

Mechanisms for Hiring a Matroid Base without Money

Emmanouil Pountourakis¹ and Guido Schäfer²

¹ Northwestern University, USA
manolis@u.northwestern.edu

² CWI and VU University Amsterdam, The Netherlands
g.schaefer@cwi.nl

Abstract. We consider the problem of designing mechanisms for hiring a matroid base without money. In our model, the elements of a given matroid correspond to agents who might misreport their actual costs that are incurred if they are hired. The goal is to hire a matroid base of minimum total cost. There are no monetary transfers involved. We assume that the reports are binding in the sense that an agent's cost is equal to the maximum of his declared and actual costs. Our model encompasses a variety of problems as special cases, such as computing a minimum cost spanning tree or finding minimum cost allocation of jobs to machines.

We derive a polynomial-time randomized mechanism that is truthful in expectation and achieves an approximation ratio of $(m - r)/2 + 1$, where m and r refer to the number of elements and the rank of the matroid, respectively. We also prove that this is best possible by showing that no mechanism that is truthful in expectation can achieve a better approximation ratio in general. If the declared costs of the agents are bounded by the cost of a socially optimal solution, we are able to derive an improved approximation ratio of $3\sqrt{m}$. For example, this condition is satisfied if the costs constitute a metric in the graphical matroid.

Our mechanism iteratively extends a partial solution by adding feasible elements at random. As it turns out, this algorithm achieves the best possible approximation ratio if it is equipped with a distribution that is optimal for the allocation of a single task to multiple machines. This seems surprising given that matroids allow for much richer combinatorial structures than the assignment of a single job.

1 Introduction

The task of designing algorithms that are resilient to manipulations of strategic agents in large, distributed systems (such as the Internet) has become a major challenge in recent years. For example, in online marketplaces (such as eBay or eBid) auction formats are desired that incentivize truth revelation of the bidders' valuations for the items on auction. In online workplaces (like Elance, oDesk or Guru) that match freelance experts to clients mechanisms are sought to prevent unjustified declarations of costs.

The classical approach to incite truth telling in strategic environments is to use *mechanism design* (see, e.g., [15, 14]). Here the basic idea is to issue payments to the agents in order to convince them to behave truthfully. Typically, these payments are used to compensate for the advantage that an agent could obtain by lying. Mechanism design is a powerful approach that gave rise to several enlightening results in the past and still is a very active research area with many intriguing open questions.

However, there are many applications in which monetary transfers (as used in the traditional setting) are infeasible. As a result, researchers have more recently started to look into what is known as *mechanism design without money*. Here the basic question one asks is: Can one incite agents to behave truthfully without the use of monetary transfers? Unfortunately, classical results in voting theory show that the answer to this question is “No!” in general. In particular, the well-known Gibbard-Satterthwaite theorem [9, 18] states that for unrestricted domains and at least three outcomes the only mechanism enforcing truthfulness without monetary transfers is *dictatorial*, i.e., the outcome is determined by a single agent. In particular, this also rules out the possibility of approximating any interesting objective in such a setting.

In light of this strong intractability result, there has recently been a large interest in studying more restrictive settings of mechanism design without money. A partial list of proposals that have been addressed in the literature includes the limitation of the agents’ preferences [17], changing the social choice model using imposition [16] or binding reports [11, 3, 2].

Our model In this paper, we study the problem of selecting a minimum cost matroid base in a strategic environment. Here the elements of the matroid correspond to agents who might misreport their actual costs. The intuition behind our model is that a certain task can be accomplished only through the collaboration of certain groups of agents. These groups correspond to the bases of the given matroid. Each agent i declares a cost c_i for performing the task, which is not necessarily equal to his actual cost. Based on the declared costs, the mechanism designer wants to “hire” a matroid base at the cheapest possible cost. There are no monetary transfers between the mechanism designer and the agents.

As an example, suppose that the mechanism designer wants to hire a spanning tree in a given network in order to establish connectivity between all nodes at the lowest possible cost. Here the agents are the edges and each edge declares a cost that it incurs for establishing connectivity between its endpoints. This problem falls into our matroid model simply by using the graphic matroid whose bases correspond to the spanning trees of the given graph.

Another example is the problem of scheduling n jobs on m unrelated machines (possibly with restrictions). Every machine i declares for each job $j \in [n]$ it can execute a processing time p_{ij} . The goal is to determine an assignment of jobs to machines such that the total processing time is minimized. It is not hard to see that this problem is a special case of finding a minimum cost basis in a partitioning matroid and is therefore captured by our matroid model.

Binding reports The latter problem was studied by Koutsoupias [11] under the assumption that the reports are *binding*. This notion was first considered by Christodoulou et al. [3] and Angel et al. [2]. Basically, this means that one can detect whether an agent overstates his actual cost. The motivation for this assumption is that in many situations costs are “observable” and thus declaring a cost that is larger than the actual one can be punished. On the other hand, if an agent understates his cost then his actual cost remains unaffected through this false declaration. For example, in the scheduling problem mentioned above binding reports means that the mechanism can enforce that the machine is busy for at least the declared processing time.

Koutsoupias [11] settles the problem of assigning one job to m machines completely. He designs a randomized algorithm that is truthful in expectation and achieves an approximation ratio of $(m + 1)/2$ (which he shows is best possible). He also extends these results to the case of scheduling n jobs on m machines. The crucial insight in [11] that enables him to derive these results is a characterization of the distributions for the assignment of a single job that guarantee truthfulness in expectation. Given this characterization, he then determines a distribution that achieves the best possible approximation ratio.

Our contributions Here we continue this line of research. We consider the problem of designing mechanisms without money for the more general model of hiring a matroid base under binding reports. Our main contributions are as follows:

1. We give a randomized algorithm that is truthful in expectation and achieves an approximation ratio of $(m - r)/2 + 1$, where m and r refer to the number of elements and the rank of the underlying matroid, respectively.
2. We prove that this approximation ratio is best possible. More specifically, we show that no (randomized) mechanism that is truthful in expectation can achieve a better approximation ratio.
3. We then show that an improved approximation ratio of $3\sqrt{m}$ can be achieved if the declared costs of the agents are bounded by the cost of a socially optimal solution. For example, this condition is satisfied if the costs constitute a metric in the graphic matroid.

Our techniques Our results are based on a natural extension of the greedy algorithm for the computation of a minimum cost basis of a matroid. The algorithm iteratively extends a partial solution by adding elements that maintain feasibility. However, because of truthfulness we cannot enforce that a minimum cost element is chosen in each iteration (as in the standard greedy algorithm). Instead, we have to ensure that in each iteration each feasible addition of an element is chosen with some positive probability such that the resulting probability of picking an element meets certain monotonicity properties.

Although we have some freedom to choose these distributions, their choice impacts the resulting approximation ratio of the mechanism. Intuitively, we would like to tailor these distributions in such a way that the minimum cost element is

chosen with some good probability, while the required monotonicity properties are still satisfied. Here the insights obtained by Koutsoupias [11] for assigning a single job to m machines turn out to be very useful.

Our findings show that an appropriate composition of the distribution that is proven to be optimal for the single task assignment in [11] also delivers the best possible results in the more general setting of hiring a matroid base. We find this somewhat surprising because matroids allow for combinatorially much richer structures than the assignment of a single job. In fact, the problem of optimally assigning a single job to m machines is equivalent to computing a minimum cost basis of a 1-uniform matroid (which is one of the most trivial matroids). For this special case our mechanism coincides with the one of Koutsoupias.

In order to bound the approximation ratio of our mechanism we crucially exploit properties of the matroid. However, there are many approximation algorithms that follow a similar design paradigm of iteratively extending a partial solution in a greedy manner (e.g., the greedy algorithm for the set cover problem). We conjecture that our findings might be extended to a broader context of greedy-like approximation algorithms which gives rise to some intriguing questions for follow-up research.

Additional related work The design of mechanism that do not use monetary transfers has recently received considerable attention in the literature on economics and computation. Procaccia and Tennenholtz [17] initiated the study of approximate mechanism design without payments for combinatorial problems by studying facility location problems. Their studies triggered several follow-up articles on this topic (see, e.g., [1, 13, 12, 6, 7]). Dughmi and Gosh [4] derived approximate mechanisms without money for several variants of the assignment problems. Guo and Conitzer [10] studied the problem of selling items without payments for the case of two agents.

The idea of binding reports is also related to *mechanisms with verification*, whose study was first proposed by Nisan and Ronen [15]. However, the notion of verification is much stronger than the notion of binding reports that we consider here. In particular, mechanisms with verification may defer the issuing of payments to the agents until they learned the actual outcome. As a result, these mechanisms can punish misreports a posteriori by imposing very high penalties for lying.

Mechanism with binding reports are related to the notion of *imposition* proposed by Nissim et al. [16]. In the context of the facility location problem, agents might be forced to connect to the facility that is closest to their declared position instead of the one that is closest to their actual position. This approach was further pursued by Fotakis and Tzamos [8].

2 Preliminaries

In this section, we give a formal definition of the model that we consider in this paper and introduce some basic concepts.

2.1 Matroids

We first formally introduce the notion of a *matroid*:

Definition 1. A matroid $\mathcal{M} = (E, F)$ is defined by a finite set E of elements and a set $F \subseteq 2^E$ of subsets of E satisfying

1. $\emptyset \in F$ (non-emptiness),
2. if $S \in F$ and $S' \subseteq S$ then $S' \in F$ (downward closure),
3. if $S, T \in F$ and $|S| > |T|$ then there exists some $i \in S \setminus T$ such that $T + i \in F$ (exchange property).³

The sets in F are called independent sets. An inclusion-wise maximal independent set $B \in F$ is a basis of \mathcal{M} .⁴ The common size of all bases of \mathcal{M} is called the rank of \mathcal{M} and will be denoted by $r(\mathcal{M})$.⁵

Throughout this paper, we assume that the matroid $\mathcal{M} = (E, F)$ is implicitly represented by an *independent set oracle*: given a set $S \subseteq E$, the oracle specifies whether S is an independent set or not. Unless specified otherwise, we identify the elements in E with the first m natural numbers, i.e., $E = [m]$. We assume that every element $i \in E$ constitutes an independent set, i.e., $\{i\} \in F$.⁶ Note that this assumption is without loss of generality because we can remove all elements from E that do not occur in any independent set.

Example 1. A typical example of a matroid is the *graphic matroid*. Given a graph $G = (V, E)$, we let the edges E of G be the elements of the matroid and each subset $S \subseteq E$ of edges that does not contain a cycle in G constitutes an independent set in F . It is easy to verify that Properties 1–3 of Definition 1 are satisfied. The bases of $\mathcal{M} = (E, F)$ correspond to the spanning trees of G . The rank of \mathcal{M} is $r(\mathcal{M}) = n - 1$, where n is the number of vertices in G .

2.2 Hiring a matroid base

Let $\mathcal{M} = (E, F)$ be a matroid. In our model, we associate an agent with each element $i \in E$ of the matroid. Each agent $i \in E$ has a non-negative cost $\bar{c}_i \in \mathbb{R}_+$. Intuitively, by choosing agent $i \in E$ a cost of \bar{c}_i is incurred. The cost \bar{c}_i is “private” in the sense that it is unknown to us. Our goal is to select (or hire) a base of the matroid of minimum total cost. The intuition behind our model is that the bases of the underlying matroid represent groups of agents that together can perform a certain task.

³ For ease of notation, for a set $T \subseteq E$ and an element $i \in E$ we also use $T + i$ and $T - i$ as a short for $T \cup \{i\}$ and $T \setminus \{i\}$, respectively.

⁴ Subsequently, by “maximal” we mean “inclusion-wise maximal”, i.e., B is maximal if for every $i \in E \setminus B$, $B + i$ is not an independent set.

⁵ Using the properties of Definition 1, it is not hard to show that all bases of a matroid \mathcal{M} have the same size.

⁶ We slightly abuse notation here and write $i \in F$ instead of $\{i\} \in F$ for notational convenience.

Example 2. In order to establish connectivity among all nodes in a given graph $G = (V, E)$ one may want to determine a minimum cost spanning tree of G . Here each edge $i \in E$ corresponds to an agent and selecting an edge incurs a cost of \bar{c}_i . Our goal then is to select a minimum cost basis of the graphic matroid.

2.3 Binding reports

We assume that agents might misreport their costs, i.e., each agent $i \in E$ declares a cost c_i , which is possibly different from his actual cost \bar{c}_i . Based on the matroid \mathcal{M} and the declared costs $\mathbf{c} = (c_1, \dots, c_m)$, the mechanism selects a basis of the underlying matroid. We consider mechanisms *without money*, i.e., the mechanism does not receive/issue any payments from/to the agents.

In order to achieve truthfulness it will turn out to be crucial to allow for random selections of agents, i.e., we consider *randomized* mechanisms. Subsequently, we use $p_i(\mathbf{c})$ to refer to the probability that our (random) mechanism picks element $i \in E$, given the reported costs \mathbf{c} .

We assume that the reports are *binding* as proposed by Koutsoupias [11]. More precisely, if agent i 's reported cost is c_i then his actual cost is $\max\{\bar{c}_i, c_i\}$. That is, if agent i overstates his actual cost by reporting $c_i > \bar{c}_i$ and agent i is selected then his actual cost becomes c_i . On the other hand, if agent i understates his actual cost \bar{c}_i and is selected then his actual cost remains \bar{c}_i because this is the cost incurred by i . Formally, we assume that each agent $i \in E$ strives to minimize his expected cost

$$C_i(\mathbf{c}) = \max\{\bar{c}_i, c_i\}p_i(\mathbf{c}).$$

Subsequently, we use $\bar{\mathbf{c}} = (\bar{c}_1, \dots, \bar{c}_m) \in \mathbb{R}_+^m$ to refer to the vector of *actual* costs and $\mathbf{c} = (c_1, \dots, c_m) \in \mathbb{R}_+^m$ to refer the vector of *declared* costs.

2.4 Truthful mechanisms

We are interested in designing mechanisms that are *truthful in expectation*, which we define next. To this aim, we first need to introduce some standard notation. Let $\mathbf{c} = (c_1, \dots, c_m) \in \mathbb{R}_+^m$ be a cost vector. Then we denote by \mathbf{c}_{-i} , $i \in [m]$, the $(m - 1)$ -dimensional vector with the i th coordinate removed, i.e.,

$$\mathbf{c}_{-i} = (c_1, c_2, \dots, c_{i-1}, c_{i+1}, \dots, c_m).$$

For a subset $T \subseteq [m]$, we will also use \mathbf{c}_T to refer to the restriction of \mathbf{c} to index set T , i.e., $\mathbf{c}_T = (c_{i_1}, c_{i_2}, \dots, c_{i_{|T|}})$ with $T = \{i_1, \dots, i_{|T|}\}$.

Definition 2 (Truthful mechanism). *A mechanism M is truthful in expectation if for every agent i and every vector \mathbf{c}_{-i} , the expected cost of i is minimized by declaring the actual cost truthfully, i.e., for every $i \in E$,*

$$C_i(\bar{c}_i, \mathbf{c}_{-i}) = \bar{c}_i p_i(\bar{c}_i, \mathbf{c}_{-i}) \leq \max\{\bar{c}_i, c_i\} p_i(\mathbf{c}) = C_i(\mathbf{c}).$$

There are stronger notions of truthfulness (e.g., truthfulness or universal truthfulness). However, it is easy to see that with these stronger notions of truthfulness no positive results are possible; see also [11].

2.5 Approximate social cost

The *social cost* function that we consider throughout this paper is the sum of the individual costs, i.e., $SC(\mathbf{c}) = \sum_{i \in E} C_i(\mathbf{c})$. We use $OPT(\mathbf{c})$ to refer to the cost of a socially optimal solution, i.e., the minimum cost of a base of $\mathcal{M} = (E, F)$ with respect to \mathbf{c} . Ideally, we would like to derive a truthful mechanism that computes a socially optimal outcome. However, this is impossible and we therefore relax the optimality condition and resort to approximate solutions.

Definition 3. *A mechanism M is α -approximate with $\alpha \geq 1$ if for every vector \mathbf{c} of declared costs, the expected social cost satisfies*

$$SC(\mathbf{c}) = \sum_{i \in E} C_i(\mathbf{c}) = \sum_{i \in E} \max\{\bar{c}_i, c_i\} p_i(\mathbf{c}) \leq \alpha OPT(\mathbf{c}).$$

2.6 Koutsoupias' characterization

Koutsoupias [11] considers the problem of scheduling one job on m available machines. The actual cost incurred by machine i to schedule job j is \bar{p}_i and each machine wants to minimize his cost. The overall objective is to determine an assignment of minimum total cost.⁷ Note that this corresponds to computing a minimum cost basis in the matroid that consists only of singletons.

Koutsoupias characterizes the set of truthful mechanisms for this problem.

Proposition 1 ([11]). *Let $p_i(\mathbf{c})$ be the probability that element i is chosen by mechanism M given the vector of declared costs \mathbf{c} . Then M is truthful in expectation if and only if for every $i \in E$:*

1. $p_i(c_i, \mathbf{c}_{-i})$ is non-increasing in c_i ,
2. $c_i p_i(c_i, \mathbf{c}_{-i})$ is non-decreasing in c_i .

Based on the above characterization result, Koutsoupias then derives a distribution that satisfies the above properties and whose expected social cost is at most $(m+1)/2$ times the optimal one. He also proves that this is best possible in the sense that no other truthful in expectation mechanism (without payments) can achieve a better approximation ratio.

3 Greedy Mechanism and Truthfulness Conditions

In this section, we provide a general framework for constructing truthful mechanisms. Our framework is based on the greedy approach which iteratively extends a partial solution (i.e., independent set) by adding a least cost element. We parameterize our mechanism with a collection of distributions: for every $T \subseteq E$ we are given a distribution $d^T = \{d_i^T(\mathbf{c}_T) \mid i \in T\}$ over the elements i in T .⁸

⁷ We note that in [11] also the objective of minimizing the makespan is considered.

⁸ The assumption that all these distributions are given is a conceptual one. Subsequently, it will become clear that we can generate the relevant distributions considered by the algorithm efficiently.

Definition 4 (Greedy Mechanism). Given a matroid $\mathcal{M} = (E, F)$ with a cost vector \mathbf{c} and a collection of distributions $(d^T)_{T \subseteq E}$, the greedy algorithm is as follows:

1. Let $S \leftarrow \emptyset$
2. While S is not a base
 - (a) Let $T = \{i \in E \setminus S \mid S + i \in F\}$
 - (b) Draw $i \in T$ with probability $d_i^T(\mathbf{c}_T)$
 - (c) Set $S \leftarrow S + i$
3. Output S

Note that the set T in Step 2(a) contains all elements that can be added to the independent set S without rendering it infeasible. The main difference of our algorithm to the standard greedy algorithm for matroids is that we do not require that the element $i \in T$ added to S in Step 2(c) is of minimum cost. Indeed, such a mechanism would not be truthful because it failed to satisfy Condition (b) of Proposition 1. Instead, here we choose an element i from T with probability $d_i^T(\mathbf{c}_T)$. In particular, our algorithm coincides with the standard greedy algorithm if $d_i^T(\mathbf{c}_T) > 0$ only for the minimum cost elements in T .

We next establish some sufficient conditions for the distributions used by our greedy algorithm that ensure truthfulness.

Theorem 1. *The greedy mechanism is truthful in expectation if for every $T \subseteq E$ and every $i \in T$ it holds:*

1. $d_j^T(c_i, \mathbf{c}_{T-i})$ is non-decreasing in c_i for every $j \in T - i$,
2. $d_i^T(c_i, \mathbf{c}_{T-i})c_i$ is non-decreasing in c_i .

Proof. Fix a collection of distributions $(d^T)_{T \subseteq E}$ that satisfies Properties (1) and (2). Let $p_i(\mathbf{c})$ be the probability of picking element $i \in E$ after the execution of the mechanism. We need to show that Properties (1) and (2) of Proposition 1 are satisfied, i.e.,

1. $p_i(c_i, \mathbf{c}_{-i})$ is non-increasing in c_i ,
2. $p_i(c_i, \mathbf{c}_{-i})c_i$ is non-decreasing in c_i .

We prove these by induction on the number m of elements in E . If $m = 1$ then there is only one element to be picked and the properties clearly hold.

Suppose that the claim holds true for all element sets of size less than m . We show that it continues to hold for sets of size m . We use $\mathcal{M}^{(j)}$ to refer to the matroid that we obtain from \mathcal{M} by contracting element j , i.e., the sub-matroid that contains only the sets that include j .

Let $p_i^{(j)}(\mathbf{c}_{-j})$ be the probability of picking element i in the matroid $\mathcal{M}^{(j)}$. Note that $p_i^{(j)}(\mathbf{c}_{-j})$ is precisely the probability of picking element i conditional on the event that player j has been picked in the first round.

First property: Using Bayes rule, we obtain

$$p_i(\mathbf{c}) = d_i^E(\mathbf{c}) + \sum_{j \in E-i} d_j^E(\mathbf{c})p_i^{(j)}(\mathbf{c}_{-j}) = 1 - \sum_{j \in E-i} d_j^E(\mathbf{c}) + \sum_{j \in E-i} d_j^E(\mathbf{c})p_i^{(j)}(\mathbf{c}_{-j})$$

$$= 1 + \sum_{j \in E-i} d_j^E(\mathbf{c}) [p_i^{(j)}(\mathbf{c}_{-j}) - 1].$$

By assumption, $d_j^E(\mathbf{c})$ is non-decreasing in c_i for every $j \neq i$. Also, by our induction hypothesis, $p_i^{(j)}(\mathbf{c}_{-j})$ is non-increasing in c_i . Thus, the product $d_j^E(\mathbf{c}) [p_i^{(j)}(\mathbf{c}_{-j}) - 1]$ is non-increasing in c_i . We conclude that $p_i(\mathbf{c})$ is non-increasing in c_i .

Second property: Using Bayes rule, we obtain

$$p_i(\mathbf{c})c_i = d_i^E(\mathbf{c})c_i + \sum_{j \in E-i} d_j^E(\mathbf{c}) [p_i^{(j)}(\mathbf{c}_{-j})c_i]$$

By assumption and our induction hypothesis, $d_j^E(\mathbf{c})$ and $p_i^{(j)}(\mathbf{c}_{-j})c_i$ are non-decreasing in c_i for every $j \neq i$. Thus, their product is non-decreasing in c_i . By assumption, also $d_i^E(\mathbf{c})c_i$ is non-decreasing in c_i . We conclude that $p_i(\mathbf{c})c_i$ is non-decreasing in c_i .

4 Optimal Distributions and Approximation Ratio

In this section we identify a distribution that satisfies Properties (1) and (2) of Theorem 1 and yields a truthful in expectation mechanism with approximation ratio $(m - r)/2 + 1$, where $r = r(\mathcal{M})$ is the rank of the underlying matroid \mathcal{M} .

4.1 Optimal Distributions

A natural choice for a collection $(d^T)_{T \subseteq E}$ of distributions to be used by the greedy algorithm is to choose each element i from a given set T with probability that is inversely proportional to its cost c_i . This distribution is also independently considered in [5]. More precisely, for every $T \subseteq E$ and every $i \in T$, we define

$$d_i^T(\mathbf{c}_T) = \frac{c_i^{-1}}{\sum_{k \in T} c_k^{-1}}. \quad (1)$$

The distribution d^T is also called the *proportional distribution*. It is not hard to show that these distributions satisfy Properties (1) and (2) of Theorem 1. However, the problem is that the greedy mechanism equipped with these distributions results in an approximation ratio which is arbitrarily close to m .

The following distribution was introduced by Koutsoupias [11] for scheduling a single job on m machines. (A similar probability distribution is considered and analyzed in the Facility Location setting in [7].)

Definition 5 (Optimal Distribution). Let $T \subseteq E$ be a subset of elements and assume without loss of generality that $T = \{1, \dots, |T|\}$ such that $c_1 \leq c_2 \leq \dots \leq c_{|T|}$. Define probabilities⁹

$$d_1^T(\mathbf{c}_T) = \frac{1}{c_1} \int_0^{c_1} \prod_{k \neq 1} \left(1 - \frac{x}{c_k}\right) dx$$

⁹ This distribution corresponds to the following experiment: We select uniformly at random and independently a number $x_i \in [0, c_i]$ for each element i . The distribution d_i^T of Definition 5 corresponds to the distribution of the minimum of these x_i 's.

$$d_j^T(\mathbf{c}_T) = \frac{1}{c_1 c_j} \int_0^{c_1} \int_0^y \prod_{k \neq 1, j} \left(1 - \frac{x}{c_k}\right) dx dy \quad \text{for } j \neq 2.$$

This generalized distribution yields a truthful in expectation greedy mechanism that achieves the best possible approximation ratio.

Theorem 2. *The greedy mechanism equipped with the distributions of Definition 5 is truthful.*

All omitted proofs from this and future sections can be found in the full version of the paper.

4.2 Approximation Ratio

Koutsoupias [11] used the distribution d^T given in Definition 5 to handle the case of allocating a single job to m machines. He showed that the resulting mechanism achieves an approximation ratio of $(m + 1)/2$. Here we prove that our greedy mechanism, equipped with the distributions in Definition 5, has an approximation ratio of $(m - r)/2 + 1$, where r is the rank of the matroid.

Theorem 3. *The greedy mechanism with distributions d^T as defined in Definition 5 has approximation ratio $(m - r)/2 + 1$.*

5 Lower bound

In this section, we provide a general lower bound on the approximation ratio of truthful in expectation mechanisms for hiring a matroid base that matches the upper bound of our greedy algorithm established in the previous section.

We show that for any given parameters m and r , we can always construct a matroid with m elements and rank r such that no mechanism that is truthful in expectation can achieve an approximation ratio better than $(m - r)/2 + 1$ ¹⁰.

Using the previous lemma we show that for every choice of m and r our upper bound is tight in the sense that there exists a matroid instance where any truthful mechanism has approximation ratio $(m - r)/2 + 1$.

Theorem 4. *Given m and r , there exists a matroid $\mathcal{M} = (E, F)$ with $|E| = m$ and $r(\mathcal{M}) = r$ for which no mechanism that is truthful in expectation can achieve an approximation ratio better than $(m - r)/2 + 1$.*

Finally, we show a weaker result regarding only graphical matroids. Specifically, we show that there is a family of graphs where the worst case bound of $(m - r)/2 + 1$ occurs.

Theorem 5. *There is no mechanism that is truthful in expectation that achieves an approximation ratio better than $(m - r)/2 + 1$ for graphical matroids.*

¹⁰ Note this result does not necessarily imply that every truthful mechanism will perform poorly given any matroid set system with these parameters.

6 Improved Approximation Ratio for Metrics

In this section, we show that we can derive an improved approximation ratio of $O(\sqrt{m})$ for our greedy algorithm if each agent's declared cost is at most the cost of a socially optimal solution, i.e., for every agent $i \in E$, $c_i \leq \text{OPT}(\mathbf{c})$. Said differently, this condition requires that the cost of an arbitrary base of the matroid is at least as large as $\max_{i \in E} c_i$. We call a vector $\mathbf{c} = (c_1, \dots, c_m)$ of declared costs *opt-bounded* if it satisfies this condition.

Note that in the case of a graphical matroid this property is trivially satisfied if the declared cost vectors \mathbf{c} are restricted to constitute a metric. If the declared cost vectors \mathbf{c} are restricted to constitute a metric then this condition is trivially satisfied. It is interesting to note that we obtain this result for the greedy algorithm using the proportional distributions.

Theorem 6. *If the declared cost vector is opt-bounded then the greedy mechanism using the proportional distributions as defined in (1) is truthful in expectation and achieves an approximation ratio of $3\sqrt{m}$.*

7 Future work

There are a lot of open problems that arise from our work. We designed an algorithm that achieves an approximation ratio based on the size of the matroid and its rank. In Section 5 we proved a lower bound that was dependent on the substitutability of elements within the matroid's bases. It could be possible to provide a more refined upper bound using this parameter.

Also there are many questions still open in the case of graphical matroids when the costs constitute a metric. We analyzed only the proportional method which generally performs worse than the distribution in Definition 5. We also have no matching lower bounds. Additionally, our iterative algorithm and generally our framework didn't depend on the matroid property of the set system to satisfy truthfulness. Thus, it will be interesting to analyze its performance in more general settings especially where the classic greedy has good approximation guarantees. Finally, we only considered social costs and not other social objectives like a minmax solution concept.

References

1. Noga Alon, Michal Feldman, Ariel D. Procaccia, and Moshe Tennenholtz. Strategyproof approximation of the minimax on networks. *Math. Oper. Res.*, 35(3):513–526, 2010.
2. Eric Angel, Evripidis Bampis, Fanny Pascual, and Alex-Ariel Tchetgnia. On truthfulness and approximation for scheduling selfish tasks. *J. Scheduling*, 12(5):437–445, 2009.
3. George Christodoulou, Laurent Gourvès, and Fanny Pascual. Scheduling selfish tasks: About the performance of truthful algorithms. In Guohui Lin, editor, *COCOON*, volume 4598 of *Lecture Notes in Computer Science*, pages 187–197. Springer, 2007.

4. Shaddin Dughmi and Arpita Ghosh. Truthful assignment without money. In David C. Parkes, Chrysanthos Dellarocas, and Moshe Tennenholtz, editors, *ACM Conference on Electronic Commerce*, pages 325–334. ACM, 2010.
5. Bruno Escoffier, Laurent Gourvès, Nguyen Kim Thang, Fanny Pascual, and Olivier Spanjaard. Strategy-proof mechanisms for facility location games with many facilities. In Ronen I. Brafman, Fred S. Roberts, and Alexis Tsoukiàs, editors, *ADT*, volume 6992 of *Lecture Notes in Computer Science*, pages 67–81. Springer, 2011.
6. Dimitris Fotakis and Christos Tzamos. On the power of deterministic mechanisms for facility location games. In Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, and David Peleg, editors, *ICALP (1)*, volume 7965 of *Lecture Notes in Computer Science*, pages 449–460. Springer, 2013.
7. Dimitris Fotakis and Christos Tzamos. Strategyproof facility location for concave cost functions. In Michael Kearns, R. Preston McAfee, and Éva Tardos, editors, *ACM Conference on Electronic Commerce*, pages 435–452. ACM, 2013.
8. Dimitris Fotakis and Christos Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. *Theor. Comput. Sci.*, 472:90–103, 2013.
9. Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, July 1973.
10. Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, editors, *AAMAS*, pages 881–888. IFAAMAS, 2010.
11. Elias Koutsoupias. Scheduling without payments. In Giuseppe Persiano, editor, *SAGT*, volume 6982 of *Lecture Notes in Computer Science*, pages 143–153. Springer, 2011.
12. Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan Allen Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM conference on Electronic commerce, EC '10*, pages 315–324, New York, NY, USA, 2010. ACM.
13. Pinyan Lu, Yajun Wang, and Yuan Zhou. Tighter bounds for facility games. In *Proceedings of the 5th International Workshop on Internet and Network Economics, WINE '09*, pages 137–148, 2009.
14. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
15. Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1-2):166–196, 2001.
16. Kobbi Nissim, Rann Smorodinsky, and Moshe Tennenholtz. Approximately optimal mechanism design via differential privacy. In Shafi Goldwasser, editor, *ITCS*, pages 203–213. ACM, 2012.
17. Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In John Chuang, Lance Fortnow, and Pearl Pu, editors, *ACM Conference on Electronic Commerce*, pages 177–186. ACM, 2009.
18. Mark Allen Satterthwaite. Strategy-proofness and arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, April 1975.