



Budget Feasible Mechanisms on Matroids

Stefano Leonardi¹ · Gianpiero Monaco² · Piotr Sankowski³ · Qiang Zhang¹

Received: 23 July 2019 / Accepted: 26 October 2020 / Published online: 13 November 2020
© The Author(s) 2020

Abstract

Motivated by many practical applications, in this paper we study *budget feasible mechanisms* with the goal of procuring an independent set of a matroid. More specifically, we are given a matroid $\mathcal{M} = (E, \mathcal{I})$. Each element of the ground set E is controlled by a selfish agent and the cost of the element is private information of the agent itself. A budget limited buyer has additive valuations over the elements of E . The goal is to design an incentive compatible budget feasible mechanism which procures an independent set of the matroid of largest possible value. We also consider the more general case of the pair $\mathcal{M} = (E, \mathcal{I})$ satisfying only the hereditary property. This includes matroids as well as matroid intersection. We show that, given a polynomial time deterministic algorithm that returns an α -approximation to the problem of finding a maximum-value independent set in \mathcal{M} , there exists an individually rational, truthful and budget feasible mechanism which is $(3\alpha + 1)$ -approximated and runs in polynomial time, thus yielding also a 4-approximation for the special case of matroids.

Keywords Mechanism design · Budget feasible mechanisms · Matroid intersection · Hereditary property

A preliminary version of this work appeared in the Proceedings of the 19th International Conference on Integer Programming and Combinatorial Optimization (IPCO) [20]. This work was partially supported by FET IP project MULTIPLEX 317532, by the polish funds for years 2011–2015 for co-financed international projects, NCN Grant UMO-2014/13/B/ST6/00770, ERC Project PAAL-POC 680912. It was also partly supported by the Google Focused Award on Algorithms for Large-scale Data Analysis, ERC Advanced Grant 788893 AMDROMA “Algorithmic and Mechanism Design Research in Online Markets”, and MIUR PRIN project ALGADIMAR “Algorithms, Games, and Digital Markets”

✉ Stefano Leonardi
leonardi@diag.uniroma1.it

Extended author information available on the last page of the article

1 Introduction

Procurement auctions (a.k.a. reverse auctions) are executed by governments or private companies to purchase commodities and services from providers. The budget that can be spent in a procurement auction is often limited thus imposing a limit on the total payments that can be handled to the providers. Motivated by the setting described above, in this work we study the problem of a budget limited buyer with additive valuations on a set of indivisible items, each item controlled from an independent strategic agent. More specifically, we assume rational agents with quasi-linear utilities, i.e., they aim to maximize the difference between the payment they receive and the true cost of the item. We consider the case of a buyer that is constrained to purchase a subset of objects that forms an independent set of an underlying matroid structure. Matroids are indeed linked to many interesting economic applications, e.g., auctions [5, 12, 18], spectrum market [24], scheduling [11] and house market [19].

One of the main challenge in procurement auctions is to incentivize sellers to declare their true costs (by assuming that those costs are *private information*). The celebrated Vickrey-Clark-Groves (VCG) mechanism [10, 15, 25] provides a solution to this problem: it returns a solution that maximizes the valuation of the buyer while charging the sellers with payments equal to their externalities to the procurement. The VCG mechanism is a *truthful* mechanism, i.e., no seller can improve her utility by manipulating her cost regardless of the costs declared by other agents. However, VCG has also a major shortcoming: the payments to the sellers cannot be controlled and could be higher than the available budget. To overcome this problem, two different approaches have been proposed and investigated. The first one is the design of *frugal mechanisms* [16] which minimize the payments handled from the buyer to the strategic sellers. The other approach, also investigated in this work, is the one of developing *budget feasible mechanisms* [23] aiming to maximize the value for the procurement under a given budget that limits the total payments to the sellers. Unfortunately, truthful budget feasible mechanisms cannot maximize the valuation of the buyer and therefore the attention was turned to the development of truthful budget feasible mechanisms that obtains an approximation of the optimal value of the procurement in a non-strategic setting. This approach was initiated by Singer [23] that presented a budget feasible mechanisms for a buyer with additive and non-decreasing submodular valuations that approximate the optimal value of the procurement computed on the true costs of the sellers.

1.1 Our Results

The goal of our work is to design budget feasible mechanisms for procuring items that form an independent set in a given matroid structure. To the best of our knowledge, this is the first time that matroid constraints are considered in budget feasible mechanism design. Nevertheless, similar results to the ones presented in our paper have been contemporaneously published in an independent work [1].

Our results are positive. In Sect. 3 we consider the general problem of procuring a set of maximum-value in a set system holding the only hereditary property. In particular, given a deterministic polynomial time α -approximation algorithm for the maximum-value set problem, we present a deterministic, polynomial time, individually rational, truthful and budget feasible mechanism which is a $(3\alpha + 1)$ -approximation to the optimal solution. Then, in Sect. 3.1, the proposed mechanism is used to obtain a deterministic, polynomial time, individually rational, truthful and budget feasible mechanism which is a 4-approximation to the maximum-value independent set of a matroid. These results hold for the case of additive valuations over the objects.

Given the results presented in [9], where the authors prove that any deterministic mechanism cannot achieve an approximation ratio better than $1 + \sqrt{2}$ for additive valuations (it is worth noticing that such lower bound does not rely on any computational or complexity assumption), we conclude that the performance of our mechanism is not far from the best achievable ones.

Our methods hold several potential applications also because of their simplicity. Budget feasible mechanism design was actually motivated [9, 23] by the application to crowdsourcing [13], and the uniform matroid constraint that we introduce allows to impose an upper limit on the total number of workers that can be hired. Our mechanisms also find application to job scheduling. Consider the problem of purchasing processing time in scheduling jobs associated with a deadline and a profit. As jobs may conflict with each other, only one job can be scheduled at the same time. Our work allows to model this setting with a matroid constraint on the set of jobs that can be scheduled. Finally, we mention an application to the spectrum market. Tse and Hanly [24] showed that the set of achievable rates in a Gaussian multiple-access channel, known as the Cover-Wyner capacity region, forms a polymatroid. It is known a pseudopolynomial reduction from integral polymatroid to matroids [22]. Therefore, our mechanism can also find application to the purchase of transmission rates in the spectrum market.

1.2 Related Work

The study of the budget feasible mechanisms was initiated in [23] that gives a 6-approximation polynomial deterministic mechanism for additive valuations. This result was improved in [9] that gives a deterministic $2 + \sqrt{2}$ -approximated mechanism and a deterministic lower bound of $1 + \sqrt{2}$. More recently, [14] proposes a new and simple randomized mechanism that yields a tight approximation ratio of 2 for additive valuations. The authors of [14] also present a simple deterministic mechanism with a tight approximation guarantee of 3 against the fractional optimum.

This study of budget feasible mechanisms was later extended to consider different buyer's valuation functions $w(\cdot)$ defined over subset of the items set: i) Additive (a.k.a. linear) $w(S) = \sum_{j \in S} w(j)$; ii) Submodular $w(S \cup T) + w(S \cap T) \leq w(S) + w(T)$; iii) Fractionally subadditive (XOS) $w(S) \leq \sum_i \alpha_i w(T_i)$, $0 \leq \alpha_i \leq 1$ if $\sum_{i: j \in T_i} \alpha_i \geq 1, \forall j \in S$. iv.) Subadditive $w(S \cup T) \leq w(S) + w(T)$.

Already in the first paper of Singer [23], it was given a deterministic inapproximability lower bound of 2 and a randomized 112-approximation for general monotone submodular functions. The result for monotone submodular functions was later improved to 5 in [17].

Bei et al. [6] provide a 436-approximation mechanism for XOS valuations and extend the study of budget feasible mechanisms to the Bayesian settings. A recent work [2] improves the approximation factor for XOS functions to 244. Chan and Chen [8] study budget feasible mechanisms when each seller holds multiple copies of the object. They give a logarithmic approximation mechanisms for concave additive and sub-additive valuations.

Budget feasible mechanisms are also attractive due to their various applications. In crowdsourcing, one of the main goal is to assign skilled workers with private costs to tasks. By injecting some characteristics in crowdsourcing, budget feasible mechanisms have been further developed and improved. For instance, Goel et al. [13] develop budget feasible mechanisms that achieve a $\frac{2e-1}{e-1}$ -approximation to the optimal social welfare by exploiting the assumption that one worker has limited contribution to the social welfare. Furthermore, Anari et al. [4] provide a budget feasible mechanism that achieves an approximation ratio of $1 - 1/e \approx 0.63$ by using the “large market” property which assumes that the cost of any worker is relatively small compared to the budget of the buyer. In the recent work of [17], the approximation for budget feasible mechanisms that optimize general monotone submodular functions in large markets has been improved to 2.58. Another related work [7] studies the “dual” problem of maximizing the revenue by selling the maximum independent set of a matroid. They propose a truthful ascending auction in which a seller is constrained to sell objects that form a basis in a matroid.

Some of the results of this work have been independently published in [1]. In particular, in [1] the authors study generalizations of matroid constraints such as finding maximum weighted matroid members, maximum weighted k -D matchings, and maximum weighted independent sets. For these problems they establish that a ρ -approximation can be converted into a deterministic (resp. randomized), truthful, budget feasible mechanism with an approximation ratio of $2\rho + 2$ (resp. $2\rho + 1$). This results into a 4-approximation deterministic and a 3-approximation randomized truthful mechanism for matroid constraints. Our approach and the one of [1] are closely related. The mechanism of [1] also considers elements in decreasing order of cost/valuation ratio. The elements are discarded till the ratio of the current element multiplied by the value of the solution is at most equal to the budget B . However, [1] does not consider more general problems having hereditary property (e.g., k -matroid intersection) while provides results for specific cases as for example maximum matching and maximum 3D-matching.

Finally, in [2] the authors investigate budget feasible mechanism design for symmetric submodular valuations which is a prominent class of non-monotone submodular functions, and provide an $O(1)$ -approximation truthful mechanism that can require an exponential number of value queries. This result has been improved in [3] where the authors present an $O(1)$ -approximation mechanism that runs in polynomial time in the value query model.

2 Preliminaries

Hereditary property A hereditary problem \mathcal{M} is a pair of (E, \mathcal{I}) where E is a finite ground set and $\mathcal{I} \subseteq 2^E$ consists of subsets of the ground set satisfying the following property:

- Hereditary property: If $I \in \mathcal{I}$, then $J \in \mathcal{I}$ for every $J \subset I$.

The set \mathcal{I} does not need to be given by an explicit list of its elements. Indeed, given that $|\mathcal{I}|$ could be exponential in $|E|$, it would be infeasible to use this representation. For this reason, we resort to an independence oracle that, given a subset of elements $T \subseteq E$, returns whether $T \in \mathcal{I}$ or not.

Matroid A matroid \mathcal{M} is a hereditary problem satisfying the following additional property:

- Exchange property: For any pair of sets $I, J \in \mathcal{I}$, if $|I| < |J|$, then there exists an element $e \in J \setminus I$ such that $I \cup \{e\} \in \mathcal{I}$.

The sets in \mathcal{I} are called *independent sets*.

Given a matroid $\mathcal{M} = (E, \mathcal{I})$ and a subset $T \subseteq E$, the *restriction* of \mathcal{M} to T , denoted by $\mathcal{M}|T$, is the matroid in which the ground set is T and the independent sets are the independent sets of \mathcal{M} that are contained in T . That is, $\mathcal{M}|T = (T, \mathcal{I}(\mathcal{M}|T))$ where $\mathcal{I}(\mathcal{M}|T) = \{I \subseteq T : I \in \mathcal{I}\}$. Similarly, the *deletion* of \mathcal{M} , denoted by $\mathcal{M} \setminus T$, is the matroid in which the ground set is $E - T$ and the independent sets are the independent sets of \mathcal{M} that do not contain any element in T . That is, $\mathcal{M} \setminus T = (E - T, \mathcal{I}(\mathcal{M} \setminus T))$ where $\mathcal{I}(\mathcal{M} \setminus T) = \{I \subseteq E - T : I \in \mathcal{I}\}$.

We use the above notation and definitions as well in the case when only hereditary property is satisfied. Such sets are called (E, \mathcal{I}) .

Matroid Intersection A possible generalization of matroids is the *matroid intersection* problem which is defined as follows. Given k matroids $\mathcal{M}_1, \dots, \mathcal{M}_k$, let $\mathcal{M} = (E, \mathcal{I})$ be the matroid where E is the set of common ground elements and $\mathcal{I} = \bigcap_{j=1}^k \mathcal{I}_j$ is the independent sets. In particular, the matroid intersection has to satisfy the hereditary property but does not need to satisfy the exchange property.

Matroid Budget Feasible Mechanisms In an instance of the matroid budget feasible mechanism design problem, we are given a matroid $\mathcal{M} = (E, \mathcal{I})$ consisting of n ground elements (i.e., $|E| = n$), each of whom is associated with a weight $w_e \in \mathbb{R}_+$. Each element $e \in E$ is also associated with a private cost $c_e \in \mathbb{R}_+$, which is only known to the element itself. Our goal is to design a truthful mechanism that gives incentives to elements for declaring their private costs truthfully, and that selects an independent set with the constraint that the total payment given to the elements does not exceed a given budget b . Given an independent set $I \in \mathcal{I}$, its value is defined by $w(I) = \sum_{e \in I} w(e)$. We compare the value of the independent set selected by the mechanism against the value of the maximum-value independent set such that the sum of the true costs of the elements does not exceed the budget.

We use $\mathbf{w} = \langle w_1, \dots, w_n \rangle$ to denote the weights of the ground elements and use $\mathbf{d} = \langle d_1, \dots, d_n \rangle$ to denote the costs declared by the ground elements. Let τ be the maximum-weight element (breaking ties arbitrarily), that is, $w_\tau = \max_{e \in E} w_e$. We assume that $d_e \in \mathbb{R}_+$ and $d_e \leq b$ for any $e \in E$ since elements with costs greater than b cannot be selected by any mechanism due to the budget constraint. This also implies that no element i can improve its utility by declaring $d_i > b$. Given a subset of element $T \subseteq E$, we use \mathbf{w}_{-T} and \mathbf{d}_{-T} to denote the weights and the costs vector excluding elements in T . Similarly, we use \mathbf{w}_T and \mathbf{d}_T to denote the weights and the costs vector only including elements in T . For each element $e \in E$, $\text{bb}(e) = \frac{d_e}{w_e}$ is called the *buck-per-bang* rate for element e .¹

A deterministic mechanism $M = (f, p)$ consists of an allocation function $f : (\mathcal{M}, \mathbf{w}, \mathbf{d}, b) \rightarrow I \in \mathcal{I}$ and a payment function $p : (\mathcal{M}, \mathbf{w}, \mathbf{d}, b) \rightarrow \mathbb{R}_+^n$. In the following, for notational convenience, we will omit braces in the domains of functions f and p . Given the weights and the declared costs of the ground elements, the allocation function returns an independent set in the matroid and the payment function indicates the payments for each element. Let $\mathbf{f}^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b)$ and $\mathbf{p}^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b)$ be the independent set and the payments returned by M , respectively. If an element e is in the independent set obtained by the mechanism M , then $f_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) = 1$. Otherwise, $f_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) = 0$. It is assumed that $p_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) = 0$ if $f_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) = 0$. The utility of an element is the difference between the payment received from the mechanism and its true cost. Formally, the utility of an element e is given by $u_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) = p_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) - f_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) \cdot c_e$.

Individual Rationality A mechanism M is *individually rational* if $p_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) - f_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) \cdot d_e \geq 0$ for any \mathcal{M} , any $\mathbf{w} \in \mathbb{R}_+^n$, any $\mathbf{d} \in \mathbb{R}_+^n$, any $b \in \mathbb{R}_+$ and any element $e \in E$. That is, no element in the selected independent set is paid less than the cost it declared.

Truthfulness A mechanism M is *truthful* if it holds that $u_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}_{-e}, c_e, b) \geq u_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}_{-e}, d_e, b)$, for any \mathcal{M} , any $\mathbf{w} \in \mathbb{R}_+^n$, any $\mathbf{d}_{-e} \in \mathbb{R}_+^{n-1}$, any $d_e \in \mathbb{R}_+$, any $c_e \in \mathbb{R}_+$, $b \in \mathbb{R}_+$ and any $e \in E$, where $\mathbf{d}_{-e} = \langle d_1, \dots, d_{e-1}, d_{e+1}, \dots, d_n \rangle$. When the context is clear, we sometimes abuse some notations. For example, here we write $u_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}_{-e}, c_e, b)$ instead of $u_e^M(\mathcal{M}, \mathbf{w}, \langle \mathbf{d}_{-e}, c_e \rangle, b)$. A truthful mechanism prevents any element to improve its utility by mis-declaring its true cost regardless the costs declared by the other elements.

Budget Feasibility A mechanism M is *budget feasible* if $\sum_{e \in E} p_e^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b) \leq b$, for any \mathcal{M} , $\mathbf{w} \in \mathbb{R}_+^n$, any $\mathbf{d} \in \mathbb{R}_+^n$ and any $b \in \mathbb{R}_+$.

Approximation A mechanism M is *x-approximation* (or *x-approximating*), for some $x \geq 1$, if $w(\mathbf{f}^M(\mathcal{M}, \mathbf{w}, \mathbf{d}, b)) \geq \frac{1}{x} w(\text{OPT}(\mathcal{M}, \mathbf{w}, \mathbf{d}, b))$ for any $\mathbf{w} \in \mathbb{R}_+^n$, $\mathbf{d} \in \mathbb{R}_+^n$ and $b \in \mathbb{R}_+$, where $\text{OPT}(\mathcal{M}, \mathbf{w}, \mathbf{d}, b)$ is the maximum-value independent set in which the total cost of the elements is at most b . We often call $\text{OPT}(\mathcal{M}, \mathbf{w}, \mathbf{d}, b)$ the optimal independent set and simplify it as $\text{OPT}(\mathcal{M}, b)$ throughout the paper when the

¹ $\frac{w_e}{c_e}$ is usually known as the *bang-per-buck* rate. To simplify the presentation, we call $\frac{d_e}{w_e}$ the *buck-per-bang* rate.

weights and the costs of elements are clear. Similarly, we use $\text{MAX}(\mathcal{M}, \mathbf{w})$, shortened by $\text{MAX}(\mathcal{M})$, to denote the maximum-value independent set in \mathcal{M} without the budget constraint. It is well known that there exists a greedy algorithm that computes the maximum-value independent set of a matroid \mathcal{M} in time $\text{poly}(n)$ which uses a $\text{poly}(n)$ number of calls to an independence oracle. Finally, we use $\text{APX}(\mathcal{M})$ to denote an α -approximation (where $\alpha \geq 1$) to the maximum-value set of a problem \mathcal{M} without the budget constraint. In particular, already in the case of matroid intersection finding the maximum-value common independent set of a matroid intersection is NP-hard when more than three matroids are involved.

Simplifying notations From now on, to avoid heavy notation, we sometimes simplify it. For example, we will write f^M, f_e^M, p^M, p_e^M when the input of the mechanism is clear. Moreover, we will use $\text{OPT}(\mathcal{M} \setminus T, b)$ instead of $\text{OPT}(\mathcal{M} \setminus T, \mathbf{w}_{-T}, \mathbf{d}_{-T}, b)$ to denote the optimal independent set in the matroid $\mathcal{M} \setminus T$. Similarly, we will use $\text{OPT}(\mathcal{M}|T, b)$ instead of $\text{OPT}(\mathcal{M}|T, \mathbf{w}_T, \mathbf{d}_T, b)$ to denote the optimal independent set in the matroid $\mathcal{M}|T$. Furthermore, we will use $\text{MAX}(\mathcal{M} \setminus T)$ instead of $\text{MAX}(\mathcal{M} \setminus T, \mathbf{w}_{-T})$ to denote the maximum-value independent set in $\mathcal{M} \setminus T$ without considering the costs of the elements and the budget. Finally, we will use $\text{APX} \mathcal{M} \setminus T$ to denote an α -approximation to the maximum-value independent set of a matroid intersection $\mathcal{M} \setminus T$ without considering the costs of the elements and the budget.

3 A Mechanism for Hereditary Problems

In this section we consider problems that have hereditary property. In particular, this set of problems includes matroid intersection, which in general is NP-hard when more than three matroids are involved. Some interesting cases of matroid intersection problems can be solved efficiently (i.e., they can be formulated as the intersection of two matroids), for example, matchings in bipartite graphs, arborescences in directed graphs, spanning forests in undirected graphs, etc.

For general hereditary problems, our main result is the following. Given a deterministic polynomial time blackbox approximation algorithm APX that achieves an α -approximation (for some $\alpha \geq 1$) to the problem of finding the maximum-value set of a hereditary problem \mathcal{M} , we provide a polynomial time, individually rational, truthful and budget feasible deterministic mechanism that is a $(3\alpha + 1)$ -approximation to the maximum-value set of the problem \mathcal{M} with the constraint that the sum of the costs of its elements is at most a given budget.

It is known that the VCG payment rule does not preserve the property of truthfulness in the presence of approximated solutions (i.e., non-optimal outcome). However, unlike the VCG mechanism, we show that Mechanism 1 is truthful when APX is used.

Before providing the mechanism, we discuss some intuition that has guided us in the design of Mechanism 1. First, imagine that there exists an element with a very high weight. Clearly, any set without this element would result in a poor value compared to the optimal set. In this case, such element may strategically declare a high cost in order to increase its utility, as it knows that any good-approximating mechanism has to select it. To avoid that this happens, we remove the element τ (i.e., the

element with the largest weight) from the problem via the deletion operation and compare it with the set computed later by the mechanism. Second, we observe that some of the existing budget feasible mechanisms ([13, 23]) adopt proportional payment schemes, where elements (i.e., agents) are paid proportionally and according to their contribution in the solution. In other words, in a proportional payment scheme there is an uniform price such that the payments for the elements in the solution are the products of their contribution and this price. In addition, greedy algorithms are commonly used in matroid systems.

Our mechanism combines these observations and works as follows. It first orders the elements from the largest to the smallest buck-per-bang. Then, it starts from a high price, i.e. the highest buck-per-bang value, and computes the (approximated) maximum-value set of the hereditary problem. If there is enough budget to pay this set at the current price then it proceeds to the final step. Otherwise, it reduces the price (i.e. it considers the next buck-per-bang value of the ordering), removes the previous element with higher buck-per-bang value from the instance of the hereditary problem and computes the (approximated) maximum-value set of the reduced instance of the hereditary problem. We notice that the current buck-per-bang value is an upper bound of the payment on each contribution in the next iteration. The mechanism performs the procedure described above until the payment of the current (approximated) maximum-value set is within the budget b . In the final step, the mechanism returns the best solution between the set of selected elements and the element τ of largest weight.

Mechanism 1: A budget feasible mechanism for procuring a set in the hereditary problem \mathcal{M}

Input: $\mathcal{M} = (E, \mathcal{I})$, w, d, b .
Output: f, p

- 1 Let τ be the maximum-weight element (breaking ties arbitrarily);
- 2 Sort elements in $E - \tau$ in a non-increasing order of buck per bang, i.e. $bb(i) \geq bb(j)$ if $i < j$, break ties arbitrarily;
- 3 Let $i = 1$ and $T = \emptyset$;
- 4 Set $r = bb(i)$;
- 5 **while** $w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) \cdot r > b$ **do**
- 6 $T = T \cup \{i\}$ and $i = i + 1$;
- 7 $r = bb(i)$;
- 8 $r = \min\{\frac{b}{w(\text{APX}(\mathcal{M} \setminus (T \cup \tau)))}, bb(i - 1)\}$;
- 9 **if** $w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) > w_\tau$ **then**
- 10 For each $e \in E$, if $e \in \text{APX}(\mathcal{M} \setminus (T \cup \tau))$, $f_e = 1$ and $p_e = r \cdot w_k$. Otherwise, $f_e = 0$ and $p_e = 0$;
- 11 **else**
- 12 $f_\tau = 1, p_\tau = b$. For each edge $e \in E - \tau$, $f_e = 0, p_e = 0$;
- 13 **return** f, p ;

Theorem 1 *Mechanism 1 is a deterministic, polynomial time, individually rational, truthful and budget feasible mechanism that is a $(3\alpha + 1)$ -approximation to the optimal set of the given hereditary problem (e.g., matroid intersection) for a given budget.*

Proof Complexity It is not difficult to see that Mechanism 1 performs a $\text{poly}(n)$ number of steps. Moreover, we assume that we are given a deterministic polynomial time (with respect to n) blackbox approximation algorithm $\text{APX}(\mathcal{M})$ that achieves an α -approximation to the problem of finding the maximum-value set of the hereditary problem \mathcal{M} . It follows that Mechanism 1 has complexity $O(\text{poly}(n))$.

Approximation Recall that T is the set of elements removed from the problem. $\text{APX}(\mathcal{M} \setminus (T \cup \tau))$ is the set found when Mechanism 1 stops, and it is an α -approximation to the maximal-value set of $\mathcal{M} \setminus (T \cup \tau)$. The roadmap of the proof is to first show that the set $\text{APX}(\mathcal{M} \setminus (T \cup \tau))$ well approximates the optimal set in $\mathcal{M} \setminus \tau$. Next, we show that returning the maximum between τ and $\text{APX}(\mathcal{M} \setminus (T \cup \tau))$ gives the $(3\alpha + 1)$ -approximation to the optimal set in \mathcal{M} with the budget constraint.

Lemma 1 *Given any \mathcal{M} , w , d , b , when Mechanism 1 stops, it holds that:*

$$w(\text{OPT}(\mathcal{M} \setminus \tau, b)) \leq 2\alpha w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) + \alpha w_\tau.$$

Proof It is trivial to see that this lemma holds when τ is the only element in \mathcal{M} . In the rest of the proof we use an idea which was also used in [13]. We distinguish between two cases depending on whether the full budget b is spent or not. Consider that the set $E - \{\tau\}$ is partitioned into two disjoint sets $E - \{\tau\} - T$ and T . We notice that the maximum-value set $w(\text{OPT}(\mathcal{M} \setminus \tau, b))$ can be bounded by:

$$w(\text{OPT}(\mathcal{M} \setminus \tau, b)) \leq w(\text{OPT}(\mathcal{M}|T, b)) + w(\text{OPT}(\mathcal{M} \setminus (T \cup \tau), b)).$$

In fact, let $A = \text{OPT}(\mathcal{M} \setminus \tau, b) \cap T$ be the subset of $\text{OPT}(\mathcal{M} \setminus \tau, b)$ containing only elements belonging to T . Moreover, let $B = \text{OPT}(\mathcal{M} \setminus \tau, b) \cap (E - \{\tau\} - T)$ be the subset of $\text{OPT}(\mathcal{M} \setminus \tau, b)$ containing only elements belonging to $E - \{\tau\} - T$. We have that $\text{OPT}(\mathcal{M} \setminus \tau, b) = A \cup B$. Moreover, by the hereditary property and the definition of restriction and deletion, we have that $A \in \mathcal{I}(\mathcal{M}|T)$ and $B \in \mathcal{I}(\mathcal{M} \setminus (T \cup \tau))$.

Since the buck-per-bang is at least r for every element in T , we have that the weight of the optimal set given a budget b in $\mathcal{M}|T$ (i.e., $w(\text{OPT}(\mathcal{M}|T, b))$) is at most b/r . Moreover, we have that $w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) \geq b/r$. Thus, when the full budget is spent, we get that:

$$\begin{aligned} w(\text{OPT}(\mathcal{M} \setminus \tau, b)) &\leq w(\text{OPT}(\mathcal{M}|T, b)) + w(\text{OPT}(\mathcal{M} \setminus (T \cup \tau), b)) \\ &\leq \frac{b}{r} + \alpha w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) \\ &\leq (\alpha + 1)w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))). \end{aligned}$$

We now turn to the case where some budget is left in Mechanism 1. The idea here is to bound the budget left. Since Mechanism 1 does not stop when $r = \text{bb}(i - 1)$, this implies that the set returned by APX was not budget feasible at the previous iteration. This further implies that the maximum-value set is not budget feasible either if the payment per weight is r . After removing the element $i - 1$, the set returned by APX becomes budget feasible when $r = \text{bb}(i - 1)$. These together imply that:

$$\begin{aligned} w(\text{MAX}(\mathcal{M} \setminus (T' \cup \tau)))bb(i - 1) &\geq w(\text{APX}(\mathcal{M} \setminus (T' \cup \tau)))bb(i - 1) \\ &> b \\ &> w(\text{APX}(\mathcal{M} \setminus (T \cup \tau)))bb(i - 1), \end{aligned}$$

where $T' = T - \{i - 1\}$. Moreover, since the sum of $w(\text{APX}(\mathcal{M} \setminus (T \cup \tau)))$ and $w(i - 1)$ is at least an $\frac{1}{\alpha}$ fraction of $w(\text{MAX}(\mathcal{M} \setminus (T' \cup \tau)))$, we get that:

$$(w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) + w_{i-1})bb(i - 1) \geq \frac{1}{\alpha}w(\text{MAX}(\mathcal{M} \setminus (T' \cup \tau)))bb(i - 1) > \frac{b}{\alpha}.$$

Hence, we get that:

$$w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) + w_{i-1} > \frac{b}{\alpha bb(i - 1)}.$$

Overall, it follows that:

$$\begin{aligned} \text{OPT}(\mathcal{M} \setminus \tau, b) &\leq w(\text{OPT}(\mathcal{M}|T, b)) + w(\text{OPT}(\mathcal{M} \setminus (T \cup \tau), b)) \\ &\leq \frac{b}{bb(i - 1)} + \alpha w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) \\ &\leq 2\alpha w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))) + \alpha w_{i-1}. \end{aligned}$$

By substituting w_{i-1} with w_τ , we get the proof. □

Next, we show that returning the maximum between τ and $\text{APX}(\mathcal{M} \setminus (T \cup \tau))$ is a $(3\alpha + 1)$ -approximation to the optimal set in \mathcal{M} .

Lemma 2 *Given any $\mathcal{M}, \mathbf{w}, \mathbf{d}, b$, the set returned by Mechanism 1, i.e., the maximum between τ and $\text{APX}(\mathcal{M} \setminus (T \cup \tau))$, is a $(3\alpha + 1)$ -approximation to the optimal set in \mathcal{M} .*

Proof First notice that the optimal set in \mathcal{M} is bounded by

$$w(\text{OPT}(\mathcal{M}, b)) \leq w_\tau + w(\text{OPT}(\mathcal{M} \setminus \tau, b)).$$

Moreover, by Lemma 1, we have that:

$$w(\text{OPT}(\mathcal{M} \setminus \tau, b)) \leq \alpha w_\tau + 2\alpha w(\text{APX}(\mathcal{M} \setminus (T \cup \tau))).$$

Therefore, the maximum between τ and $\text{APX}(\mathcal{M} \setminus (T \cup \tau))$ approximates the optimal set within a factor of $(3\alpha + 1)$. □

Truthfulness Truthfulness of the mechanism relies on Myerson’s Lemma for single parameter environments [21]. We show that the allocation function of Mechanism 1 is monotone when APX is used by proving that the allocation function is non-decreasing if the bid of the agent is decreased. The reason behind is that our mechanism works in a greedy fashion and at each iteration the cost declared by elements *does not* affect the set computed by the mechanism.

We start by showing that the element τ cannot benefit by mis-declaring its true cost.

Lemma 3 *The element with the maximum weight, i.e., the element τ , could not improve its utility by declaring a cost $d_\tau \neq c_\tau$.*

Proof On the one hand, if Mechanism 1 returns the element τ when it declares its true cost, then τ gets a payment of b . Thus, there is no incentive for it to declare a different cost. On the other hand, when Mechanism 1 returns $\text{APX}(\mathcal{M} \setminus (T \cup \tau))$, even if τ mis-declares its true cost it does not change the outcome of the mechanism because it would be still the element with the largest weight and $w_\tau < w(\text{APX}(\mathcal{M} \setminus (T \cup \tau)))$. \square

Next, we show that no element in $E - \tau$ can improve its utility by mis-declaring its true cost. The proof relies on the analysis of different cases. The first case shows that no element which is removed by Mechanism 1 can benefit by mis-declaring its true cost.

Lemma 4 *Assume that an element $k \in T$ (i.e., the element k is removed by the mechanism) when it declares its cost truthfully. Then, the element k could not improve its utility by declaring a cost $d_k \neq c_k$.*

Proof Since $k \in T$, it means that the element k is not in the set returned by Mechanism 1. Hence, its utility is zero. This implies that, when the element k is considered, that is, $r = \text{bb}(k)$, it holds that $w(\text{APX}(\mathcal{M} \setminus (T^k \cup \tau))) \cdot r > b$, where T^k denotes the set of elements removed until k is considered. We now distinguish between two different cases.

- Suppose that the element k declares a higher cost $d_k > c_k$, and it is considered earlier at the h -th iteration with $h \leq k$. Equivalently speaking, k becomes the element with the h -th largest buck-per-bang rate. In this case, Mechanism 1 would not stop until k is considered. Notice that the declared costs are not involved in the computation of the approximated maximum-value sets. Indeed, the computed approximated maximum-value sets are exactly the same as declaring truthfully until the h -th iteration. Moreover, the approximated maximum-value set is also the same in the h -th iteration since the remaining elements in the problem are the same. Given that r in the h -th iteration is equal to or greater than in the previous iterations, it follows that the approximated maximum-value set returned by APX is not budget feasible. This implies that element k would be removed from the problem and it would never be included in a set by Mechanism 1. Therefore, its utility is zero.

We use a similar argument to show that the element k could not improve its utility by declaring a smaller cost.

- Suppose that the element k declares a smaller cost $d_k < c_k$ and it is considered at the h -th iteration with $h \geq k$. Let us focus on how Mechanism 1 performs. Until the k -th iteration, the approximated maximum-value sets are the same as k declaring its cost truthfully and they are not budget feasible. Next, the approximated maximum-value set in the k -th iteration is also the same as before. If the set is budget feasible, then we know that $\frac{b}{w(\text{APX}(\mathcal{M} \setminus (T \cup \tau)))}$ is strictly less than $bb(k) = \frac{c_k}{w_k}$. This holds because the mechanism does not terminate at $r = \frac{c_k}{w_k}$ when the element k declares truthfully. Therefore, even if the element k is in this set, the payment will be strictly less than its true cost. On the other hand, if the set is not budget feasible, then the mechanism would update its upper bound of payment for each contribution in the next iteration. The new upper bound is at most $bb(k) = \frac{c_k}{w_k}$. This implies that the element k would not get a payment greater than its true cost.

□

We now show that the no remaining element (i.e., an element that is not put in the set T by Mechanism 1) that is not included in the set returned by the mechanism can benefit by mis-declaring its true cost.

Lemma 5 *Assume that an element k is in $E - \tau - T - \text{APX}(\mathcal{M} \setminus (T \cup \tau))$ when it declares its cost truthfully. Then, the element k could not improve its utility by declaring a cost $d_k \neq c_k$.*

Proof Since $k \notin \text{APX}(\mathcal{M} \setminus (T \cup \tau))$, we know that the utility of the element k is zero. Suppose that the mechanism terminates at the h -th round when the element k declares its cost truthfully. We now distinguish among three different cases.

- Suppose that the element k declares a higher cost $d_k > c_k$ and it is considered by Mechanism 1 at an l -th iteration with $l < h$. In this case, Mechanism 1 would not stop before or at the l -th iteration since the computed approximated maximum-value sets are exactly the same, and thus they are not budget feasible, as in the case where k declares its cost truthfully. This implies that the element k would be never included in a set by Mechanism 1 when it declares a larger cost. Therefore, its utility is still zero.
- Suppose that the element k declares a cost $d_k \neq c_k$ and it is considered by Mechanism 1 at the h -th iteration. In this case, Mechanism 1 would not stop before the h -th iteration because the approximated maximum-value sets are not feasible. The approximated maximum-value set in the h -th iteration is the same as in the case where k declares its cost truthfully since the remaining elements are the same. Therefore, if the set is budget feasible, then the mechanism computes the same set and payments. Otherwise, the element k is removed since it must be the element with the h -th largest buck-per-bang rate. It follows that the element k cannot benefit also in this case.

- Finally, suppose that the element k declares a cost $d_k \neq c_k$ and it is considered after the h -th iteration. In this case, Mechanism 1 would compute the same set and payments and again the element k would get utility zero. □

Finally, it remains to show that no element belonging to the set returned by Mechanism 1 can benefit by mis-declaring its true cost. The proof of the following lemma uses a similar argument to the one used in Lemma 5.

Lemma 6 *Assume that an element k is in $\text{APX}(\mathcal{M} \setminus (T \cup \tau))$ when it declares its cost truthfully. Then, the element k could not improve its utility by declaring a cost $d_k \neq c_k$.*

Proof Suppose that the mechanism terminates at the h -th round when the element k declares its cost truthfully. Since $k \in \text{APX}(\mathcal{M} \setminus (T \cup \tau))$, we know that the payment of the element k is $r \cdot w_k$. Hence, its utility is $r \cdot w_k - c_k$. We distinguish among three different cases.

- Suppose that the element k declares a higher cost $d_k > c_k$ and it is considered at the l -th iteration, with $l < h$. In a similar way to Lemma 5, in this case Mechanism 1 would not stop before or at the l -th iteration since the computed approximated maximum-value sets are exactly the same and they are not budget feasible. This implies that the element k would be removed from the matroid and it will never be included in a set by Mechanism 1. Therefore, its utility becomes zero.
- Suppose that the element k declares a cost $d_k \neq c_k$ and it is considered at the h -th iteration. In this case, Mechanism 1 would not stop before the h -th iteration because the approximated maximum-value sets are not feasible. The approximated maximum-value set in the h -th round is the same as in the case where k declares its cost truthfully since the remaining elements are the same. Therefore, if the set is budget feasible, the mechanism would compute the same set and payments. Otherwise, the element k would be removed. It follows that the element k cannot benefit in any case.
- Finally, suppose that the element k declares a cost $d_k \neq c_k$ and it is considered after the h -th iteration. In this case, Mechanism 1 would compute the same set and payments, and again the element k would not benefit. □

This concludes the proof of the theorem. □

3.1 Budget Feasible Mechanisms for Simple Matroids

We now show how to use the mechanism defined in the previous section to obtain a deterministic, polynomial time, individually rational, truthful and budget feasible mechanism that is a 4-approximation to the optimal independent set of a matroid. The mechanism is directly obtained from Mechanism 1 by exchanging APX with

MAX, the optimal algorithm for matroids. We point out that a very similar mechanism achieves a 4-approximation also for the case when, instead of a matroid, we are given an undirected weighted (general) graph where the selfish agents are the edges of the graph and the buyer wants to procure a matching under the given budget that yields the largest possible value for him.

Theorem 2 *Mechanism 1 is a deterministic, polynomial time, individually rational, truthful and budget feasible mechanism that is a 4-approximation to the optimal independent set of a matroid given a budget.*

Proof We can prove the truthfulness of Mechanism 1 for matroids by applying exactly the same argument used in Theorem 1. Moreover, we get the 4-approximation just by using the same argument used in Theorem 1 with $\alpha = 1$. We only briefly discuss the complexity.

Complexity It is not difficult to see that Mechanism 1 performs a $poly(n)$ number of steps. Moreover, as described in Sect. 2, there exists a greedy algorithm that computes the maximum-value independent set of a matroid \mathcal{M} in time $poly(n)$ which uses a $poly(n)$ number of calls to an independence oracle, that is, $MAX(\mathcal{M})$ can be computed in polynomial time (with respect to n) by using an independence oracle. It follows that Mechanism 1 for matroids has complexity $O(poly(n))$. \square

4 Conclusions

We present in this paper a budget feasible mechanism for general hereditary problems including intersection of matroids. Our main result is a $(3\alpha + 1)$ -approximation for matroid intersection given an α -approximation algorithm for the optimization version of the problem. Our algorithm is simple and we hope it will provide guidelines for designing novel and more efficient algorithms for this class of problems.

Budget feasible mechanism design was an area of active research with huge advancement in the last few years but there is still a number of problems that resist to elegant solutions with small approximation factor. We mention, for example, XOS valuations and non-monotone submodular functions, and any combination with matroid constraints of valuation functions other than additive.

Funding Open access funding provided by Università degli Studi di Roma La Sapienza within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is

not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Amanatidis, G., Birmpas, G., Markakis, E.: Coverage, matching, and beyond: new results on budgeted mechanism design. In: Proceedings of the 12th International Conference on Web and Internet Economics (WINE), pp 414–428 (2016)
2. Amanatidis, G., Birmpas, G., Markakis, E.: On budget-feasible mechanism design for symmetric submodular objectives. In: Proceedings of the 13th International Conference on Web and Internet Economics (WINE), pp 1–15 (2017)
3. Amanatidis, G., Kleer, P., Schäfer, G.: Budget-feasible mechanism design for non-monotone submodular objectives: offline and online. In: Proceedings of the twentieth ACM conference on Economics and Computation (EC), pp 901–919 (2019)
4. Anari, N., Goel, G., Nikzad, A.: Mechanism design for crowdsourcing: an optimal $1-1/e$ competitive budget-feasible mechanism for large markets. In: 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp 266–275 (2014)
5. Ausubel, L.M.: An efficient ascending-bid auction for multiple objects. *Am. Econ. Rev.* **94**, 1452–1475 (2004)
6. Bei, X., Chen, N., Gravin, N., Lu, P.: Budget feasible mechanism design: from prior-free to bayesian. In: Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing (STOC), pp 449–458 (2012)
7. Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R.V.: An ascending Vickrey auction for selling bases of a matroid. *Oper. Res.* **59**(2), 400–413 (2011)
8. Chan, H., Chen, J.: Truthful multi-unit procurements with budgets. In: The proceedings of the 10th International Conference on Web and Internet Economics (WINE), pp 89–105 (2014)
9. Chen, Ning., Gravin, Nick., Lu, Pinyan.: On the approximability of budget feasible mechanisms. In: Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp 685–699 (2011)
10. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* **11**(1), 17–33 (1971)
11. Demange, G., Gale, D., Sotomayor, M.: Multi-item auctions. *J. Polit. Econ.* **94**, 863–872 (1986)
12. Goel, G., Mirrokni, V., Leme, R.P.: Polyhedral clinching auctions and the adwords polytope. *J. ACM (JACM)* **62**(3), 18 (2015)
13. Goel, G., Nikzad, A., Singla, A.: Allocating tasks to workers with matching constraints: truthful mechanisms for crowdsourcing markets. In: Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web Companion, pp 279–280 (2014)
14. Gravin, N., Jin, Y., Lu, P., Zhang, C.: Optimal budget-feasible mechanisms for additive valuations. In: Proceedings of the 2019 ACM Conference on Economics and Computation, EC '19, pp 887–900, New York, NY, USA, Association for Computing Machinery (2019)
15. Groves, T.: Incentives in teams. *Econom. J. Econom. Soc.* **41**, 617–631 (1973)
16. Karlin, A.R., Kempe, D.: Beyond VCG: frugality of truthful mechanisms. In: 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp 615–624. IEEE, (2005)
17. Khalilabadi, P.J., Tardos, É.: Simple and efficient budget feasible mechanisms for monotone submodular valuations. In: Proceedings of Web and Internet Economics—14th International Conference, WINE 2018, Oxford, UK, December 15–17, 2018, pp 246–263 (2018)
18. Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: Proceedings of the Forty-fourth annual ACM symposium on Theory of computing, pp 123–136. ACM, (2012)
19. Krysta, P., Zhang, J.: House markets with matroid and knapsack constraints. In: Proceedings of The 43rd International Colloquium on Automata, Languages and Programming (ICALP) (2016)
20. Leonardi, S., Monaco, G., Sankowski, P., Zhang, Q.: Budget feasible mechanisms on matroids. In: Proceedings of the 19th International Conference on Integer Programming and Combinatorial Optimization (IPCO), pp 368–379 (2017)
21. Myerson, R.B.: Optimal auction design. *Math. Oper. Res.* **6**(1), 58–73 (1981)

22. Schrijver, A.: Combinatorial Optimization. Algorithms and Combinatorics, vol. 24. Springer, Berlin (2003)
23. Singer, Y.: Budget feasible mechanisms. In: 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp 765–774 (2010)
24. Tse, D.N.C., Hanly, S.V.: Multiaccess fading channels. I. Polymatroid structure, optimal resource allocation and throughput capacities. *IEEE Trans. Inf. Theory* **44**(7), 2796–2815 (1998)
25. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16**(1), 8–37 (1961)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Stefano Leonardi¹  · Gianpiero Monaco² · Piotr Sankowski³ · Qiang Zhang¹

Gianpiero Monaco
gianpiero.monaco@univaq.it

Piotr Sankowski
sank@mimuw.edu.pl

Qiang Zhang
csqzhang@gmail.com

¹ Sapienza University of Rome, Rome, Italy

² University of L'Aquila, L'Aquila, Italy

³ University of Warsaw, Warsaw, Poland