pages 402–413. Springer, 2008.
- [24] Ahuva Mu’alem and Michael Schapira. Setting lower bounds on truthfulness. *Games and Economic Behavior*, 110:174–193, 2018.
- [25] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.
- [26] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [27] Michael E. Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In *EC*, pages 286–293. ACM, 2005.
- [28] William Vickrey. Counterspeculations, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [29] Angelina Vidali. The geometry of truthfulness. In *WINE*, volume 5929 of *Lecture Notes in Computer Science*, pages 340–350. Springer, 2009.
- [30] Changyuan Yu. Truthful mechanisms for two-range-values variant of unrelated scheduling. *Theoretical Computer Science*, 410(21-23):2196–2206, 2009.