

Inapproximability of Combinatorial Public Projects

Michael Schapira¹ and Yaron Singer²

¹ The School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel, mikesch@cs.huji.ac.il

² Computer Science Division, UC Berkeley, USA, yaron@cs.berkeley.edu

Abstract. We study the Combinatorial Public Project Problem (CPPP) in which n agents are assigned a subset of m resources of size k so as to maximize the social welfare. Combinatorial public projects are an abstraction of many resource-assignment problems (Internet-related network design, elections, etc.). It is known that if all agents have submodular valuations then a constant approximation is achievable in polynomial time. However, submodularity is a strong assumption that does not always hold in practice. We show that (unlike similar problems such as combinatorial auctions) even slight relaxations of the submodularity assumption result in non-constant lower bounds for approximation.

1 Introduction

There are various real-world settings in which a set of resources is chosen to collectively serve an entire community: In elections, for instance, candidates are chosen to serve a community of voters. States choose which roads to build for the benefit of their residents. Another interesting example is that of choosing overlay networks in the Internet [3, 2] (for instance, in the context of inter-domain routing): A node in the network chooses the optimal subset of nodes to route traffic through, in a manner which is most beneficial to its clients³. Such an overlay - the subset of chosen nodes through which traffic is routed - is beneficial to different clients in different degrees. This creates the difficulty of maximizing the social welfare in this setting (as in the settings above).

The Combinatorial Public Project Problem (CPPP), recently presented and studied in [18], is an abstraction of such settings. In CPPP there are m resources and n agents, each with a valuation function defined over all subsets of resources, and the objective is to choose the k resources which maximize the social welfare of the agents. It is easy to see that CPPP is NP hard. However, in the case where each valuation function is submodular CPPP allows for good approximations [15, 18] (within $1 - \frac{1}{e}$).⁴ A valuation function v is said to be *submodular* if for every two subsets of resources S, T we have $v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$.

³ A formal description of the overlay networks setting is presented in Section 2.

⁴ This bound is tight, see [8].

In this paper we study CPPP, focusing on computational and communication complexity lower bounds which result from relaxations of the above submodularity property. Indeed, out of the myriads of problems which can be depicted as combinatorial public projects, in many instances the submodularity constraint does not apply. We illustrate this point in Section 2, where we show how the difference between submodular and general valuations separates between two seemingly close Internet-related network design problems (overlay networks).

Our main result is showing that even small relaxations of the submodularity property in combinatorial public projects result in strong inapproximability bounds:

Theorem: Obtaining an approximation ratio of $m^{\frac{1}{4}-\epsilon}$ for CPPP with subadditive valuation functions requires exponential communication in m (for every constant $\epsilon > 0$).

A valuation function v is subadditive if for every two subsets of resources S, T we have that $v(S \cup T) \leq v(S) + v(T)$. In fact, we prove our lower bound for a more restricted class of valuations called fractionally-subadditive [8] (introduced in [17] and termed “XOS” there). This lower bound establishes that in terms of approximability there is a huge gap between the submodular and subadditive case. This is in stark contrast to what is known about the related combinatorial auctions problem, for which a 2-approximation is achievable for subadditive valuations [8]. We show that our lower bound is nearly tight by presenting a simple $O(m^{\frac{1}{2}})$ approximation algorithm for subadditive valuations. We leave the question of closing the gap between $m^{\frac{1}{4}}$ and $m^{\frac{1}{2}}$ open. Also, a big open question is showing that a similar lower bound is achievable in the computational-complexity model.

In [18] it is shown, for the case of CPPP with submodular valuations, that while a constant approximation ratio is possible, no such ratio is achievable via *truthful* algorithms. The $O(m^{\frac{1}{2}})$ approximation algorithm for subadditive valuations presented in this paper is truthful. Hence, our results show that the hardness of CPPP with subadditive valuations is due to computational reasons and not to the truthfulness requirement. That is, truthful approximation algorithms perform just, or nearly, as well as unrestricted algorithms.

Finally, we study the approximability of CPPP with general valuations. We prove two complementary lower bounds, one in the computational-complexity model and one in the communication-complexity model:

Theorem: Obtaining an approximation of $O(n^{\frac{1}{2}-\epsilon})$ in CPPP with general valuations, for any $\epsilon > 0$, is impossible unless $P = NP$. Obtaining an approximation of $O(n^{1-\epsilon})$ is impossible unless $P = ZPP$.

Theorem: Obtaining an approximation of $(1 - \epsilon)n$ for CPPP with general valuations requires exponential communication in m (for any $\epsilon > 0$ and for any $n \ll 2^m$).

2 Model and Motivation

In this section we formally present CPPP and our model. We motivate CPPP by presenting an Internet-related network design setting which is naturally formulated as a combinatorial public project. We use this example to illustrate the importance of considering relaxations of submodularity.

2.1 The Model

In CPPP there is a set of n agents $N = \{1, \dots, n\}$, a set of m resources $M = \{1, \dots, m\}$, and a parameter $1 \leq k \leq m$. Each agent has a valuation function (sometimes simply referred to as a *valuation*) $v_i : 2^{[m]} \rightarrow \mathcal{R}_{\geq 0}$. We make two standard assumptions regarding each valuation function v_i : That $v_i(\emptyset) = 0$ (normalized) and that for all $S \subseteq T \subseteq [m]$ it holds that $v_i(S) \leq v_i(T)$ (non-decreasing). The objective in CPPP is to find a subset of resources S^* of size k which maximizes the social-welfare. That is, we wish to find $S^* \in \mathit{argmax}_{S \subseteq [m], |S|=k} \sum_i v_i(S)$.

We require algorithms for CPPP to run in time that is polynomial in the natural parameters of the problem – n and m . In some cases agents’ valuations can be concisely represented (encoded in space polynomial in m and n). In these cases we wish to explore the *computational complexity* of the problem.

However, in general, the size of the “input” (the valuations) can be exponential in m , and so we must specify how it can be accessed. We take a “black box” (concrete complexity) approach (see [5]): Every valuation function is assumed to be represented by an oracle which can answer a certain type of queries, and we restrict our algorithms to ask polynomially (in both m and n) many such queries. We consider two standard types of queries:

- **Value queries:** A value query to a valuation v_i is in the form of a subset of resources $S \subseteq [m]$. The answer is simply $v_i(S)$. This is a natural type of query to use when designing algorithms for CPPP.
- **General queries:** A general query allows any type of questions (even computationally intractable ones), as long as each question is addressed to a single v_i . We only require that the “size” of the query not be too “large”. This is equivalent to Yao’s communication model [20] in which the different parties are computationally omnipotent and we measure the number of bits they must exchange to compute a given function (see an introduction to communication complexity in [12]). The immense strength of general queries is useful for showing impossibility results.

2.2 Overlay Networks

Consider the following Internet-inspired setting [3, 2] we call the *overlay network problem*: We are given a network graph $G = (V, E)$, where $V = S \cup M \cup D$ (S, M, D are disjoint). We shall refer to the nodes in S as *source-nodes*, to the

nodes in M as *potential overlay nodes*, and to the nodes in D as *destination-nodes*. Assume that there is some metric function m that assigns a non-negative real value m_e to every $e \in E$ (assume that m 's value for a non-existent edge is ∞).

Each node $i \in S$ has a valuation function v_i that assigns a non-negative real value to every subset of potential overlays $M' \subseteq M$. Intuitively, for every $M' \subseteq M$, $v_i(M')$ specifies i 's desire to route through M' to the nodes in D . Each v_i is consistent with the metric m in the following sense: For every $M', M'' \subseteq M$, $v_i(M') \geq v_i(M'')$ iff the sum of the lengths of the shortest routes from i to all nodes in D that only go through nodes in M' is at most that obtained from routing only through nodes in M'' .⁵

The goal is to choose an overlay network from the potential overlay nodes. For a given parameter $1 \leq k \leq |M|$, we want to find a subset of M of size k that maximizes the social welfare of the source nodes. That is, we wish to find a set $M^* = \operatorname{argmax}_{M' \subseteq M, |M'|=k} \sum_{i \in S} v_i(M')$. It is easy to see that the overlay network problem is a special case of CPPP in which the agents are the source nodes, the resources are the potential overlays and the agents' valuations are induced by the network graph and the metric function⁶. We distinguish between two versions of this problem, which illustrate the big differences between instances for which submodularity does, and does not, apply:

- **The submodular case:** Observe that if no two nodes in M are directly connected via an edge in E then the valuation functions of the source nodes in the overlay network problem are submodular. This is due to the fact that they exhibit the decreasing marginal values property that is known to be equivalent to submodularity: For every node $i \in S$, for every $M' \subseteq M'' \subseteq M$, and for every $a \in M$, it holds that $v_i(M' \cup \{a\}) - v_i(M') \geq v_i(M'' \cup \{a\}) - v_i(M'')$.
- **The non-submodular case:** In the more general case, in which nodes in M can be connected to one another, the valuation functions are no longer guaranteed to be submodular. For example, think of a network graph $G = (V, E)$ with four nodes: a single source node i , a single destination node d , and two potential overlay nodes a, b . Assume that $E = \{(i, a), (a, b), (b, d)\}$ and that m assigns a value of 0 to every edge in E . Let $v_i(\{a\}) = v_i(\{b\}) = 0$ and $v_i(\{a, b\}) = 1$. Observe that v_i is indeed consistent with m , but is neither submodular nor subadditive.

In this model the difference between general and submodular valuations originates from whether or not the nodes in M are connected to each other. These

⁵ Formally, for every node $i \in S$, and for every subset $M' \subseteq M$ we define $c_i(M')$ to be the sum of the lengths of the shortest paths from i to all nodes in D that only go through nodes in M' . That is, if we define $(\forall i \in S, d \in D)$ $c_{id}(M')$ to be the length of the shortest route (given m) that has i as its first node, d as its last node, and all intermediate nodes in M' , then $c_i(M') = \sum_{d \in D} c_{id} M'$. $v_i(M') \geq v_i(M'')$ iff $c_i(M') \leq c_i(M'')$.

⁶ The artificial division of V into S , M , and D is only required for ease of exposition.

seemingly subtle and insignificant differences lead to violations of submodularity which motivate our study of CPPP with valuations which are not submodular.

3 Subadditive Valuations

In this section we study CPPP with subadditive valuations. We prove our main result, which is showing a lower bound of $m^{\frac{1}{4}}$ for a class strictly contained in subadditive valuations. This result shows that relaxations of submodularity that may seem small at first glance, and are not too costly for other problems (e.g., combinatorial auctions [7, 8]), lead to unreasonable approximation ratios for CPPP.

We show that this lower bound is nearly tight by presenting a simple \sqrt{m} approximation algorithm. This algorithm also has the advantage of being *truthful*. Hence, we show that for CPPP with subadditive valuations the gap between truthful and unrestricted algorithms in terms of approximation ratio is insignificant (this is in contrast to the submodular case in which the gap is huge [18]).

3.1 Lower Bound for Subadditive Valuations

We prove our result for fractionally-subadditive valuations [8] (defined in [17] and termed “XOS” there). This result is achieved in the communication complexity (general queries) model. Thus, we show that even if agents are computationally unlimited the number of bits they must exchange to obtain a good approximation is unreasonable. The proof uses a probabilistic construction of a collection of subsets of resources, which has useful combinatorial properties. These properties are then exploited in a reduction from a well known problem in communication complexity.

Definition 1. *A valuation function v is said to be additive (linear) if for every $S \subseteq M$ $v(S) = \sum_{j \in S} v(\{j\})$.*

Informally, a valuation function is fractionally-subadditive if it is the point-wise maximum over a set of additive (linear) valuations.

Definition 2. *A valuation function v is said to be fractionally subadditive if there is a set of additive (linear) valuations $\{a_1, \dots, a_l\}$ such that for every $S \subseteq M$ $v(S) = \max_{r \in [l]} a_r(S)$.*

The class of fractionally-subadditive valuation functions is known to be strictly contained in the class of subadditive valuations and to strictly contain all submodular valuations [17, 13].

Theorem 31 *Obtaining an approximation ratio of $m^{\frac{1}{4}-\epsilon}$ for fractionally-subadditive valuation functions requires exponential communication in m (for every $\epsilon > 0$).*

Proof. Fix a small $\epsilon > 0$. We prove the theorem for the case $n = k = \sqrt{m}$. The proof is by reduction from the SET DISJOINTNESS problem. In the SET DISJOINTNESS problem, we have n parties. Each party $i = 1, 2, \dots, n$ holds a t -bit string which specifies a subset $A_i \subseteq \{1, \dots, t\}$. The parties are required to distinguish between the two following extreme cases:

1. $\bigcap_{i=1}^n A_i \neq \emptyset$
2. for every $i \neq j$ it holds that $A_i \cap A_j = \emptyset$

The SET DISJOINTNESS problem was studied in [1] where it was shown to require $\Omega(\frac{t}{n^4})$ communication complexity. In [16] Nisan shows a lower bound due to Radhakrishnan and Srinivasan of $\Omega(\frac{t}{n})$.

For our reduction, we wish to first show the existence of an exponentially large family of sets of items with a useful combinatorial property:

Definition 3. A collection F of subsets of $[m]$ is said to have the r -intersection property if, for every $S, T \in F$, $|S \cap T| \leq r$.

Lemma 32 There exists collection F of subsets of $[m]$ such that:

- F has the $2m^\epsilon$ -intersection property.
- For each $S \in F$ it holds that $|S| \geq \sqrt{m}$.
- $|F| \geq e^{\alpha m^\epsilon}$ for some constant $\alpha > 0$.

Proof. We prove the lemma via a probabilistic construction. Each set $S \in F$ is constructed by choosing uniformly at random each element $l \in [m]$ to be in S with probability $m^{\frac{\epsilon-1}{2}}$. Let S and T be two such sets in F . We wish to show that each such set is “large” and that the intersection between them is “small”, with very high probability. We will use the Chernoff bound.

Claim. (Chernoff Bound) Let X_1, \dots, X_m be a set of m independent random variables that take values in $\{0, 1\}$ such that for every l , $Pr[X_l = 1] = p$. Then, for any δ is in the range $[0, 2e - 1]$ we have that:

$$Pr\left[\sum_{l=1}^m X_l > (1 + \delta)pm\right] \leq e^{-\frac{\delta^2 pm}{3}} \quad (1)$$

$$Pr\left[\sum_{l=1}^m X_l < (1 - \delta)pm\right] \leq e^{-\frac{\delta^2 pm}{3}} \quad (2)$$

For every resource $l \in [m]$ we define a random variable X_l that is assigned a value of 1 if $l \in S \cap T$ and of 0 otherwise. Observe, that the probability that $X_l = 1$ is $m^{\epsilon-1}$, and by 1:

$$Pr[|S \cap T| > 2m^\epsilon] = Pr[\sum_l X_l > 2m^\epsilon] < e^{-\frac{m^\epsilon}{3}}$$

Similarly, for every $S \in F$, define a random variable X_l that is assigned a value of 1 if $l \in S$ and 0 otherwise, by 2 we have that for any $\delta \in (0, 1)$:

$$Pr[|S| < (1 - \delta)m^{\frac{1+\epsilon}{2}}] = Pr[\sum_l X_l < (1 - \delta)m^{\frac{1+\epsilon}{2}}] < e^{-\frac{\delta^2 m^{\frac{1+\epsilon}{2}}}{3}}$$

We thus have that there is some constant $\alpha > 0$ such that:

$$Pr[|S \cap T| > 2m^\epsilon \text{ or } |S| < m^{\frac{1}{2}} \text{ or } |T| < m^{\frac{1}{2}}] < e^{-2\alpha m^\epsilon}$$

Since these bounds must hold for any S, T we get that as long as $|F|^2 \leq e^{2\alpha m^\epsilon}$ there is such a collection F . Thus, we know there exists a family of sets $F = \{S_1, \dots, S_t\}$ where $t = e^{\alpha m^\epsilon}$ for some constant $\alpha > 0$, with the property that every set is at least of size \sqrt{m} and $|S_i \cap S_j| \leq 2m^\epsilon \forall i \neq j \in t$.

Now, we show the reduction from the SET DISJOINTNESS problem. Let $1, 2, \dots, \sqrt{m}$ be the parties, and set $t = e^{\alpha m^\epsilon}$. Let A_i be the subset of $[t]$ held by party i . We construct an instance of CPPP with m resources, $n = \sqrt{m}$ agents (corresponding to the \sqrt{m} parties in the SET DISJOINTNESS problem), and set the number of resources to be chosen to $k = \sqrt{m}$. We identify each element $r \in [t]$ with a set S_r in the family F of subsets of $[m]$ described above. Each agent i constructs a valuation function v_i in the following manner: Let a_S denote the additive valuation that assigns a value of 1 to every resource in S and a value of 0 to every resource $j \notin S$. Let $v_i = \max\{a_{S_r} | r \in A_i\}$.

Observe that if $\bigcap_i A_i \neq \emptyset$ then there is a set S_r that has a corresponding additive valuation in all of the v_i 's. Hence, assigning a subset of S_r of size \sqrt{m} to the agents (simulated by the SET DISJOINTNESS parties) results in a social welfare value of m . What happens if for every two $i \neq j \in [t]$ $S_i \cap S_j = \emptyset$? We shall now show that in this case the optimal social welfare is $O(m^{\frac{3}{4}+\epsilon})$. This would mean that an approximation of $O(m^{\frac{1}{4}-\epsilon})$ to the CPPP with fractionally-subadditive valuations enables the distinction between the two extreme cases in the SET DISJOINTNESS problem. Therefore, we will then be able to conclude that $\Omega(\frac{t}{n})$ bits are required to do so (a number exponential in both n and m).

So, we are left with showing that if for every two $i \neq j \in [t]$ $S_i \cap S_j = \emptyset$ then the optimal social welfare is $O(m^{\frac{3}{4}+\epsilon})$. Assume, for the purpose of contradiction, that there is some set T of size \sqrt{m} such that the social welfare derived from T , $SW(T)$, is greater than $2m^{\frac{3}{4}+\epsilon}$. Let a_i be an additive valuation function of i for which v_i is maximized (for T). Observe that $SW(T) = \sum_{i \in [n]} a_i(T)$. Assume, w.l.o.g., that $T = \{1, \dots, \sqrt{m}\}$. For every resource $l \in T$, let x_l be the number of the a_i 's that assign a value of 1 to l . Observe that $SW(T) = \sum_{i \in [n]} a_i(T) = \sum_{l \in T} x_l$. Also observe that $\sum_{l \in T} x_l(x_l - 1) = \sum_{i \neq j} |S_i \cap S_j \cap T|$. Since we know that the cardinality of the intersection of every two sets cannot exceed $2m^\epsilon$ we now have that:

$$2m^{1+\epsilon} = 2n^2 m^\epsilon \geq \sum_{i \neq j} |S_i \cap S_j \cap T| = \sum_{l \in T} x_l(x_l - 1)$$

Using elementary calculus, it is easy to show that $\sum_{l \in T} x_l^2 \geq m^{\frac{1}{4}} \sum_{l \in T} x_l$. (This is due to the fact that the worst case ratio is achieved when all x_l 's are equal.)

Combining the last two equations gives us that

$$SW(T) = \sum_{l \in T} x_l \leq 2m^{\frac{3}{4} + \epsilon}.$$

A contradiction.

3.2 A Truthful \sqrt{m} Approximation Algorithm

We show that the result stated in the above theorem is nearly tight by presenting a simple *truthful* algorithm which obtains a $\min\{k, \sqrt{m}\}$ approximation ratio (for any value of k and n) and requires at most $n \cdot m$ value queries.

The Algorithm:

- Arbitrarily partition $[m]$ into $r = \max\{\frac{m}{k}, \sqrt{m}\}$ disjoint subsets of equal size S_1, \dots, S_r .
- Ask each agent to specify her value for each of the different subsets S_t .
- Choose the subset S_t that maximizes the social welfare $\sum_i v_i(S_t)$.

This algorithm, similar to that presented in [5], is a simple maximal-in-range algorithm [5, 6], and so it can be made truthful via VCG payments [19, 4, 11]. Observe that the algorithm indeed requires at most m value queries to be addressed to each of the n agents. Therefore, all that is left to show is that the algorithm provides the required approximation-ratio. We show this for subadditive valuations.

Proposition 33 *If v_1, \dots, v_n are subadditive then the algorithm provides an approximation ratio of $\min\{k, \sqrt{m}\}$.*

Proof. Let $k \leq \sqrt{m}$. Let O be a set of size k that maximizes the social welfare. Then, by (iterative use of) subadditivity, for every $i \in [n]$, $v_i(O) \leq \sum_{j \in O} v_i(\{j\})$. Hence, $\sum_{i \in [n]} v_i(O) \leq \sum_{i \in [n]} \sum_{j \in O} v_i(\{j\}) = \sum_{j \in [O]} \sum_{i \in [n]} v_i(\{j\})$. This implies that there is an element $j \in [m]$ such that the social welfare derived from j is at least $\frac{1}{|O|} = \frac{1}{k}$ of the optimal social welfare. This item j appears in one of the S_t 's, and so, because the valuations are non-decreasing, the social welfare derived from that S_t is also at least $\frac{1}{k}$ of the optimal social welfare. Since the algorithm optimizes over all the S_t 's it is bound to achieve the desired approximation ratio.

Let $k > \sqrt{m}$. Let O be the set of size k that maximizes the social welfare. Because the valuations are non-decreasing, $\sum_i v_i(O) \leq \sum_i v_i([m])$. Let $S_1, \dots, S_{\sqrt{m}}$ be some arbitrary partition of $[m]$ into \sqrt{m} disjoint subsets of size \sqrt{m} . Exploiting subadditivity in a way similar to that shown above implies that for one of these sets the social welfare is at least a $\frac{1}{\sqrt{m}}$ fraction of the social welfare for the entire set $[m]$. This concludes the proof of the proposition.

4 Inapproximability of CPPP with General Valuations

In this section we study CPPP with general valuations (but still normalized and non-decreasing). As the overlay networks example (Section 2), and others (elections etc.), illustrate, in many cases submodularity, and even subadditivity, do not apply. We prove strong inapproximability results for CPPP with general valuations in both the computational- and the communication-complexity models. In the communication-complexity model our lower bound is tight (a trivial matching upper bound exists).

Theorem 41 *Obtaining an approximation of $O(n^{\frac{1}{2}-\epsilon})$ to the social welfare in CPPP with general valuations, for any $\epsilon > 0$, is impossible unless $P = NP$. Obtaining an approximation of $O(n^{1-\epsilon})$ to the social welfare is impossible unless $P = ZPP$.*

Proof. We reduce from the Maximal Welfare Tree (MWT) problem studied in the context of distributed algorithmic mechanism design [10, 9, 14]. Our reduction preserves the hardness results for this problem as shown in [14]. In the MWT problem we are given a graph $G = (N, L)$ with a set of nodes N and links L . A unique destination node d is given and each node $a \in N \setminus \{d\}$ has a valuation function $v_a : P_a \rightarrow \mathcal{R}_{\geq 0}$, where P_a is used to denote the set of all simple paths from a to the destination d . The objective in MWT is to form a tree rooted in d which maximizes the social welfare, i.e., choose the tree T^* such that $T^* \in \operatorname{argmax}_{T \in T_L^d} \sum_{a \in N \setminus \{d\}} v_a(T)$, where T_L^d is the set of all possible trees in L rooted in d . We consider the special case of MWT in which for all $a \in N \setminus \{d\}$ we have $v_a : P_a \rightarrow \{0, 1\}$. It is known that for any $\epsilon > 0$ approximating MWT, even for this special case, within a factor of $O(n^{\frac{1}{2}-\epsilon})$ is impossible unless $P = NP$ and approximating within a factor of $O(n^{1-\epsilon})$ is impossible unless $P = ZPP$ [14].

The reduction from MWT is as follows: Given an instance of MWT such that the range of all valuation functions is $\{0, 1\}$, for each link $l \in L$ we associate a resource l' in CPPP and each node $a \in N \setminus \{d\}$ in MWT will correspond to an agent a' in CPPP. It remains to define the valuation function of a' . Note that since our interest is in showing a lower bound, we can adversarially set the number of chosen items to be $k = |N| - 1$. Now, let P_a^+ be the set of paths for which $v_a = 1$. For all $E \subseteq L$, the valuation function for the corresponding agent a' in CPPP is defined by:

$$v_{a'}(E) = \begin{cases} 1 & \exists P \in P_a^+ : P \subseteq E \text{ and } E \in T_L^d \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Observe that choosing a tree T in MWT with social welfare value $SW(T) = c$ corresponds to choosing a set of resources that induces the same social welfare value in CPPP. Conversely, choosing a set of resources T' in CPPP s.t. $SW(T') = c'$ necessarily means that we can trim T' to a set of edges T which forms a routing tree with d as its source, and that we have exactly c' nodes which have routes to d in T , and hence $SW(T) = c'$ in MWT.

Theorem 42 *Obtaining an approximation ratio of $(1 - \epsilon)n$ for general valuations requires exponential communication in m (for any $\epsilon > 0$ and for any $n \ll 2^m$).*

Proof. For CPPP with general valuation functions, n agents, m items and a parameter $1 \leq k \leq m$ we show a lower bound of $\Omega(\binom{m}{k} \cdot n^{-1})$ again by reducing from the SET DISJOINTNESS problem.

We construct an instance of CPPP with n agents in which no restrictions (except for being normalized and non-decreasing) apply to the agents' valuation functions. Let S_1, \dots, S_t be the (ordered) sets in the range of all possible allocations of size k . For each party i in SET DISJOINTNESS with the set $A_i \subseteq \{1, \dots, t\}$, we associate an agent i in CPPP with the following valuation function:

$$v_i(S_r) = \begin{cases} 1 & r \in A_i \\ 0 & \text{otherwise} \end{cases}$$

Observe that these valuation functions are indeed normalized and non-decreasing. Let S_l be the set which maximizes the social welfare, i.e., $l \in \operatorname{argmax}_{l \in [d]} |\{A_i \mid l \in A_i\}|$. To approximate the social welfare within a factor of $(1 - \epsilon)n$ for any $\epsilon > 0$, one must allocate some set S for which there are at least two agents i and j such that $v_i(S) = v_j(S) = 1$. Due to the above construction of the agents' valuation functions this necessarily implies deciding between the two extreme cases of the SET DISJOINTNESS problem. Thus, for $d = \binom{m}{k}$ we get a lower bound of $\Theta(\binom{m}{k} \cdot n^{-1})$ for CPPP with general valuation functions.

In the communication model a trivial matching upper bound of n exists: Query each agent i for her most valued set S_i of size k , and assign the agents a set $T \in \operatorname{argmax}_i v_i(S_i)$. It is easy to see that this indeed guarantees an n -approximation.

5 Acknowledgements

We warmly thank Christos Papadimitriou for many helpful and interesting discussions.

References

1. N. Alon, Y. Matias, and M. Szegedy. Journal of computer and system sciences. *The space complexity of approximating the frequency moment*, (58(1)):137–137, 1999.
2. David Andersen, Hari Balakrishnan, Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *SOSP '01: Proceedings of the eighteenth ACM symposium on Operating systems principles*, pages 131–145, New York, NY, USA, 2001. ACM.
3. Byung-Gon Chun, Rodrigo Fonseca, Ion Stoica, and John Kubiawicz. Characterizing selfishly constructed overlay routing networks. In *INFOCOM*, 2004.
4. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
5. S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *STOC*, 2005.

6. Shahar Dobzinski and Noam Nisan. Limitations of VCG-based mechanisms. In *STOC*, 2007.
7. Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *SODA 06*'.
8. Uriel Feige. On maximizing welfare when the utility functions are subadditive. In *STOC 06*'.
9. Joan Feigenbaum, Christos H. Papadimitriou, Rahul Sami, and Scott Shenker. A BGP-based mechanism for lowest-cost routing. *Distributed Computing*, 18(1):61–72, 2005.
10. Joan Feigenbaum, Rahul Sami, and Scott Shenker. Mechanism design for policy routing. In *PODC*, pages 11–20, 2004.
11. T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.
12. Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.
13. Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *ACM conference on electronic commerce*, 2001.
14. Hagay Levin, Michael Schapira, and Aviv Zohar. Interdomain routing and games. In *Proceedings of STOC 08*'.
15. G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions ii. *Math. Programming Study* 8, pages 73–87, 1978.
16. Noam Nisan. The communication complexity of approximate set packing and covering. In *ICALP 2002*.
17. Noam Nisan. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce*, 2000.
18. Christos H. Papadimitriou, Michael Schairer, and Yaron Singer. On the hardness of being truthful. In *FOCS*, 2008.
19. W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.
20. Andrew Chi-Chih Yao. Some complexity questions related to distributive computing. In *ACM Symposium on Theory of Computing (STOC)*, pages 209–213, 1979.