Mechanisms for Complement-Free Procurement

Shahar Dobzinski*
Department of Computer Science
Cornell University, Ithaca NY 14853
shahar@cs.cornell.edu

Christos Papadimitriou†
Computer Science Division
University of California at Berkeley, CA, 94720
christos@cs.berkeley.edu

Yaron Singer‡
Computer Science Division
University of California at Berkeley, CA, 94720
yaron@cs.berkeley.edu

ABSTRACT

We study procurement auctions when the buyer has complement-free (subadditive) objectives in the budget feasibility model [18]. For general subadditive functions we give a randomized universally truthful mechanism which is an $O(\log^2 n)$ approximation, and an $O(\log^3 n)$ deterministic truthful approximation mechanism; both mechanisms are in the demand oracle model. For cut functions, an interesting case of nonincreasing objectives, we give both randomized and deterministic truthful and budget feasible approximation mechanisms that achieve a constant approximation factor.

Categories and Subject Descriptors

F.2.8 [**Analysis of Algorithms and Problem complexity**]: Miscellaneous

General Terms

Theory

Keywords

Procurement Auctions, Incentive Compatibility, Truthfulness, Budget Feasibility

1. INTRODUCTION

When a principal wishes to buy items or services provided by strategic agents, her goal is to maximize an objective that assigns a valuation to any set of items. Since the agents may exaggerate their costs one is interested in designing *truthful mechanisms,* in which the allocation and prices are set in such a way so that it is provably in each agent's interest to reveal her true cost.

Such problems are extensively studied in the framework of *frugality* (see, for example, [1]), where the mechanism designer's aim is to minimize payments. The frugality approach deals with buyer objectives with 0-1 values (e.g., the buyer wants the set of edges to be a spanning tree of a graph, and all spanning trees are equally desirable).[1]

An alternative to frugality which encompasses much more general objectives, is the *budget feasibility* framework introduced in [18] and subsequently studied in [6, 13, 10]. In budget feasibility, our goal is to optimize the buyer's objective under a budget constraint.

A budget feasibility setting is fully defined by the buyer's objective, a function mapping subsets of items to the reals. The important question to ask here is, for which class objectives one can obtain truthful mechanisms with reasonable approximation guarantees? A first step in this direction was taken recently in [18], where truthful mechanisms with constant-factor approximation are developed for the important class of *submodular objective functions. In this paper we extend this result to much more general classes of objectives.*

What are the limits of our ambition here? One first observation is that even for very simple *superadditive* objectives (that is, objectives in which $V(S \cup T) > V(S) + V(T)$ for some sets S, T) mechanisms of the desired kind (truthful of bounded approximation ratio) do not exist. Consider for example the objective in which there is a crucial item a^* and the valuation $V(S)$ is 1 if $a^* \in S$ and 0 otherwise. Then it is easy to see that for this objective there can be no approximation guarantee, because the seller of the crucial item can always extract the whole budget (publicly known). In view of this, we consider objectives that are *complement-free*, or *subadditive*, valuations obeying $V(S \cup T) \leq V(S) + V(T)$. This is a class that is a substantial generalization of submodular valuation functions. We seek truthful mechanisms with reasonable approximation guarantees for complement-free objectives. As the example above shows, in procure-

*Supported by an Alfred P. Sloan Foundation Fellowship and a Microsoft Research New Faculty Fellowship.

†Supported in part by NSF grant ccf-0635319.

‡Supported in part by a Microsoft Research graduate fellowship.

EC'11, June 5–9, 2011, San Jose, California, USA.

[1]It is problematic to extend the frugality framework beyond 0-1 objectives, because that would bring us in the tricky realm of *multi-objective optimization.* This reduces the problem to budget feasibility: fixing the budget and maximizing the objective function.

ment auctions often desirable mechanisms are impossible *even with unbounded computational resources.* Therefore, we focus on the question of existence of a mechanism of the desired kind, with less attention on computational efficiency (even though, as it happens, our algorithms are polynomial time).

1.1 Our Results

For complement-free (subadditive) objectives we first give a randomized $O(\log^2 n)$ approximation mechanism that is universally truthful and budget feasible. We derandomize this mechanism and present a deterministic mechanism for this domain with an $O(\log^3 n)$ approximation guarantee. Since subadditive functions may require representation that is exponential in the number of agents, we assume we are given access to a *demand* oracle. Since the weaker value query oracles result in inapproximability as shown in [18], a stronger oracle model such as the one used here is required.

We also examine another class of objectives, namely *cut functions:* when agents represent vertices on a graph, a cut function is the cardinality of the cut procured. The optimization problem is a variant of the celebrated MAX-CUT problem, and the setting provides a nontrivial case valuation functions that can decrease in value as items are added. For such functions we give a randomized constant factor approximation mechanism that is universally truthful and budget feasible, as well as a deterministic constant factor approximation mechanism.

The deterministic mechanisms, for both classes of objectives, are based on their randomized analogues. In both cases derandomization in a manner that preserves truthfulness is nontrivial and requires constructing "monotone estimators" that are crucial for obtaining bounded approximation guarantees. Each construction uses a different technique, though their underlying principle is similar. We discuss these issues in the appropriate sections.

1.2 Related Work

In the past decade procurement auctions have been extensively studied in algorithmic game theory. The first line of work focused on minimizing total *cost*, not payments [17, 11, 12] using variants of the VCG mechanism. To address the issue of overpayments, the frugality framework was initiated in [1] where the problems typically studied assume there is some possible feasible set of solutions (e.g. all spanning trees in a graph) and that the mechanism designer's objective is to obtain a feasible set (without any preference among the sets) at minimal cost. The framework suggests designing truthful mechanisms with payment schemes that are guaranteed to be within a certain factor of a given benchmark (e.g. minimal payments at Nash equilibrium or a variant as suggested in [14]). In the past decade procurement has been studied in the frugality framework in various domains such as path auctions, spanning trees and vertex cover [19, 9, 14, 8, 4]. Recently, a general scheme has been independently suggested by [5] and [15], which uses spectral techniques to obtain desirable guarantees against the different benchmarks.

The budget feasibility framework has been recently initiated in [18], and this is the model we follow in this paper. The framework addresses the question of designing truthful mechanisms under budget constraint on payments that have provable guarantees in term of the buyer's objective function. The main result in [18] shows that for any increasing submodular function there is universally truthful constant factor approximation mechanism that is budget feasible. In [6] improved approximation ratios were achieved for the submodular case as well as various problems within the submodular class through use of sophisticated analysis. The framework has been adopted in the study of designing auctions for procuring private data [13] dynamic auctions [10] as well as other domains [6].

1.3 Open Questions

In this work we show that there exists an $O(\log^2 n)$ randomized mechanism that uses a polynomial number of demand queries. We conjecture that a constant approximation ratio requires using an exponential number of demand queries. A fundamental question is whether, regardless of computational constraints, a constant-factor budget feasible mechanism exists for subadditive function.

For cut functions, our goal in this work is in proving the existence of a mechanism with a constant factor approximation factor. It would be interesting to determine the best approximation factor achievable for this class. Another interesting open question is to determine the largest class for which there exist constant ratio approximation mechanisms. Can our mechanism be extended to handle all (not necessarily increasing) submodular functions?

Finally, we present techniques for derandomizing truthful mechanisms in a manner that preserves truthfulness. Extending our methods to derandomize mechanisms in other classes is an intriguing open question.

2. THE MODEL

We study procurement auctions where there are n items, each held by a single agent that associates a private cost $c_i \in \mathcal{R}_{\geq 0}$ with the item. The set of agents is denoted by $\mathcal{N}$ and we refer to agents and items interchangeably. The buyer has a public budget $B \in \mathcal{R}_{\geq 0}$ and a public valuation function $V : 2^{[n]} \to \mathcal{R}_{\geq 0}$ over the subsets of items.

The goal in this setting is to design mechanisms that approximate the optimal solution of the valuation function under the budget constraint. A mechanism $\mathcal{M} = (f, p)$ consists of an allocation function $f : \mathcal{R}^n_{\geq 0} \to 2^{[n]}$ and a payment function $p : \mathcal{R}^n_{\geq 0} \to \mathcal{R}^n_{\geq 0}$. The allocation function f maps a set of n bids to a selected subset of agents. The payment function p returns a vector $p_1, \ldots, p_n$ of payments to the agents. We seek normalized ($i \notin S$ implies $p_i = 0$), individually rational ($p_i \geq c_i$) mechanisms with no positive transfers ($p_i \geq 0$).

Truthfulness. In our model we assume the participating agents are strategic and may report false costs if it is in their benefit. We therefore seek truthful mechanisms so that reporting the true costs is a dominant strategy for agents. Formally, a mechanism $\mathcal{M} = (f, p)$ is *truthful* (*incentive compatible*) if for every $i \in \mathcal{N}$ with cost c_i and bid c'_i, and every set of bids by $\mathcal{N} \setminus \{i\}$ we have $p_i - s_i \cdot c_i \geq p'_i - s'_i \cdot c_i$, where (s_i, p_i) and (s'_i, p'_i) are allocation ($s_i, s'_i \in \{0, 1\}$) and payment pairs when the bidding is c_i and c'_i, respectively. A mechanism that is a randomization over truthful mechanisms is *universally truthful*.

As each bidder has a single private value, we shall rely on Myerson's well-known characterization for truthfulness in single parameter domains [16] that states that a mechanism $\mathcal{M} = (f, p)$ is truthful if and only if it is *monotone* and uses *threshold payments.* The mechanism is monotone if $\forall i \in [n]$, if $c'_i \leq c_i$ then $i \in f(c_i, c_{-i})$ implies $i \in f(c'_i, c_{-i})$ for every c_{-i}; winners are paid threshold payments if payment to each selected agent is $\inf\{c_i : i \notin f(c_i, c_{-i})\}$.

Budget Feasibility. We require that the mechanism is budget feasible: the mechanism's payments should not exceed the budget: $\sum_i p_i \leq B$.

The objective is to maximize the function under the budget, i.e. find the subset $S \in \{T | \sum_{i\in T} c_i \leq B\}$ for which $V(S)$ is maximized, under the constraint that the *payments* (not costs) are within the budget. We want the allocated subset to yield a high value for the buyer. When we are unable to output the optimal solution, we are interested in finding a set that has a value that is close to the optimal (when all costs are known) solution. Formally, for $\alpha \geq 1$ we say that a mechanism is α-approximate if in every instance the mechanism allocates to a set S such that $V(S^*) \leq \alpha V(S)$, where S^* is the optimal solution when all costs are known. As usual, when dealing with randomization we seek mechanisms that yield a good approximation in expectation over the internal random coins of the mechanism.

The mechanisms we construct have the additional property that the functions f and p can be computed in polynomial time. In cases where the valuation function requires exponential data to be represented, we take the common "black-box" approach and assume that V is represented by an oracle. We will define the exact model in the relevant sections, but as shown our mechanisms make polynomially many queries to the oracle.

3. SUBADDITIVE FUNCTIONS

A valuation function $V : 2^{[n]} \to R$ is subadditive if $V(S \cup T) \leq V(S) + V(T)$. A naive representation of subadditive functions requires exponential space in n, and thus we assume our mechanisms are given access to an oracle which enables evaluating the function V. A *value oracle* receives a subset S and returns $V(S)$. Since we know value oracles cannot obtain a reasonable approximation [18, 7], we investigate whether stronger types of queries enable us to do so. A *demand oracle* receives a vector of prices $p_1, \ldots, p_n$ and returns a subset S s.t. $S \in \text{argmax}_{T\in[n]} V(T) - \sum_{i\in T} p_i$ is maximal. Let $\mathcal{V} = \{1, 2, 4, \ldots 2^{\log(V(N))}\}$. A value query can be simulated by a polynomial number of demand queries [3] and we therefore allow our algorithms to use value queries. We assume, without loss of generality, that 1 is the smallest non-zero value of V. We note that in this section we assume that V is non-decreasing.

We start by describing a procedure for finding a bundle of size t with value close to the bundle of size t with the highest value. We show that polynomially many demand queries suffice to achieve a 4-approximation, even if the function is subadditive[2].

[2]In fact, our algorithm can be slightly modified to achieve a $(2 + \epsilon)$-approximation, but the improved approximation algorithm does not suffice for constructing truthful budget feasible mechanisms.

A Procedure for Finding an Approximate Bundle of Size t

For each v in $\mathcal{V}$:
a. Find the bundle S that maximizes the demand when the price per item is $\frac{v}{2t}$
b. If $V(S) - |S| \cdot t < \frac{v}{2}$ then set $S_v = \emptyset$ and continue to the next v
c. Else, if $|S| > t$, let S_v be some bundle of size t such that $S_v \subseteq S$. Else, let $S_v = S$

Output: (v, S_v) for the maximal $v \in \mathcal{V}$ such that S_v is not empty

The procedure clearly uses a polynomial number of demand queries. Before proving some properties of the procedure, including its correctness, we prepare a lemma.

CLAIM 3.1. *Let $S \in \arg\max_T V(T) - p \cdot |T|$. Then for each $S' \subseteq S$ we have that $V(S') \geq p \cdot |S'|$.*

PROOF. From subadditivity we have that

$$V(S) - |T| \cdot p \leq V(S \setminus S') - |S \setminus S'| \cdot p + V(S') - |S'| \cdot p$$

which implies that $V(S') - |S'| \cdot p \geq 0$ as otherwise S does not maximize the demand. We therefore have that:

$$V(S') \geq |S'| \cdot p$$

□

LEMMA 3.2. *Let $S^* \in argmax_{|S|=t} V(S)$. The procedure finds a subset S_v such that $V(S_v) > \frac{V(S^*)}{4}$. Furthermore, v is either the maximal $v \in \mathcal{V}$ such that $v \leq V(S^*)$ or v is the minimal $v \in \mathcal{V}$ such that $v > V(S^*)$.*

PROOF. First, consider an iteration for which $v \leq V(S^*)$. Setting a price $p = v/2t$ for all items, the demand query oracle finds a subset S that maximizes the demand, i.e. $S \in \text{argmax}_T V(T) - |T| \cdot p$. In particular:

$$\begin{aligned} V(S) &- |S| \cdot p \geq V(S^*) - |S^*| \cdot p \\ &= V(S^*) - t \cdot \frac{v}{2t} = V(S^*) - \frac{v}{2} \\ &\geq \frac{V(S^*)}{2} \end{aligned}$$

Clearly, $V(S) - |S| \cdot p \geq \frac{V(S^*)}{2}$ implies $V(S) \geq \frac{V(S^*)}{2}$.

If $|S| \leq t$ then $S_v| = S$ and we are already done. If $|S| > t$, by Claim 3.1: $V(S_v) \geq p \cdot t = t \cdot \frac{v}{2t} = \frac{v}{2}$.

Notice that for the maximal $v \in \mathcal{V}$ where $v \leq V(S^*)$, we have that $v \geq \frac{V(S^*)}{2}$, thus $V(S_v) > \frac{V(S^*)}{4}$. To finish the proof, consider $v > 2V(S^*)$. Observe that $V(S) - |S| \cdot \frac{v}{2t}$, thus any bundle of size less than t will have a negative profit in Step (b) and the iteration will fail. Thus, assume towards contradiction that the iteration passes Step (b). By our discussion, we have that $|S| > t$. In this case, Claim 3.1 gives us that $V(S_v) \geq t \cdot \frac{v}{2t} = \frac{v}{2} > V(S^*)$. A contradiction. □

3.1 A Randomized Mechanism

We first use the above procedure to construct a randomized mechanism. In the next subsection we will derandomize this construction to obtain a deterministic mechanism. For this section, let $\alpha = 2 \cdot \lceil \log n \rceil$ and $i^* \in \arg\max_k V(\{k\})$.

A Randomized Budget Feasible Approximation Mechanism

For each t in $\mathcal{T} = \{1, 2 \ldots 2^{\lceil \log n \rceil}\}$ in decreasing order:

a. Let $\mathcal{N}'$ be the set of items with cost at most $B/(\alpha \cdot t)$ that are different than i^*

b. Using the procedure, find (v_t, S_t) among items in $\mathcal{N}'$

Output: Choose u.a.r between $\cup_t S_t$ and agent i^*

THEOREM 3.3. *The mechanism is universally truthful, budget feasible, and provides an expected approximation ratio of $O(\log^2 n)$.*

LEMMA 3.4. *The mechanism provides an approximation ratio of $O(\log^2 n)$.*

PROOF. Denote by S^* the set of agents participating in the optimal solution. Let $\overline{S^*} = \{i | i \in S^*, c_i > \frac{B}{\alpha}\}$, and $\underline{S^*} = S^* \setminus \overline{S^*}$.

Suppose that $V(\overline{S^*}) \geq \frac{V(S^*)}{2}$. Since the payment for each of the agents in $\overline{S^*}$ is at least $\frac{B}{\alpha}$, $|\overline{S^*}| \leq \alpha$. By subadditivity there must be one bidder $i \in \overline{S^*}$ such that $V\{i\} \geq \frac{V(\overline{S^*})}{\alpha}$. In particular we have that $V(\{i^*\}) \geq V(\{i\})$. Thus, if i^* is chosen this gives us an $O(\log n)$ approximation. Since i^* is chosen with probability $\frac{1}{2}$ the expected approximation guarantee is $O(\log n)$ in this case.

Assume now that $V(\underline{S^*}) \geq \frac{V(S^*)}{2}$ (this is the last case since by subadditivity $V(\overline{S^*})+V(\underline{S^*}) \geq V(S^*)$). If $V(\{i^*\}) \geq \frac{V(\underline{S^*})}{2}$, then similarly to before we have a constant approximation if i^* is chosen (which happens with probability $\frac{1}{2}$). Let $\underline{S^*}' = \underline{S^*} \setminus \{i^*\}$. We have that $V(\underline{S^*}') \geq \frac{V(S^*)}{4}$. Now put agents in $\underline{S^*}$ in bins according to their cost s.t. agent $i \in \underline{S^*}$ is in bin j if and only if $B/(\alpha \cdot 2^{j+1}) \leq c_i < B/(\alpha \cdot 2^j)$ where $j \in \mathcal{T}$.

Since there are $O(\log n)$ bins, subadditivity implies that there is a single bin k with value that is at least $O(\log n)$-fraction of $V(\underline{S^*})$. It follows that the optimal solution of size $\alpha \cdot 2^k$ over all items with cost at most $B/(\alpha \cdot 2^k)$ has value at least $O(V(\underline{S^*})/\log n)$. Since one of the iterations of the procedure gives us a set S_k of size 2^k that is a 4-approximation to the solution of size 2^k, and by subadditivity the solution of size 2^k is an $O(\alpha)$-approximation to the solution of size $O(\alpha \cdot 2^k)$, we have that $V(\cup_t S_t) \geq V(\underline{S^*})/4 \log^2 n$. Therefore, with probability $\frac{1}{2}$ we have an $O(\log^2 n)$ approximation in this case. □

LEMMA 3.5. *The mechanism is universally truthful.*

PROOF. To prove the lemma we will show that the mechanism is monotone: an agent that is selected and reduces his cost is still selected.

Now fix the random coin and suppose that agent i^* wins. Notice that he will remain the winner regardless of his cost, and in particular will remain the winner if he reduces his cost.

Assume now that the selected set is $\cup_t S_t$, and consider some agent $i \neq i^*$ that wins and reduces his cost from c_i to c'_i. Let $t \in \mathcal{T}$ be the maximal such that $i \in S_t$. Notice that i will be selected to S_t also if he reduces his cost (since the procedure is oblivious to the actual cost of the agent and takes into account only that the cost is smaller than some threshold). Therefore we have that the mechanism is monotone. □

LEMMA 3.6. *The mechanism is budget feasible.*

PROOF. Recall that the payment for an agent is the maximal cost that he can declare and still win[3] (the *threshold cost*). It is not hard to see that if i^* is chosen then his payment is B. For each other agent i that is chosen we claim that the threshold cost is at most $\frac{B}{\alpha t_i}$ where t_i is the maximal index such that $i \in S_{t_i}$. This implies that the mechanism is budget feasible: observe that for each $t \in \mathcal{T}$, at most t agents may receive a payment of $B/(\alpha \cdot t)$. Thus the total payment in that case is $\sum_{t \in \mathcal{T}} t \cdot B/(\alpha \cdot t) \leq |\mathcal{T}| \frac{B}{\alpha} \leq B$.

We now show that for each other agent i that is chosen the threshold cost is at most $\frac{B}{\alpha t_i}$. Suppose agent i has cost bigger than $\frac{B}{\alpha t_i}$. In this case i will not be selected when the size of the bundle considered is t_i or smaller, because his cost is too high. Also i will not be selected in iterations in which the bundle size is larger than t_i: in these iterations either the procedure runs on the exact same set $\mathcal{N}'$ of items (if his cost is still small enough), and for this set of items we know that i is not selected, or i has a cost that is too high and thus is not considered for selection at all. □

3.2 A Deterministic Mechanism

Our next goal is to construct a *deterministic* mechanism with a good approximation ratio. The randomized mechanism uses only one random coin, so a first natural attempt is to select the highest-value outcome of the two possible ones. Unfortunately, this does not work: when an agent reduces his cost he might increase $V(\cup_t S_t)$[4]. Therefore, we use a "monotone estimator" for the value of the union that has the property that if the cost of an agent decreases the value of the monotone estimator increases. We make sure that the value of the monotone estimator is "close" to $V(\cup_t S_t)$. Comparing the value of the monotone estimator to the value of single agent with the best value gives us a monotone $O(\log^3 n)$ mechanism. We note that, ignoring computational constraints, one can use the optimal algorithm as a monotone estimator an obtain a deterministic $O(\log^2 n)$ approximation. Let $\alpha = 2 \cdot \lceil \log n \rceil$ and $i^* \in \arg\max_k V(\{k\})$.

[3]The mechanism is universally truthful so we fix the random coin and prove the budget feasibility of the mechanism for every outcome of the random coin.

[4]This may happen since we do not have exact mechanisms to find the best bundle of size t but rather use an approximation mechanism to do so.

A Deterministic Budget Feasible Approximation Mechanism

For each t in $\mathcal{T} = \{1, 2 \ldots 2^{\lceil \log n \rceil}\}$ in decreasing order:

a. Let $\mathcal{N}'$ be the set of items with cost at most $B/(\alpha \cdot t)$ that are different than i^*

b. Using the procedure, find (v, S_t), $|S_t| = t$, among items in $\mathcal{N}'$

c. Using the procedure, find (v_t, S), $|S| = \alpha t$, among items in $\mathcal{N}''$

Output: If $\sum_t v_t \geq V(\{i^*\})$ then output $\cup_t S_t$. Else choose only agent i^*

Notice that the mechanism uses only a polynomial number of demand queries.

LEMMA 3.7. *The mechanism provides an approximation ratio of $O(\log^3 n)$.*

PROOF. Denote by S^* the set of agents participating in the optimal solution. Let $\overline{S^*} = \{i | i \in S^*, c_i > \frac{B}{\alpha}\}$, and $\underline{S^*} = S^* \setminus \overline{S^*}$. We start by showing two useful facts. The first one is:

$$2 \cdot \Sigma_t v_t \geq V(\underline{S^*}) \quad (1)$$

The first inequality follows since for every t, $2 \cdot v_t \geq V(S_t)$ and from the subadditivity of V. For each t, let S_t^* be the optimal solution of items of size at most $B/(\alpha \cdot t)$. The second inequality is:

$$\begin{aligned} \Sigma_t v_t &\leq 4\Sigma_t V(S_t^*) \\ &\leq 4 \log n \Sigma_t \cdot V(S_t) \\ &\leq 4 \log^2 n \cdot V(\cup_t S_t) \end{aligned} \quad (2)$$

The first inequality follows since v_t is a 4 approximation to S_t^*. The second one follows since V is subadditive and since in S_t^* there are at most a multiplicative factor of $\log n$ more items than in S_t. The third inequality is because t can take at most $\log n$ values.

We now proceed to the proof itself. Suppose that $V(\overline{S^*}) \geq \frac{V(S^*)}{2}$. The payment of each of the agents in $\overline{S^*}$ is at least $\frac{B}{\alpha}$, hence $|\overline{S^*}| \leq \alpha$. Therefore, by subadditivity there must be one bidder $i \in \overline{S^*}$ such that $V\{i\} \geq \frac{V(\overline{S^*})}{\alpha}$. In particular we have that $V(\{i^*\}) \geq V(\{i\})$. Thus, if i^* is chosen this gives us an $O(\log n)$ approximation. Else, we have that:

$$\begin{aligned} 4 \log^2 n \cdot V(\cup_t S_t) &\geq \Sigma_t v_t \\ &\geq V(\{i^*\}) \\ &\geq \frac{V(\overline{S^*})}{\alpha} \\ &\geq \frac{V(S^*)}{2 \log n} \end{aligned}$$

The first inequality follows from (2). This proves that in the case that i^* is not chosen we have an $O(\log^3 n)$ approximation.

Assume now that $V(\underline{S^*}) \geq \frac{V(S^*)}{2}$ (this is the last case since by subadditivity $V(\overline{S^*}) + V(\underline{S^*}) \geq V(S^*)$). If $V(\{i^*\}) \geq \frac{V(S^*)}{2}$, then similarly to the previous case we have a constant approximation if i^* is chosen. If $\cup_t S_t$ is selected then arguments very similar to (3) prove that the approximation ratio is $O(\log^2 n)$ in this case.

Therefore, let $\underline{S^*}' = \underline{S^*} \setminus \{i^*\}$ and assume that $V(\underline{S^*}') \geq \frac{V(S^*)}{4}$. Using (1) and (2) together we have that $8 \log^2 n \cdot V(\cup_t S_t) \geq V(\underline{S^*})$, i.e., an $O(\log^2 n)$ approximation. If i^* is selected then, using (1), $V(i^*) \geq \Sigma_t v_t \geq \frac{V(\underline{S^*})}{4}$, i.e., a constant approximation in this case. □

LEMMA 3.8. *The mechanism is truthful.*

PROOF. To prove the lemma we will show that the mechanism is monotone: an agent that is selected and reduces his cost is still selected. Suppose that agent i^* wins. If he reduces his cost he clearly remains the winner, since the expression $\sum_t v_t$ remains the same. Therefore, assume that some agent $i \neq i^*$ wins and reduces his cost from c_i to c_i'. We will show that the expression $\sum_t v_t$ does not decrease in this case, and monotonicity will follow. Furthermore, we will show that for every t the value of v_t cannot decrease.

To see this, fix some t. The only case where the output of the procedure might change is when $c_i > \frac{B}{t}$ but $c_i' \leq \frac{B}{t}$. Notice that the difference is that the procedure considers one more item i. In this case it might happen that the value of the set $V(S_t)$ will go down, but we will show that the value v_t cannot decrease. This follows from the observation that the procedure succeeds for a certain value of $v \in \mathcal{V}$ if it succeeds in step b: i.e., if there exists a profit maximizing bundle with large enough value. However, if such bundle exists, it will still exist when considering more items. □

LEMMA 3.9. *The mechanism is budget feasible.*

The proof of this lemma is almost identical to the proof of Lemma 3.6. In conclusion:

THEOREM 3.10. *The mechanism uses a polynomial number of demand queries, is truthful, individually rational, and budget feasible. Its approximation ratio is $O(\log^3 n)$.*

4. CUT FUNCTIONS

An interesting class of objective functions within the complement-free domain are functions for which $S \subseteq T$ does not necessarily imply that $V(S) \leq V(T)$, a property we refer to as nonincreasing utilities[5]. Naturally, we are interested in investigating whether budget feasible mechanisms can be obtained for these classes as well. Here we take a first step in this direction by examining *cut functions*: valuation functions where the value of a subset of agents can be represented as a cut on a graph. This class of functions is a representative of the class of nonincreasing valuation functions, which, as we now show, has constant factor approximation mechanisms that are budget feasible. The results in this section lead us to conjecture that there are broader classes of nonincreasing valuation functions with desirable guarantees in the budget feasibility model.

4.1 Mechanisms for Cut Functions

A cut function $V : 2^{[n]} \rightarrow \mathcal{R}_{\geq 0}$ is a valuation function for which there exists a graph $G = (\mathcal{N}, E)$ s.t. $V(S) = |C(S)|$,

[5]In optimization literature this property is referred to as nonmonotnicity. To avoid confusion in discussion of monotonicity of the allocation function, we use the nonincreasing utilities term.

where C is the cut induced by S, i.e. $C(S) = \{(u,v) \in E : u \in S, v \in \mathcal{N} \setminus S\}$. We note that maximizing this function is a variant of the classic computationally intractable MAX CUT problem. We denote the degree of a vertex $v_i \in \mathcal{N}$ by d_i (when it will be clear from the context we will use i instead of v_i) and for any $T \subseteq \mathcal{N}$, we use $E(T) = \{(u,v) \in E : u \in T\}$ to denote the set of edges that have at least one vertex in T. In our setting each vertex is held by a single strategic agent with a private cost and our objective is to maximize $V(S) = |C(S)|$ under the budget constraint.

We first show a randomized mechanism that is budget feasible and obtains a constant factor approximation. We will then discuss its derandomization which also guarantees a constant factor approximation.

4.2 A Randomized Mechanism

THEOREM 4.1. *For cut functions there is a randomized $O(1)$-approximation mechanism that is universally truthful and budget feasible.*

In proof, consider the following mechanism:

A Randomized Budget Feasible Approx. Mechanism for Cut Functions

1. Set $\mathcal{N}' = \{i \in \mathcal{N} : c_i \le B/2\}$, $i^* \in \text{argmax}_{i \in \mathcal{N}'} d_i$, $S = \{i^*\}$ and $i \in \text{argmax}_{i \in \mathcal{N}' \setminus \{i^*\}} \frac{d_i}{c_i}$
2. While $c_i \le \frac{B}{24} \cdot \left(\frac{|C(S \cup \{v_i\})| - |C(S)|}{C(S \cup \{v_i\})|}\right)$:
   a. Add i to S
   b. Set $\bar{\mathcal{N}}$ to be all agents in $j \in \mathcal{N}' \setminus \{i^*\}$ for which $|C(S \cup \{j\})| - |C(S)| \ge \frac{2}{3} d_j$
   c. Set $i \in \text{argmax}_{j \in \bar{\mathcal{N}} \setminus \{i^*\}} \frac{|C(S \cup \{j\})| - |C(S)|}{c_j}$

Output: Choose u.a.r between S and $\text{argmax}_{i \in \mathcal{N} \setminus \mathcal{N}'} d_i$

It is easy to verify that the above mechanism is monotone in the agents' costs and thus truthful. We first prove its approximation guarantee before showing that it is indeed budget feasible. In the following proofs we will use S_i to denote the subset of agents selected by the mechanism after the ith stage.

LEMMA 4.2. *At each stage j we have $|C(S_j)| \ge \frac{1}{3}|E(S_j)|$.*

PROOF. We will show by induction on the stage of the mechanism that $|C(S_j)| \ge \frac{1}{3}\sum_{i \in S_j} d_i$ which suffices to prove the lemma since $\sum_{i \in S_j} d_i$ is an upper bound on $|E(S_j)|$.

In the first stage of the mechanism the inequality trivially holds. For a general step j, the vertex v_j that is selected must respect the condition $|C(S_{j-1} \cup \{v_j\})| - |C(S_{j-1})| \ge \frac{2}{3}d_j$. This condition implies that when adding v_j there are at most $\frac{1}{3}d_j$ edges between v_j and vertices in S_{j-1} and thus by adding v_j to S_{j-1} at most $\frac{1}{3}d_j$ edges will be removed from the cut and $\frac{2}{3}d_j$ edges will be added. Thus, together with the inductive hypothesis we have that:

$$\begin{aligned}|C(S_j)| \ge& \Big(|C_{j-1}| - \frac{1}{3}d_j\Big) + \frac{2}{3}d_j \\ \ge& \Big(\frac{1}{3}\sum_{i \in S_{j-1}} d_i - \frac{1}{3}d_j\Big) + \frac{2}{3}d_j \\ =& \frac{1}{3}\sum_{i \in S_j} d_i\end{aligned}$$

□

LEMMA 4.3. *Let S^* be the optimal solution over agents in $\mathcal{N}'$ with budget $B' = \frac{B}{24}$, then $|C(S)| \ge \frac{|C(S^*)|}{6}$.*

PROOF. Partition the set of edges in $C(S^*)$ to the following disjoint subsets of edges: $S_1^* = \{(u,v) \in C(S^*) : u, v \in S\}, S_2^* = \{(u,v) \in C(S^*) : u \in S, v \notin S\}, S_3^* = \{(u,v) \in C(S^*) : u, v \notin S\}$. First, as implied by Lemma 4.2, we have that $|E(S) \setminus C(S)| \le 2|C(S)|$ and thus:

$$|C(S_1^*)| \le |E(S) \setminus C(S)| \le 2|C(S)| \tag{3}$$

In the case of S_2^*, since each vertex has an endpoint in S, it must be that $|C(S_2^*)| \le |C(S)|$, and thus:

$$|C(S_2^*)| \le |C(S)| \tag{4}$$

In the case where $S_3^* = \emptyset$ the above inequalities suffice to prove our lemma. Otherwise, to bound the ratio between $|C(S_3^*)|$ and $|C(S)|$, assume $S_3^* \ne \emptyset$ and w.l.o.g assume its vertices are labeled s.t. vertex v_i has the greatest ratio between marginal contribution to the cut and cost, given the cut induced by vertices $v_1, \ldots, v_{i-1}$. In such an ordering, for all $i < r = |S_3^*|$ we get:

$$\frac{|C(T_i)| - |C(T_{i-1})|}{c_i} \ge \frac{|C(T_{i+1})| - |C(T_i)|}{c_{i+1}}$$

where T_i is the subset that includes the first i vertices taken according to the ordering and $T_0 = \emptyset$.

Let v_k be the first vertex not selected by our mechanism to be in S. In this case we have that:

$$c_k > B' \cdot \Big(\frac{|C(S \cup \{v_k\})| - |C(S)|}{|C(S \cup \{v_k\})|}\Big) \tag{5}$$

Since we assume $S_3^* \ne \emptyset$ and $S \cap S_3^* = \emptyset$, all vertices in S_3^* respect the condition of having at least $2/3$ of their edges not connected to vertices in S, and thus such a vertex must exist. Also, since $S_3^* \cap S = \emptyset$, it follows that v_k is either in S_3^*, or that every vertex in S_3^* has smaller marginal contribution ratio to cost than that of v_k. Thus, in either case for any $i \in [r]$ we have that:

$$\frac{c_k}{(|C(S \cup \{v_k\})| - |C(S)|)} \le \frac{c_i}{d_i} \le \frac{c_i}{(|C(T_i)| - |C(T_{i-1})|)} \tag{6}$$

where the first inequality is due to the fact that the vertices in S_3^* do not have endpoints in S and thus their marginal contribution equals their degree, and the second inequality is due to the decreasing marginal utilities property of the cut function.

Since S_3^* is feasible we have that $\sum_{i=1}^{r} c_i \le B'$, which we can write as:

$$\sum_{i=1}^{r}\Big(\frac{c_i}{|C(T_i)|-|C(T_{i-1})|}\Big)\cdot\Big(|C(T_i)|-|C(T_{i-1})|\Big)\le B'$$

Since $S_3^* = T_r$ we have that $|C(S_3^*)| = \sum_{i=1}^{r}\Big(|C(T_i)| - |C(T_{i-1})|\Big)$ and therefore together with (6) above we have:

$$c_k\cdot\Big(\frac{|C(S_3^*)|}{|C(S\cup\{v_k\})|-|C(S)|}\Big)\le B'$$

The above inequality, together with (5) implies that $|C(S\cup\{v_k\})| > |C(S_3^*)|$. Since S includes the vertex with largest degree it follows that $|C(S)| \ge d_i$, for any $i \in \mathcal{N}'$, and thus:

$$2|C(S)| \ge |C(S)| + d_k \ge |C(S)\cup\{v_k\}| \ge C(S_3^*) \qquad (7)$$

To conclude, let $\alpha_i = \frac{|C(S_i^*)|}{|C(S^*)|}$ for $i \in \{1,2,3\}$. Since the sets are disjoint $\alpha_1+\alpha_2+\alpha_3 = 1$. Since there must exist an $\alpha_i \ge 1/3$, from (3),(4),and (7) it follows that $|C(S)| \ge \frac{|C(S^*)|}{6}$. □

The above lemma implies the desired constant factor approximation ratio. Let $OPT(c,\mathcal{N},B)$ denote the value of the optimal solution over the set of agents $\mathcal{N}$ with bid profile c and budget B. First, observe that since V is subadditive, for any natural $\alpha > 1$ when the agent i^* with largest value is selected is in $OPT(c,\mathcal{N}',B/\alpha)$, then we have that: $\alpha\cdot OPT(c,\mathcal{N},B/\alpha) + (\alpha-1)V(\{i^*\}) \ge OPT(c,\mathcal{N},B)$, and therefore $(2\alpha-1)OPT(c,\mathcal{N},B/\alpha) \ge OPT(c,\mathcal{N},B)$.

Since the vertex with largest degree in $\mathcal{N}'$ is included in S, we have that $47\cdot OPT(c,\mathcal{N}',B/24) \ge OPT(c,\mathcal{N}',B)$ and thus by Lemma 4.3 we have that $282|C(S)| \ge OPT(c,\mathcal{N}',B)$. Since we selected the optimal solution in $\mathcal{N}\setminus\mathcal{N}'$ with probability 1/2 and $OPT(c,\mathcal{N},B) \le OPT(c,\mathcal{N}',B)+OPT(c,\mathcal{N}\setminus\mathcal{N}',B)$ the mechanism is a 564-approximation.

We will complete the proof of our theorem by showing the mechanism is indeed budget feasible. Unlike the approach taken in [18] where a characterization of payments was shown, we will prove budget feasibility by directly showing a bound on threshold payments.

LEMMA 4.4. *The mechanism is budget feasible.*

PROOF. It's easy to see that the threshold payment for i^* with largest degree in $\mathcal{N}'$ is $\frac{B}{2}$, as it is always selected by the mechanism. To bound the threshold payment of the other agents in S by the remaining budget $\frac{B}{2}$, for a given bidding profile $c = (c_1 \dots c_n)$, let v_j be a selected vertex with bid c_j, and let $c_j' > c_j$ be the maximum bid that v_j can declare and remain selected when all other agents declare the same cost, and let $c' = (c_1,\dots,c_{j-1},c_j',c_{j+1}\dots,c_n)$. Let S' and S_i' denote the set of selected agents by the mechanism and agents selected at stage i, respectively, when the bid profile is c'.

First, observe that:

$$\begin{aligned}2\cdot OPT(c',\mathcal{N}',B') &\ge OPT(c,\mathcal{N}'\setminus\{j\},B') + d_j\\ &\ge OPT(c,\mathcal{N}',B')\\ &\ge |C(S)|\end{aligned}$$

From Lemma 4.3 it follows that $|C(S')| \ge \frac{OPT(c',\mathcal{N}',B')}{6}$ and thus we have that $|C(S')| \ge \frac{|C(S)|}{12}$.

W.l.o.g. assume v_j is selected at stage j when the bid profile is c. Notice that when running the mechanism with bid profile c', the same first $j-1$ agents are as when running the the mechanism with the profile c, and we have that $|C(S_{j-1})\cup\{v_j\}| - |C(S_{j-1})| \ge |C(S_{r-1}')\cup\{v_j\}| - |C(S_{r-1}')|$ where r is the stage in which v_j is selected when bidding c_j'. This implies:

$$\frac{c_j'}{(|C(S_j)|-C(S_{j-1})|)} \le \frac{c_j'}{(|C(S_r\cup\{v_j\})|-C(S_r)|)} \le \frac{B'}{|C(S')|}$$

.

where the second inequality is due to the fact that every agent $i \in S'$ respects the condition $c_i \le B'\cdot\Big(\frac{|C(S_i')|-|C(S_{i-1}')|}{|C(S')|}\Big)$. The above inequality implies:

$$\begin{aligned}c_j' &\le B'\Big(\frac{|C(S_{j-1}\cup\{v_j\})|-|C(S_{j-1})|}{|C(S')|}\Big)\\ &\le 12\cdot B'\Big(\frac{|C(S_{j-1}\cup\{v_j\})|-|C(S_{j-1})|}{|C(S)|}\Big)\end{aligned}$$

Since c_j' is the maximum bid an agent can declare, it follows that the threshold payments for any agent j are bounded from above by: $12\cdot B'\Big(\frac{|C(S_j|-|C(S_{j-1})|}{|C(S)|}\Big)$. Since $\sum_{j\in S'}(|C(S_j|-|C(S_{j-1})|) = |C(S)|$ and $B' = B/24$, the total payments to agents in $S\setminus\{i^*\}$ are bounded by $B/2$ which implies budget feasibility. □

4.3 A Deterministic Mechanism for Cut Functions

The approximation guarantee provided by our mechanism depends on randomizing between the subset S selected in steps (1) and (2) of the above mechanism and the vertex with largest degree in $\mathcal{N}\setminus\mathcal{N}'$. In order to derandomize the mechanism and provide a bounded approximation guarantee we need to select between the two solutions, based on the value of the cuts they produce. The problem with using a direct comparison between the values of two solutions is that it breaks monotonicity: when lowering her cost, an agent that is selected to S may change the order of the vertices that are selected and decrease the size of the cut so that it is smaller than the cut induced by the vertex with largest degree in $\mathcal{N}\setminus\mathcal{N}'$. If the mechanism would select based on a direct comparison between the two solutions an agent could loose her allocation by reducing her cost. We give a concrete example of such a case below.

EXAMPLE 1. *A direct comparison breaks monotonicity.*

PROOF. Consider a graph with the disjoint sets of vertices: $\{v_1,v_2,v_3,v_4\},\mathcal{N}_1,\mathcal{N}_2,\mathcal{N}_3,\mathcal{N}_4$, where $|\mathcal{N}_1| = n+2$, $|\mathcal{N}_2| = n, |\mathcal{N}_3| = n-1, |\mathcal{N}_4| = 3n$, v_2 is connected v_3 and each vertex $v \in \mathcal{N}_i$ is only connected to v_i, for all $i \le 4$.

We would like to show that there is a cost profile s.t. the mechanism allocates to $\{v_1,v_2,v_3\}$, but as v_2 slightly decreases her cost declaration, v_3 is no longer allocated and since $d_4 = 3n \ge |C(\{v_1,v_2\})| = 2n+2$, a direct comparison between will result in v_2 not being selected, and thus breaking monotonicity.

For a budget $B = 2$ and costs $c_1 = \epsilon, c_2 = (\frac{n}{n+1})c_3+\epsilon, c_3 = \frac{n-1}{3n+2}+\epsilon, c_4 = B$, using a small ϵ and sufficiently large n (say $\epsilon = 2^{-n}, n > 100$) serves as such an example. In this instance, the mechanism runs procedure only on $\{v_1,v_2,v_3\}$

and compares its solution with d_4. Under these costs one can verify that the procedure selects $\{v_1, v_2, v_3\}$ which produce a cut with $3n+1 > d_4$. If v_2 reduces her cost to ϵ however, then v_1, v_2 are selected, v_3 is not selected as her marginal contribution now dropped, and since $C(\{v_1, v_2\}) = 2n+2 < d_4$, v_4 will be allocated instead. □

It is important to emphasize that such examples are not unique to our specific mechanism and not even to cut functions. This type of problem arises when applying greedy procedures in other setting like coverage, submodular, and subadditive valuation functions.

4.3.1 Derandomization via Relaxation

To derandomize the mechanism in a manner that preserves monotonicity and provides a constant factor approximation guarantee we suggest the following approach, which is inspired by [2]: rather than a direct comparison between the two solutions, we will compute a linear programming relaxation over $\mathcal{N}'$, compare between the value returned by this solution and the largest degree in $\mathcal{N} \setminus \mathcal{N}'$, and select S if and only if the solution returned by the relaxation is greater.

Since the solution returned by the relaxation will be an optimal *fractional* solution, such a scheme guarantees monotonicity: an agent in $\mathcal{N}'$ that reduces her cost can only increase the value of the optimal fractional solution, thus avoiding the problem discussed in the above example. As long as we can guarantee that the fractional solution returned by the relaxation is a constant factor away from the optimal integral solution over $\mathcal{N}'$, implementing such a scheme will guarantee a constant factor approximation.

More concretely, the optimization problem can be reformulated as the following integer program:

$$\max \sum_{i<j} z_{ij} \tag{8}$$

$$s.t.\ z_{ij} \le x_i + x_j, \qquad i<j, \tag{9}$$

$$z_{ij} \le 2 - x_i - x_j \qquad i<j, \tag{10}$$

$$\sum_{i\in\mathcal{N}} x_i c_i \le B, \tag{11}$$

$$x_i, z_{ij} \in \{0,1\}, i \in \mathcal{N}, i<j \tag{12}$$

where x_i, c_i are variables representing the vertices and their costs, respectively and z_{ij} represent the edges. As discussed above, we would like to compute a fractional solution in polynomial time that will be a constant factor away from the optimal integral solution of the above program. We will do this by showing that the linear program relaxation has a constant integrality gap. To bound the integrality gap of the LP relaxation of the problem, we assume that $c_i \le \frac{B}{2}$, for every i.

Theorem 4.5. *The LP has a constant integrality gap.*

Proof. To prove the theorem we consider the following randomized rounding algorithm:

Randomized Rounding for Cut Functions
1. Add each vertex v_i to the cut S with probability $\frac{x_i}{4}$ 2. If $\sum_{i\in S} x_i \cdot c_i > B$ then set $S' = \emptyset$ else $S' = S$ **Output:** S'

We will show that $\Pr[(i,j) \in S'] \ge \frac{7z_{ij}}{64}$. Thus, if we let T_{ij} be the indicator variable that gets a value of 1 when $(i,j) \in S'$ and 0 otherwise, we have that $\mathbb{E}[T_{i,j}] \ge \frac{7z_{ij}}{64}$. By linearity of expectation, $\mathbb{E}[C(S')] = \sum_{(i,j)\in E} \mathbb{E}[T_{ij}] \ge \sum_{(i,j)\in E} \frac{7z_{ij}}{64}$. Therefore there must always be an integral solution that has a value of at least $\frac{7}{64}$ of the value of the optimal fractional solution.

We first calculate the probability that $(i,j) \in S$:

$$\begin{aligned}\Pr[(i,j)\in S] &= (1-\frac{x_i}{4})\frac{x_j}{4} + (1-\frac{x_j}{4})\frac{x_i}{4} \\ &= \frac{x_i}{4} + \frac{x_j}{4} - 2\cdot\frac{x_i}{4}\cdot\frac{x_j}{4} \\ &\ge \frac{z_{ij}}{4} - \frac{x_i \cdot x_j}{8} \\ &\ge \frac{z_{ij}}{4} - \frac{(\frac{z_{ij}}{2})^2}{8} \\ &\ge \frac{7z_{ij}}{32}\end{aligned}$$

where the first equality is by the properties of the randomized rounding, the first inequality is by the LP constraints, and the second inequality by basic analysis and using $z_{ij} \le x_i + x_j$. The last inequality uses the fact that $z_{ij} \in [0,1]$ and thus $z_{ij} > z_{ij}^2$.

Next we calculate $\Pr[S' = \emptyset | (i,j) \in S]$. If $(i,j) \in S$ this implies that exactly one of the vertices i and j is in S. Assume without loss of generality that $i \in S$. Now $S' = \emptyset$ only if the total budget exceeds B. Observe that $\mathbb{E}[\sum_{i'\in S, i'\ne j, i'\ne i} c_{i'} | i \in S, j \notin S] \le \frac{B}{2}$ since each i' is selected into S with probability exactly $\frac{x_{i'}}{4}$ and that $\sum_{i\in\mathcal{N}} x_i c_i \le B$ by the LP constraints. By Markov's inequality:

$$\Pr[\sum_{i'\in S, i'\ne j, i'\ne i} c_{i'} \ge \frac{B}{2} | i \in S, j \notin S] \le \frac{1}{2}$$

Taking into account that $c_i \le \frac{B}{2}$, by our assumption, we can now bound the probability that the budget used by S does not exceed B:

$$\Pr[S' \ne \emptyset | i \in S, j \notin S]$$

$$\ge \Pr[\sum_{i'\in S, i'\ne j, i'\ne i} c_{i'} \le \frac{B}{2} | i \in S, j \notin S] \ge \frac{1}{2}$$

To conclude, $\Pr[(i,j) \in S] \ge \frac{7z_{ij}}{32}$ and also $\Pr[(i,j) \in S' | (i,j) \in S] \ge \frac{1}{2}$. Thus $\Pr[(i,j) \in S'] \ge \frac{7z_{ij}}{64}$, for every $(i,j) \in E$, as needed by the discussion above. □

We can now formally state the deterministic mechanism. We use A to denote the allocation rule in steps (1) and (2) of the randomized mechanism, f_{LP} to be the optimal fractional solution, and $LP(x)$ to be the value of the LP evaluated on x.

A Deterministic Mechanism for Cut Functions
1. Let $\mathcal{N}' = \{i :\in \mathcal{N} : c_i \le B/2\}$, $i' = \text{argmax}_{i\in N\setminus\mathcal{N}'} d_i$ 2. Compute $S = A(\mathcal{N}')$ and $x^* = f_{LP}(\mathcal{N}')$ **Output:** if $LP(x^*) \ge d_{i'}$ return S o.w. return $\{i'\}$

Theorem 4.6. *There is a $O(1)$-approximation polynomial time mechanism for cut functions which is truthful and budget feasible.*

PROOF. Truthfulness and budget feasibility follow from the arguments in the case of the randomized mechanism. For the approximation guarantee, note that $OPT(c, \mathcal{N}, B) \leq OPT(c, \mathcal{N}', B) + OPT(c, \mathcal{N} \setminus \mathcal{N}', B)$. Since i' is the optimal solution in $\mathcal{N} \setminus \mathcal{N}'$ if its value is larger than $LP(x^*)$ it must be larger than the optimal integral solution as well, and thus choosing i' guarantees a 2-approximation in this case. Otherwise we have:

$$|C(S)| \geq \frac{OPT(c, \mathcal{N}', B)}{282} \geq \frac{7 \cdot L(x^*)}{64 \cdot 282} \geq \frac{|C(\{a\})|}{2579}$$

Therefore in this case, we are guaranteed that $|C(S)| \geq \frac{OPT(c, \mathcal{N}, B)}{5158}$. □

Acknowledgments

We thank Ashwinkumar B. V. for several helpful comments.

5. REFERENCES

[1] Aaron Archer and Éva Tardos. Frugal path mechanisms. In *SODA*, pages 991–999, 2002.

[2] Yossi Azar and Iftah Gamzu. Truthful unification framework for packing integer programs with choices. In *ICALP (1)*, pages 833–844, 2008.

[3] Liad Blumrosen and Noam Nisan. Combinatorial auctions. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[4] Matthew Cary, Abraham D. Flaxman, Jason D. Hartline, and Anna R. Karlin. Auctions for structured procurement. In *SODA*, pages 304–313, 2008.

[5] Ning Chen, Edith Elkind, Nick Gravin, and Fedor Petrov. Frugal mechanism design via spectral techniques. In *FOCS*, pages 755–764, 2010.

[6] Ning Chen, Nick Gravin, and Pinyan Lu. On the approximability of budget feasible mechanisms. 2011.

[7] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, 2010.

[8] Edith Elkind, Leslie Ann Goldberg, and Paul W. Goldberg. Frugality ratios and improved truthful mechanisms for vertex cover. In *ACM Conference on Electronic Commerce*, pages 336–345, 2007.

[9] Edith Elkind, Amit Sahai, and Kenneth Steiglitz. Frugality in path auctions. In *SODA*, pages 701–709, 2004.

[10] Ludwig Ensthaler and Thomas Giebe. A dynamic auction for multi-object procurement under a hard budget constraint. Discussion Papers 310, 2010.

[11] Joan Feigenbaum, Christos H. Papadimitriou, Rahul Sami, and Scott Shenker. A bgp-based mechanism for lowest-cost routing. In *PODC*, pages 173–182, 2002.

[12] Joan Feigenbaum, Rahul Sami, and Scott Shenker. Mechanism design for policy routing. In *PODC*, pages 11–20, 2004.

[13] Arpita Ghosh and Aaron Roth. Selling privacy at auction. *ARXIV*, abs/1011.1375, 2010.

[14] Anna R. Karlin, David Kempe, and Tami Tamir. Beyond VCG: Frugality of truthful mechanisms. In *FOCS*, pages 615–626, 2005.

[15] David Kempe, Mahyar Salek, and Cristopher Moore. Frugal and truthful auctions for vertex covers, flows and cuts. In *FOCS*, pages 745–754, 2010.

[16] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1), 1981.

[17] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001. A preliminary version appeared in STOC 1999.

[18] Yaron Singer. Budget feasible mechanisms. In *FOCS*, pages 765–774, 2010.

[19] Kunal Talwar. The price of truth: Frugality in truthful mechanisms. In *STACS*, pages 608–619, 2003.