Truthful Mechanisms for Combinatorial AC Electric Power Allocation

Chi-Kin Chau, Khaled Elbassioni, Majid Khonji
Masdar Institute of Science and Technology, Abu Dhabi, UAE
{ckchau,kelbassioni,mkhonji}@masdar.ac.ae

ABSTRACT

Traditional studies of combinatorial auctions often only consider linear constraints (by which the demands for certain goods are limited by the corresponding supplies). The rise of smart grid presents a new class of auctions, characterized by quadratic constraints. Yu and Chau [AAMAS 13'] introduced the *complex-demand knapsack problem*, in which the demands are complex-valued and the capacity of supplies is described by the magnitude of total complex-valued demand. This naturally captures the power constraints in AC electric systems. In this paper, we provide a more complete study and generalize the problem to the multi-minded version, beyond the previously known $\frac{1}{2}$-approximation algorithm for only a subclass of the problem. More precisely, we give a truthful PTAS for the case $\phi \in [0, \frac{\pi}{2}]$, and a truthful bi-criteria FPTAS for the general case $\phi \in (\frac{\pi}{2}, \pi - \epsilon]$, where ϕ is the maximum angle between any two complex-valued demands.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

Keywords

Combinatorial Power Allocation; Multi-unit Combinatorial Auctions; Complex-Demand Knapsack Problem; Mechanism Design; Smart Grid

1. INTRODUCTION

Auctions are vital venues for the interactions of multi-agent systems, and their computational efficiency is critical for agent-based automation. Nonetheless, many practical auction problems are combinatorial in nature, requiring carefully designed time-efficient approximation algorithms. Although there have been decades of research in approximating combinatorial auction problems, traditional studies of combinatorial auctions often only consider linear constraints. Namely, the demands for certain goods are limited by the respective supplies, described by linear constraints.

Recently, the rise of smart grid presents a new class of auction problems. One of the salient characteristics is the presence of periodic time-varying entities (e.g., power, voltage, current) in AC (alternating current) electric systems, which are often expressed in terms of complex numbers[1]. In AC electric systems, it is natural to use a quadratic constraint, namely the magnitude of complex numbers, to describe the system capacity. Yu and Chau [12] introduced the *complex-demand knapsack problem* (CKP) to model a one-shot auction for combinatorial AC electric power allocation, which is a quadratic programming variant of the classical knapsack problem.

Furthermore, future smart grids will be automated by agents representing individual users. Hence, one might expect these agents to be self-interested and may untruthfully report their utilities or demands. This motivates us to consider truthful (aka. incentive-compatible) approximation mechanisms, in which it is in the best interest of the agents to report their true parameters. In [12] a monotone $\frac{1}{2}$-approximation algorithm that induces a deterministic truthful mechanism was devised for the complex-demand knapsack problem, which however assumes that all complex-valued demands lie in the positive quadrant.

In this paper, we provide a complete study and generalize the complex-demand knapsack problem to the multi-minded version, beyond the previously known $\frac{1}{2}$-approximation algorithm. More precisely, we give a (deterministic) truthful PTAS for the case $\phi \in [0, \frac{\pi}{2}]$, and a truthful bi-criteria FPTAS for the general case $\phi \in (\frac{\pi}{2}, \pi - \epsilon]$, where ϕ is the maximum angle between any two complex-valued demands. In an extended version of this paper, we also show that, unless P=NP, neither a PTAS can exist for the latter case nor any bi-criteria approximation algorithm with polynomial guarantees for the case when ϕ is arbitrarily close to π. Our results completely settle the open questions in [12].

Because of the paucity of space, some proofs are deferred to the extended paper.

2. RELATED WORK

Linear combinatorial auctions can be formulated as variants of the classical knapsack problem [3,6,8]. Notably, these include the *one-dimensional knapsack problem* (1DKP) where each indivisible item has only one single copy, and its multi-dimensional generalization, the *m-dimensional knapsack problem* (mDKP). There is an FPTAS for 1DKP [8].

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

[1] In the common terminology of power systems [7], the real part of complex-valued power is known as *active* power, the imaginary part is *reactive* power, whereas the magnitude is *apparent* power. Electric equipment has various active and reactive power requirements, whereas power transmission systems and generators are restricted by the supported apparent power.

In mechanism design setting, where each customer may untruthfully report her utility and demand, it is desirable to design *truthful* or *incentive-compatible* approximation mechanisms, in which it is in the best interest of each customer to reveal her true utility and demand [4]. In the so-called *single-minded case*, a *monotone* procedure can guarantee incentive compatibility [10]. While the straightforward FPTAS for 1DKP is not monotone, since the scaling factor involves the maximum item value, [2] gave a monotone FPTAS, by performing the same procedure with a series of different scaling factors irrelevant to the item values and taking the best solution out of them. Hence, 1DKP admits an truthful FPTAS. More recently, a truthful PTAS, based on dynamic programming and the notion of the so-called *maximal-in-range* mechanism, was given in [5] for the *multi-minded* case.

As to mDKP with $m \geq 2$, a PTAS is given in [6] based on the integer programming formulation, but it is not evident to see whether it is monotone. On the other hand, 2DKP is already inapproximable by an FPTAS unless P = NP, by a reduction from EQUIPARTITION [8]. Very recently, [9] gave a truthful FPTAS with $(1+\epsilon)$-violation for multi-unit combinatorial auctions with a constant number of distinct goods (including mDKP), and its generalization to the multiple-choice version, when m is fixed. Their technique is based on applying the VCG-mechanism to a rounded problem. Based on the PTAS for the multi-minded 1DKP developed in [5], they also obtained a truthful PTAS for the multiple-choice multidimensional knapsack problem.

In contrast, non-linear combinatorial auctions were explored to a little extent. Yu and Chau [12] introduced complex-demand knapsack problem, which models auctions with a quadratic constraint.

3. PROBLEM DEFINITIONS AND NOTATIONS

3.1 Complex-demand Knapsack Problem

We adopt the notations from [12]. Our study concerns power allocation under a capacity constraint on the magnitude of the total satisfiable demand (i.e., apparent power). Throughout this paper, we sometimes denote $\nu^{\rm R} \triangleq {\rm Re}(\nu)$ as the real part and $\nu^{\rm I} \triangleq {\rm Im}(\nu)$ as the imaginary part of a given complex number ν. We also interchangeably denote a complex number by a 2D-vector as well as a point in the complex plane. $|\nu|$ denotes the magnitude of ν.

We define the single-minded complex-demand knapsack problem (CKP) as follows:

$$\text{(CKP)} \qquad \max_{x_k \in \{0,1\}} \sum_{k \in \mathcal{N}} u_k x_k \tag{1}$$

$$\text{subject to} \qquad \Big| \sum_{k \in \mathcal{N}} d_k x_k \Big| \leq C. \tag{2}$$

where $d_k = d_k^{\rm R} + {\bf i} d_k^{\rm I} \in \mathbb{C}$ is the *complex-valued* demand of power for k-th user, $C \in \mathbb{R}_+$ is a real-valued capacity of total satisfiable demand in apparent power. Evidently, CKP is also NP-complete, because the classical 1-dimensional knapsack problem (1DKP) is a special case.

We note that the problem is invariant, when the arguments of all demands are rotated by the same angle. Without loss of generality, we assume that one of the demands, say d_1, is aligned along the positive real axis, and define a class of sub-problems for CKP, by restricting the maximum phase angle (i.e., the argument) that any other demand makes with d_1. In particular, we will write CKP$[\phi_1, \phi_2]$ for the restriction of problem CKP subject to $\phi_1 \leq \max_{k \in \mathcal{N}} \arg(d_k) \leq \phi_2$, where $\arg(d_k) \geq 0$ for all $k \in \mathcal{N}$. We remark that in realistic setting of power systems, the active power demand is positive (i.e., $d_k^{\rm R} \geq 0$), but the power factor (i.e., $\frac{d_k^{\rm R}}{|d_k|}$) is bounded by a certain threshold [1], which is equivalent to restricting the argument of complex-valued demands.

For complexity issues, we will need to specify how the inputs are described. Throughout the paper we will assume that each of the demands is given by its real and imaginary components, represented as rational numbers.

3.2 Non-single-minded Complex Knapsack Prob.

In this paper, we extend the single-minded CKP to general *non-single-minded* version, and then we apply the well-known *VCG-mechanism*, or equivalently the framework of *maximal-in-range* mechanisms [11]. The non-single-minded version is defined as follows. As above we assume a set $\mathcal{N}$ of n users: user k has a valuation function $v_k : \mathcal{D} \to \mathbb{R}_+$ over a (possibly infinite) set of demands $\mathcal{D} \subseteq \mathbb{C}$. We assume that $\mathbf{0} \in \mathcal{D}$ and $v_k(\mathbf{0}) = 0$ for all $k \in \mathcal{N}$. We further assume that each v_k is *monotone* with respect to a partial order "$\preceq$" defined on the elements of $\mathbb{C}$ as follows: for $d, f \in \mathbb{C}$, $d \succeq f$ if and only if

$$|d^{\rm R}| \geq |f^{\rm R}|, |d^{\rm I}| \geq |f^{\rm I}|, {\rm sgn}(d^{\rm R}) = {\rm sgn}(f^{\rm R}), {\rm sgn}(d^{\rm I}) = {\rm sgn}(f^{\rm I}).$$

(We assume $\mathbf{0} \preceq d$ for all $d \in \mathcal{D}$.) Then for all $k \in \mathcal{N}$, the monotonicity of $v_k(\cdot)$ means that $v_k(d) \geq v_k(f)$ whenever $d \succeq f$.

The non-single-minded problem can be described by the following program:

$$\text{(NsmCKP)} \qquad \max \sum_{k \in \mathcal{N}} v_k(d_k) \tag{3}$$

$$\text{s.t.} \quad \Big(\sum_{k \in \mathcal{N}} d_k^{\rm R}\Big)^2 + \Big(\sum_{k \in \mathcal{N}} d_k^{\rm I}\Big)^2 \leq C^2 \tag{4}$$

$$d_k \in \mathcal{D} \text{ for all } k \in \mathcal{N}. \tag{5}$$

Of particular interest is the *multi-minded* version of the problem (MultiCKP), defined as follows. Each user $k \in \mathcal{N}$ is interested only in a *polynomial-size* subset of demands $D_k \subseteq \mathcal{D}$ and declares her valuation only over this set. Note that the multi-minded problem can be modeled in the form (NsmCKP) by assuming w.l.o.g. that $\mathbf{0} \in D_k$, for each user $k \in \mathcal{N}$, and defining the valuation function $v_k : \mathcal{D} \to \mathbb{R}_+$ as follows:

$$v_k(d) = \max_{d_k \in D_k} \{v_k(d_k) : \ d_k \preceq d\}. \tag{6}$$

We shall assume that the demand set of each user lies completely in one of the quadrants, namely, either $d^{\rm R} \geq 0$ for all $d \in D_k$, or $d^{\rm R} < 0$ for all $d \in D_k$. Note that the single-minded version (which is CKP) is special case, where $|D_k| = 1$ for all k.

We will write MultiCKP$[\phi_1, \phi_2]$ for the restriction of the problem subject to $\phi_1 \leq \phi \leq \phi_2$ for all $d \in \mathcal{D}$ where $\phi \triangleq \max_{d \in \mathcal{D}} \arg(d)$ (and as before we assume $\arg(d) \geq 0$).

3.3 Multiple-choice Multidimensional Knapsack Problem

To design truthful mechanisms for NsmCKS, it will be useful to consider the *multiple-choice multidimensional knapsack* problem (Multi-mDKS) defined as follows, where we assume more generally that $\mathcal{D} \subseteq \mathbb{R}_+^m$ and a *capacity vector* $c \in \mathbb{R}_+^m$ is given. As before, a valuation function for each user k is given by (6). An *allocation* is given an assignment of a demand $d_k = (d_k^1, ..., d_k^m) \in \mathcal{D}$ for each user k, so as to satisfy the m-dimensional capacity constraint $\sum_{k\in\mathcal{N}} d_k \le c$. The objective is to find an allocation $\mathbf{d} = (d_1, \ldots, d_n) \in \mathcal{D}^n$ so as to maximize the sum of the valuations $\sum_{k\in\mathcal{N}} v_k(d_k)$. The problem can be described by the following program:

$$\max \quad \sum_{k\in\mathcal{N}} v_k(d_k) \tag{7}$$

$$\text{s.t.} \quad \sum_{k\in\mathcal{N}} d_k \quad \le c \tag{8}$$

$$d_k \in D_k, \qquad \forall k \in \mathcal{N}. \tag{9}$$

3.4 Approximation Algorithms

We present an explicit definition of approximation algorithms for our problem. Given a feasible allocation $\mathbf{d} = (d_1, \ldots, d_n) \in \mathcal{D}^n$ satisfying (4), we write $v(\mathbf{d}) \triangleq \sum_{k\in\mathcal{N}} v_k(d_k)$. Let $\mathbf{d}^*$ be an optimal allocation of NsmCKP (or (MultiCKP)) and $\text{Opt} \triangleq v(\mathbf{d}^*)$ be the corresponding total valuation. We are interested in an algorithm that outputs an allocation that is within a factor α of the optimum total valuation, but may violate the capacity constraint by at most a factor of β:

Definition 3.1. *For $\alpha \in (0, 1]$ and $\beta \ge 1$, a bi-criteria (α, β)-approximation to* NsmCKP *is an allocation $(d_k)_k \in \mathcal{D}^n$ satisfying*

$$\Big|\sum_{k\in\mathcal{N}} d_k\Big| \le \beta \cdot C \tag{10}$$

$$\textit{such that} \qquad \sum_{k\in\mathcal{N}} v_k(d_k) \ge \alpha \cdot \text{Opt}. \tag{11}$$

*Similarly we can define an (α, β)-approximation to (*MultiCKP*).*

In particular, *polynomial-time approximation scheme* (PTAS) is a $(1 - \epsilon, 1)$-approximation algorithm for any $\epsilon > 0$. The running time of a PTAS is polynomial in the input size for every fixed ϵ, but the exponent of the polynomial may depend on $1/\epsilon$. An even stronger notion is a *fully polynomial-time approximation scheme* (FPTAS), which requires the running time to be polynomial in both input size and $1/\epsilon$.

3.5 Truthful Mechanisms

This section follows the terminology of [10]. We define truthful (aka. incentive-compatible) approximation mechanisms for our problem. We denote by $\mathcal{X} \subseteq \mathcal{D}^n$ the set of *feasible allocations* in our problem (NsmCKP or Multi-mDKP).

Definition 3.2 (Mechanisms). *Let $\mathcal{V} \triangleq \mathcal{V}_1 \times \cdots \times \mathcal{V}_n$, where $\mathcal{V}_k$ is the set of all possible valuations of agent k. A mechanism $(\mathcal{A}, \mathcal{P})$ is defined by an allocation rule $\mathcal{A} : \mathcal{V} \to \mathcal{X}$ and a payment rule $\mathcal{P} : \mathcal{V} \to \mathbb{R}_+^n$. We assume that the utility of player k, under the mechanism, when it receives the vector of bids $v \triangleq (v_1, \ldots, v_n) \in \mathcal{V}$, is defined as $U_k(v) \triangleq \bar{v}_k(d_k(v)) - p_k(v)$, where $\mathcal{A}(v) = (d_1(v), \ldots, d_n(v))$, and $\mathcal{P}(v) = (p_1(v), \ldots, p_n(v))$ and $\bar{v}_k$ denotes the true valuation of player k.*

Namely, a mechanism defines an allocation rule and payment scheme, and the utility of a player is defined as the difference between her valuation over her allocated demand and her payment.

Definition 3.3 (Truthful Mechanisms). *A mechanism is said to be truthful if for all k and all $v_k \in \mathcal{V}_k$, and $v_{-k} \in \mathcal{V}_{-k}$, it guarantees that $U_k(\bar{v}_k, v_{-k}) \ge U_k(v_k, v_{-k})$.*

Namely, the utility of any player is maximized, when she reports the true valuation.

Definition 3.4 (Social Efficiency). *A mechanism is said to be α-socially efficient if for any $v \in \mathcal{V}$, it returns an allocation $\mathbf{d} \in \mathcal{X}$ such that the total valuation (also called social welfare) obtained is at least an α-fraction of the optimum: $v(\mathbf{d}) \ge \alpha \cdot \text{Opt}$.*

As in [5, 9, 11], our truthful mechanisms are based on using VCG payments with Maximal-in-Range (MIR) allocation rules:

Definition 3.5 (MIR). *An allocation rule $\mathcal{A} : \mathcal{V} \to \mathcal{X}$ is an MIR, if there is a range $\mathcal{R} \subseteq \mathcal{X}$, such that for any $v \in \mathcal{V}$, $\mathcal{A}(v) \in \text{argmax}_{\mathbf{d}\in\mathcal{R}}\, v(\mathbf{d})$.*

Namely, $\mathcal{A}$ is an MIR if it maximizes the social welfare over a fixed range $\mathcal{R}$ of feasible allocations. It is well-known (and also easy to prove by a VCG-based argument) that an MIR, combined with VCG payments (computed with respect to range $\mathcal{R}$), yields a truthful mechanism. If, additionally, the range $\mathcal{R}$ satisfies: $\max_{\mathbf{d}\in\mathcal{R}} v(\mathbf{d}) \ge \alpha \cdot \max_{\mathbf{d}\in\mathcal{X}} v(\mathbf{d})$, then such a mechanism is also α–socially efficient.

Finally a mechanism is *computationally efficient* if it can be implemented in polynomial time (in the size of the input).

4. A TRUTHFUL PTAS FOR MULTICKP$[0, \frac{\pi}{2}]$

The multi-minded mDKP problem was shown in [9] to have a $(1 - \epsilon)$-socially efficient truthful PTAS in the setting of *multi-unit auctions with a few distinct goods*, based on generalizing the result for the case $m = 1$ in [5]. We explain this result first in our setting, and then use it the Section 4.3 to derive a truthful PTAS for MultiCKP$[0, \frac{\pi}{2}]$.

4.1 A Truthful PTAS for Multi-mDKP

Let $c = (c^1, \ldots, c^m)$ be the capacity vector, and for any $d \in \mathcal{D}$, write $d = (d^1, \ldots, d^m)$. For any subset of users $N \subseteq \mathcal{N}$ and a partial selection of demands $\bar{\mathbf{d}} = (d_k \in \mathcal{D} :\ k \in N)$, such that $\sum_{k\in N} d_k \le c$, define the vector $b_{N,\bar{\mathbf{d}}} = (b^1_{N,\bar{\mathbf{d}}}, \ldots, b^m_{N,\bar{\mathbf{d}}}) \in \mathbb{R}_+^m$ as follows

$$b^i_{N,\bar{\mathbf{d}}} = \frac{c^i - \sum_{k\in N} d^i_k}{(n - |N|)^2}. \tag{12}$$

Following [9,11], we consider a restricted range of allocations defined as follows:

$$\mathcal{S} \triangleq \bigcup_{\substack{N\subseteq\mathcal{N},\ \bar{\mathbf{d}}=(d_k:\ k\in N):\ |N|\le\frac{m}{\epsilon},\\ d_k\in\mathcal{D}\ \forall k\in N}} \mathcal{S}_{N,\bar{\mathbf{d}}}, \tag{13}$$

where, for a set $N \subseteq \mathcal{N}$ and a partial selection of demands $\bar{\mathbf{d}} = (\bar{d}_k \in \mathcal{D} :\ k \in N)$,

$$\mathcal{S}_{N,\bar{\mathbf{d}}} \triangleq \Big\{(d_1, \ldots, d_n) \in \mathcal{D}^n \mid \sum_{k\in\mathcal{N}} d_k \le c, d_k = \bar{d}_k\ \forall k \in N,$$
$$d^i_k = r^i_k \cdot b^i_{N,\bar{\mathbf{d}}}\ \forall k \notin N\ \forall i\ \exists r^i_k \in \mathbb{Z}_+ \text{ such that } \sum_{k\notin N} r^i_k \le (n - |N|)^2\ \Big\}.$$

Note that the range $\mathcal{S}$ *does not* depend on the declarations $D_1, \ldots, D_n$. The following two lemmas establish that the range $\mathcal{S}$ is a good approximation of the set of all feasible allocations and that it can be optimized over in polynomial time. The first lemma is essentially a generalization of similar one for multi-unit auctions in [5], with the simplifying difference that we do not insist here on demands to be integral. The second lemma is also a generalization of similar result in [5], which was stated for the multi-unit auctions with a few distinct goods in [9]. For completeness, we give the details of a slightly simplified version here.

LEMMA 4.1. $\max_{\mathbf{d}\in\mathcal{S}} v(\mathbf{d}) \geq (1-\epsilon)\text{OPT}$.

LEMMA 4.2 ([5, 9]). *We can find* $\mathbf{d}^* \in \operatorname{argmax}_{\mathbf{d}\in\mathcal{S}} v(\mathbf{d})$ *using dynamic programming in time* $\left|\bigcup_k D_k\right|^{O(m/\epsilon)}$.

PROOF. We first observe that, due to the way the valuations are defined in (6), we may assume for the purpose of computing an optimal allocation $\mathbf{d}^*$ that $\mathcal{D} = \bigcup_k D_k$. Indeed, suppose that $\mathbf{d}^* = (d_1^*, \ldots, d_n^*) \in \mathcal{S}_{N,\bar{\mathbf{d}}^*}$, where $\bar{\mathbf{d}}^* = (d_k^* : k \in N)$, $d_k \notin D_k$ for some $k \in N$, and $d_k^* = (r_k^i \cdot b_{N,\mathbf{d}^*}^i : i \in [m])$. Then let us define a new allocation $\widetilde{d}$ as follows: for each $k \in N$, we choose $\widetilde{d}_k \in D_k$ such that $\widetilde{d}_k \preceq d_k^*$ and $v_k(\widetilde{d}_k) = v_k(d_k^*)$; we set $\bar{\mathbf{d}} = (\widetilde{d}_k : k \in N)$, and for $k \notin N$, define $\widetilde{d}_k = (r_k^i \cdot b_{N,\bar{\mathbf{d}}}^i : i \in [m])$. Note by (12) that $b_{N,\bar{\mathbf{d}}} \geq b_{N,\bar{\mathbf{d}}^*}$, and hence $v(\widetilde{\mathbf{d}}) \geq v(\mathbf{d}^*)$. We note furthermore that $\widetilde{\mathbf{d}} \in \mathcal{S}_{N,\bar{\mathbf{d}}}$, since for all i, we have

$$\sum_k (\widetilde{d}_k^i - d_k^{*,i}) = \sum_{k\in N} (\widetilde{d}_k^i - d_k^{*,i}) + \sum_{k'\notin N} \frac{r_{k'}^i}{(n-|N|)^2} \sum_{k\in N} (d_k^{*,i} - \widetilde{d}_k^i)$$
$$= \sum_{k\in N} (\widetilde{d}_k^i - d_k^{*,i}) \left(1 - \frac{\sum_{k'\notin N} r_{k'}^i}{(n-|N|)^2}\right) \leq 0,$$

since $\widetilde{d}_k^i \leq d_k^{*,i}$ and $\sum_{k'\notin N} r_{k'}^i \leq (n-|N|)^2$, for all i. It follows that $\sum_k \widetilde{d}_k \leq \sum_k d_k{}^* \leq c$, and hence $\widetilde{d} \in \mathcal{S}_{N,\bar{\mathbf{d}}}$ as claimed.

To maximize over $\mathcal{S}$, with the restriction that $\mathcal{D} = \bigcup_k D_k$, we iterate over all subsets $N \subseteq \mathcal{N}$ of size at most $\frac{m}{\epsilon}$ and all partial selections $\bar{\mathbf{d}} = (d_k \in D_k : k \in N)$. For each such choice $(N, \bar{\mathbf{d}})$, we use dynamic programming to find $\operatorname{argmax}_{\mathbf{d}\in\mathcal{S}_{N,\bar{\mathbf{d}}}} v(\mathbf{d})$. Let $b_{N,\bar{\mathbf{d}}}$ be as defined in (12). Without loss of generality, assume $\mathcal{N} \setminus N = \{1, \ldots, n-t\}$. For $k \in \mathcal{N} \setminus N$ and $r = (r^1, \ldots, r^m) \in \{0, 1, \ldots, (n-|N|)^2\}^m$, define $U(k, r)$ to be the maximum value obtained from a subset of users $\{1, 2, \ldots, k\} \subseteq \mathcal{N} \setminus N$, with user $j \in [k]$ having demand $\widehat{d}_j^i = r_j^i \cdot b_{N,\bar{\mathbf{d}}}^i$ for $i \in [m]$, where $r_j^i \in \{0, 1, \ldots, (n-|N|)^2\}$, and such that $\sum_{j\in[k]} r_k^i \leq r^i$. For two vectors $x, y \in \mathbb{R}^m$, let us denote by $x * y$ the vector with components $(x_1y_1, \ldots, x_my_m)$. Define $U(1, r) = -\infty$, if $r \ngeq \mathbf{0}$. Then we can use the following recurrence to compute $U(k, r)$:

$$U(1, r) \;=\; \max_r v_1(b_{N,\bar{d}} * r)$$
$$U(k+1, r) \;=\; \max_{r_{k+1}\leq r} \left\{ v_k(b_{N,\bar{d}} * r_{k+1}) + U(k, r - r_{k+1}) \right\}.$$

Note that the number of possible choices for r is at most n^{2m}, and hence the total time required by the dynamic program is $n^{O(m)}$. Finally, given the vector r that maximizes $U(n - |N|, r)$, we can obtain (by tracing back the optimal choices in the table) an optimal vector $r_k = (r_k^1, \ldots, r_k^m)$, for each $k \in \mathcal{N} \backslash N$. From this, we get an allocation $\widetilde{\mathbf{d}} \in \mathcal{S}$, by defining $\widetilde{d}_k = d_k$ for $k \in N$ and, for $k \notin N$, we choose $\widetilde{d}_k \in D_k$ such that $\widetilde{d}_k \preceq r_k * b_{N,\bar{\mathbf{d}}}$ and $v_k(\widetilde{d}_k) = v_k(r_k * b_{N,\bar{\mathbf{d}}})$. $\square$

It follows that an allocation rule defined as an MIR over range $\mathcal{S}$ yields a $(1-\epsilon)$-socially efficient truthful mechanism for MULTI-mDKP.

4.2 PTAS for MULTICKP$[0, \frac{\pi}{2}]$

We now apply the result in the previous section to the multi-minded complex-demand knapsack problem, when all agents are restricted to report their demands in the positive quadrant. We begin first by presenting a PTAS without strategic considerations in Section 4.2; then we show next in Section 4.3 how to use this PTAS within the aforementioned framework of MIR's to obtain a truthful mechanism.

In this section we assume that $\arg(d) \leq \frac{\pi}{2}$, that is, $d^{\mathrm{R}} \geq 0$ and $d^{\mathrm{I}} \geq 0$ for all $d \in \mathcal{D}$. As we shall see in Section 5, it is possible to get a $(1-\epsilon, 1+\epsilon)$-approximation by a reduction to the (MULTI-2DKP) problem. We note further that although there is a PTAS for mDKP with constant m [6], such a PTAS cannot be directly applied to MULTICKP$[0, \frac{\pi}{2}]$ by polygonizing the circular feasible region for MULTICKP$[0, \frac{\pi}{2}]$, because one can show that such an approximation ratio is at least a constant factor. This is the case, for instance, if the optimal solution consists of a few large (in magnitude) demands together with many small demands, and it is not clear at what level of accuracy we should polygonize the region to be able to capture these small demands. To overcome this difficulty, we first guess the large demands, then we construct a grid (or a lattice) on the remaining part of the circular region, defining a polygonal region in which we try to pack the maximum-utility set of demands. The latter problem is easily seen to be a special case of the MULTI-mDKP problem. The main challenge is to choose the granularity of the grid small enough to well-approximate the optimal, but also large enough so that the number of sides of the polygon, and hence m is a constant only depending on $1/\epsilon$.

Without loss of generality, we assume $\epsilon < \frac{1}{4}$ where $\frac{1}{\epsilon} \in \mathbb{Z}_+$. Given a feasible set of vectors $T \subseteq \mathcal{D}$ to MULTICKP$[0, \frac{\pi}{2}]$ (that is, $\left|\sum_{d\in T} d\right| \leq C$), we define $\mathcal{R}_T$ as the conic region bounded as the following (see Fig. 1a for an illustration).

$$\mathcal{R}_T \triangleq \left\{ \nu \in \mathbb{C} : |\nu| \leq C, \operatorname{Re}(\nu) \geq \operatorname{Re}(d_T) \text{ and } \operatorname{Im}(\nu) \geq \operatorname{Im}(d_T) \right\},$$

where $d_T \triangleq \sum_{d\in T} d$. Given $\mathcal{R}_T$, we define four points in the complex plane $(\pi'^1_T, \pi^1_T, \pi^2_T, \pi'^2_T)$ such that

$$\pi'^1_T = \left(0, \sqrt{C^2 - \operatorname{Re}(d_T)^2}\right), \quad \pi^1_T = \left(\operatorname{Re}(d_T), \sqrt{C^2 - \operatorname{Re}(d_T)^2}\right),$$
$$\pi'^2_T = \left(\sqrt{C^2 - \operatorname{Im}(d_T)^2}, 0\right), \quad \pi^2_T = \left(\sqrt{C^2 - \operatorname{Im}(d_T)^2}, \operatorname{Im}(d_T)\right).$$

Moreover, we define a grid in the region $\mathcal{R}_T$ by interlacing equidistant horizontal and vertical lines with separation $\frac{\epsilon}{4} w_T^{\mathrm{I}}$ and $\frac{\epsilon}{4} w_T^{\mathrm{R}}$, where

$$w_T^{\mathrm{I}} \triangleq \sqrt{C^2 - \operatorname{Re}(d_T)^2} - \operatorname{Im}(d_T),\; w_T^{\mathrm{R}} \triangleq \sqrt{C^2 - \operatorname{Im}(d_T)^2} - \operatorname{Re}(d_T).$$

Thus, the lines of the grid intersect the circular boundary of region $\mathcal{R}_T$ at a set of points $P_T(\epsilon)$, and we let $m_T(\epsilon) \triangleq |P_T(\epsilon)| \leq \frac{8}{\epsilon} + 2$. The convex hull of the set of points $P_T(\epsilon) \cup \{\pi'^1_T, \pi^1_T, \pi^2_T, \pi'^2_T, 0\}$ defines a polygonized region, which we denote by $\mathcal{P}_T(\epsilon)$ (see Fig. 1a for an illustration).

Remark 1: For simplicity of presentation, in this section, we will ignore the issue of finite precision needed to represent intermediate calculations (such as the square roots above, or the intersection points of the lines of the gird with the boundary of the circle); we will deal such issues in the next section.

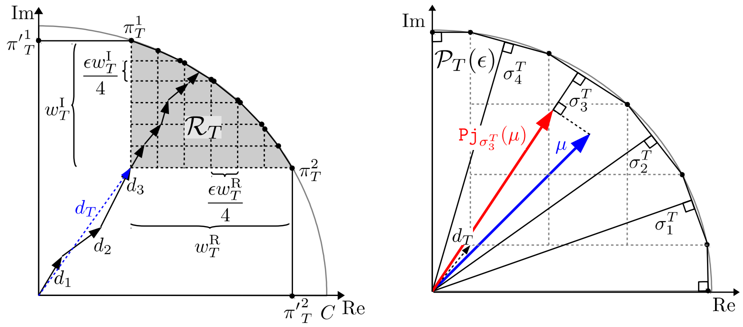

(a) We illustrate the region $\mathcal{R}_T$ by the shaded area and $P_T(\epsilon)$ by the black dots.

(b) Each in $\{\sigma_i^T\}$ is a vector (starting at the origin) perpendicular to each boundary edge of $\mathcal{P}_T(\epsilon)$.

Definition 4.3. *Consider a feasible set T to* MultiCKP$[0, \frac{\pi}{2}]$. *We define an approximate problem* (PGZ$_T$) *by polygonizing* MultiCKP$[0, \frac{\pi}{2}]$*:*

$$(\mathrm{PGZ}_T) \qquad \max \sum_{k\in\mathcal{N}} v_k(d_k)$$
$$\text{s.t.} \qquad \sum_{k\in\mathcal{N}} d_k \in \mathcal{P}_T(\epsilon)$$
$$d_k \in \mathcal{D}, \qquad \forall k \in \mathcal{N}.$$

Given two complex numbers μ and ν, we denote the projection of μ on ν by $\mathtt{Pj}_\nu(\mu) \triangleq \frac{\nu}{|\nu|}(\mu^{\mathrm{R}}\nu^{\mathrm{R}} + \mu^{\mathrm{I}}\nu^{\mathrm{I}})$. Given the convex hull $\mathcal{P}_T(\epsilon)$, we define a set of $m_T(\epsilon)$ vectors $\{\sigma_i^T\}$, each of which is perpendicular to each boundary edge of $\mathcal{P}_T(\epsilon)$ and starting at the origin (see Fig. 1b for an illustration).

Definition 4.4. *Consider a feasible set T to* MultiCKP$[0, \frac{\pi}{2}]$. *We define a* Multi-mDKP *problem based on* $\{\sigma_i^T\}$*:*

$$(\textsc{Multi-}m\mathrm{DKP}\{\sigma_i^T\}) \qquad \max \sum_{k\in\mathcal{N}} v_k(d_k) \tag{14}$$
$$\text{s.t.} \quad \sum_{k\in\mathcal{N}} \mathtt{Pj}_{\sigma_i^T}(d_k) \le |\sigma_i^T|, \quad \forall i = 1, \ldots, m_T(\epsilon), \tag{15}$$
$$d_k \in \mathcal{D}, \qquad \forall k \in \mathcal{N}. \tag{16}$$

Lemma 4.5. *Given a feasible set T to* MultiCKP$[0, \frac{\pi}{2}]$, PGZ$_T$ *and* Multi-mDKP$\{\sigma_i^T\}$ *are equivalent.*

Our PTAS for MutliCKP$[0, \frac{\pi}{2}]$ is described in Algorithm MultiCKP-PTAS, which enumerates every subset partial selection T of at most $\frac{1}{\epsilon}$ demands, then finds a near optimal allocation for each polygonized region $\mathcal{P}_T(\epsilon)$ using the PTAS of Multi-mDKP from Section 4.1, which we denote by Multi-mDKP-PTAS$[\cdot]$.

Theorem 4.6. *For any $\epsilon > 0$, Algorithm* MultiCKP-PTAS *finds a $(1-3\epsilon, 1)$-approximation to* MultiCKP$[0, \frac{\pi}{2}]$. *The running time of the algorithm is* $\left|\bigcup_k D_k\right|^{O(\frac{1}{\epsilon^2})}$.

Algorithm 1 MultiCKP-PTAS$(\{v_k, D_k\}_{k\in\mathcal{N}}, C, \epsilon)$

Require: Users' multi-minded valuations $\{v_k, D_k\}_{k\in\mathcal{N}}$; capacity C; accuracy parameter ϵ
Ensure: $(1-3\epsilon)$-allocation $(\widehat{d}_1, \ldots, \widehat{d}_n)$ to MultiCKP$[0, \frac{\pi}{2}]$
1: $(\widehat{d}_1, \ldots, \widehat{d}_n) \leftarrow (\mathbf{0}, \ldots, \mathbf{0})$
2: **for** each subset $T \subseteq \bigcup_k D_k$ of size at most $\frac{1}{\epsilon}$ s.t. $\left|\sum_{d\in T} d\right| \le C$ **do**
3: Set $d_T \leftarrow \sum_{d\in T} d$, and define the corresponding vectors $\{\sigma_i^T\}$
4: Obtain $(d_1, \ldots, d_n) \leftarrow$ Multi-mDKP-PTAS [Multi-mDKP$\{\sigma_i^T\}$] within accuracy ϵ
5: **if** $\sum_k v_k(\widehat{d}_k) < \sum_k v_k(d_k)$ **then**
6: $(\widehat{d}_1, \ldots, \widehat{d}_n) \leftarrow (d_1, \ldots, d_n)$
7: **end if**
8: **end for**
9: **return** $(\widehat{d}_1, \ldots, \widehat{d}_n)$

Proof. First, the upper bound on the running time of Algorithm MultiCKP-PTAS is due to the fact that each of the $O\left(\left|\bigcup_k D_k\right|^{\frac{1}{\epsilon}}\right)$ iterations in line 2 requires invoking the PTAS of Multi-mDKP, which in turn takes $\left|\bigcup_k D_K\right|^{O(m/\epsilon)}$ time, by Lemma 4.2, where $m = O(\frac{1}{\epsilon})$.

Clearly the algorithm outputs a feasible allocation by Lemma 4.5. To prove the approximation ratio, we show in Lemma 4.7 below that, for any optimal (or feasible) allocation $(d_1^*, \ldots, d_n^*)$, we can construct another feasible allocation $(\widetilde{d}_1, \ldots, \widetilde{d}_n)$ such that $\sum_k v_k(\widetilde{d}_k) \ge (1-2\epsilon)\sum_k v_k(d_k^*)$ and $(\widetilde{d}_1, \ldots, \widetilde{d}_n)$ is feasible to PGZ$_T$ for some T of size at most $\frac{1}{\epsilon}$. By Lemma 4.5, invoking the PTAS of Multi-mDKP$\{\sigma_i^T\}$ gives a $(1-\epsilon)$-approximation $(\widehat{d}_1, \ldots, \widehat{d}_k)$ to PGZ$_T$. Then

$$\sum_k v_k(\widehat{d}_k) \ge (1-\epsilon)\sum_k v_k(\widetilde{d}_k) \ge (1-\epsilon)(1-2\epsilon)\textsc{Opt} \ge (1-3\epsilon)\textsc{Opt}.$$

We provide an explicit construction of the allocation $(\widetilde{d}_1, \ldots, \widetilde{d}_n)$ in Algorithm 2, thus completing the proof by Lemma 4.7. □

Lemma 4.7. *Consider a feasible allocation* $\mathbf{d} = (d_1, \ldots, d_n)$ *to* MultiCKP$[0, \frac{\pi}{2}]$. *Then we can find a set* $T \subseteq \{d_1, \ldots, d_n\}$ *and construct an allocation* $\widetilde{\mathbf{d}} = (\widetilde{d}_1, \ldots, \widetilde{d}_n)$, *such that* $|T| \le \frac{1}{\epsilon}$ *and* $\widetilde{\mathbf{d}}$ *is a feasible solution to* PGZ$_T$ *and* $v(\widetilde{\mathbf{d}}) \ge (1-2\epsilon)v(\mathbf{d})$.

Proof. In Algorithm 2, let $\bar{\ell}$ and $T_{\bar{\ell}}$ be the values of ℓ and T_ℓ at the end of the repeat-until loop (line 9).

The basic idea of Algorithm 2 is that we first construct a nested sequence of sets of demands $T_0 \subset T_1 \subset \ldots \subset T_{\bar{\ell}}$, such that a demand is included in each iteration if it has either a large real component or a large imaginary component. The iteration proceeds until a sufficiently large number of demands have been summed up (namely, $|T_{\bar{\ell}}| \ge \frac{1}{\epsilon}$), or no demands with large components remain. At the end of the iteration, if the condition in line 11 holds, then $S = T_\ell$, i.e., the whole set S can be packed within the polygonized region $\mathcal{P}_{T_{\bar{\ell}}}(\epsilon)$. Otherwise, we find a subset of S that is feasible to PGZ$_{T_{\bar{\ell}}}$.

To do so, we partition $S \backslash T_{\bar{\ell}}$ into at least $\frac{1}{\epsilon}-1$ groups, each having a large component along either the real or the imaginary axes, with respect to the boundaries of the region $\mathcal{R}_{T_{\bar{\ell}}}$. Then removing the group with smallest utility among these, or removing one of the large demands with smallest utility will ensure that remaining demands have a large utility and can be packed within $\mathcal{P}_{T_{\bar{\ell}}}(\epsilon)$.

We then have to consider two cases (line 14): (i) $|T_{\bar{\ell}}|$ becomes at least $\frac{1}{\epsilon}$, or (ii) $S^{\mathrm{R}}_{\bar{\ell}} \cup S^{\mathrm{I}}_{\bar{\ell}} = \varnothing$. For case (i), we proceed to line 16 – we combine the demands in $S \backslash T_{\bar{\ell}}$ into a group V_1. Note that removing any one demand $k \in T_{\bar{\ell}}$ will make $T_{\bar{\ell}} \backslash \{k\}$ as a feasible solution to $\mathrm{PGZ}_{T_{\bar{\ell}}}$ (since the lengths $w^{\mathrm{R}}_{T_i}$ and $w^{\mathrm{I}}_{T_i}$ are monotone decreasing for $i = 1, 2, \ldots$). For case (ii), we can apply Lemma 4.8 below to partition $S \backslash T_{\bar{\ell}}$ into at least $\frac{1}{\epsilon} - 1$ groups $\{V_1, \ldots, V_h\}$, where each group V_j has a large total component along either the real or the imaginary axes (precisely, greater than $\frac{\epsilon}{4} w^{\mathrm{R}}_{\bar{\ell}}$ or $\frac{\epsilon}{4} w^{\mathrm{I}}_{\bar{\ell}}$ respectively). This implies that removing any group V_j will make $T_{\bar{\ell}} \cup \bigcup_{j' \neq j} V_{j'}$ a feasible solution to $\mathrm{PGZ}_{T_{\bar{\ell}}}$.

To conclude, there are either (i) at least $\frac{1}{\epsilon}$ demands in $|T_{\bar{\ell}}|$, or (ii) $S^{\mathrm{R}}_{\bar{\ell}} \cup S^{\mathrm{I}}_{\bar{\ell}} = \varnothing$. We define S' by deleting a minimum utility demand or group of demands from S (lines 24 and 27). Then, we set $\widetilde{d}_k = d_k$ if $k \in S'$ and $\widetilde{d}_k = 0$ if $k \notin S'$. Hence, in case (i), $v(\widetilde{\mathbf{d}}) \geq (1 - \epsilon)\textsc{Opt}$, and in case (ii), $v(\widetilde{\mathbf{d}}) \geq (1 - \frac{1}{h})\textsc{Opt} \geq \frac{1-2\epsilon}{1-\epsilon} \cdot \textsc{Opt} \geq (1 - 2\epsilon)\textsc{Opt}$. □

Remark 2: It is important to note in algorithm Construct that, when we drop a single vector from T_ℓ (when the condition $v_{\widehat{k}}(d_{\widehat{k}}) < \sum_{k: d_k \in V_{\widehat{j}}} v_k(d_k)$ in line 22 holds), then we can redefine $T = T_\ell$ and use this to define the polygon $\mathcal{P}_T(\epsilon)$. In particular, we may assume when solving problem Multi-mDKP$\{\sigma_i^T\}$ that all the vectors in T are *included* in the solution.

Lemma 4.8. *Consider a set of demands $S \subseteq \mathcal{D}$ and $T \subseteq S$, such that*

- *S is feasible solution to* MultiCKP$[0, \frac{\pi}{2}]$, *but S is not a feasible solution to* PGZ_T
- *$d^{\mathrm{R}} \leq \frac{\epsilon}{4} w^{\mathrm{R}}_T$ and $d^{\mathrm{I}} \leq \frac{\epsilon}{4} w^{\mathrm{I}}_T$, for all $d \in S \backslash T$.*

Then there exists a partition $\{V_1, \ldots, V_h\}$ of $S \backslash T$ such that

- *either (i) $\sum_{d \in V_j} d^{\mathrm{R}} \geq \frac{\epsilon}{4} w^{\mathrm{R}}_T$ for all $j = 1, \ldots, h$,*
- *or (ii) $\sum_{d \in V_j} d^{\mathrm{I}} \geq \frac{\epsilon}{4} w^{\mathrm{I}}_T$ for all $j = 1, \ldots, h$.*

where $h \in [\frac{1}{\epsilon} - 1, \frac{4}{\epsilon})$.

Algorithm 2 Construct$(\{v_k\}_{k \in \mathcal{N}}, \mathbf{d}, C, \epsilon)$

Require: Users' valuations $\{v_k\}_{k \in \mathcal{N}}$; a feasible allocation $\mathbf{d}$; capacity C; accuracy parameter ϵ
Ensure: A set of demands $T \subseteq \{d_1, \ldots, d_n\}$ and a feasible allocation $\widetilde{\mathbf{d}}$

1: $S \leftarrow \{d_1, \ldots, d_n\}$; $\widetilde{\mathbf{d}} = \mathbf{d}$; $\ell \leftarrow 0$; $T_\ell \leftarrow \varnothing$; $\eta_\ell \leftarrow 0$
  ▷ *Find a subset of large demands T*
2: **repeat**
3: $\ell \leftarrow \ell + 1$
4: $d_{T_\ell} \leftarrow \sum_{d \in T_{\ell-1}} d$
5: $w^{\mathrm{I}}_{T_\ell} \leftarrow \sqrt{C^2 - \mathrm{Re}(d_{T_\ell})^2} - \mathrm{Im}(d_{T_\ell})$; $w^{\mathrm{R}}_{T_\ell} \leftarrow \sqrt{C^2 - \mathrm{Im}(d_{T_\ell})^2} - \mathrm{Re}(d_{T_\ell})$
6: $S^{\mathrm{R}}_\ell \leftarrow \{d \in S \backslash T_\ell \mid d^{\mathrm{R}} > \frac{\epsilon}{4} w^{\mathrm{R}}_{T_\ell}\}$; $S^{\mathrm{I}}_\ell \leftarrow \{d \in S \backslash T_\ell \mid d^{\mathrm{I}} > \frac{\epsilon}{4} w^{\mathrm{I}}_{T_\ell}\}$
7: $T_\ell \leftarrow T_\ell \cup S^{\mathrm{R}}_\ell \cup S^{\mathrm{I}}_\ell$
8: $\eta_\ell \leftarrow \eta_{\ell-1} + \sum_{d \in S^{\mathrm{R}}_\ell \cup S^{\mathrm{I}}_\ell} d$
9: **until** $|T_\ell| \geq \frac{1}{\epsilon}$ or $S^{\mathrm{R}}_\ell \cup S^{\mathrm{I}}_\ell = \varnothing$ or $S \backslash T_\ell = \varnothing$
10: $\kappa \leftarrow \sum_{d \in S \backslash T_\ell} d$
11: **if** $S \backslash T_\ell = \varnothing$ or $\eta_\ell + \kappa \in \mathcal{P}_{T_\ell}(\epsilon)$ **then**
12: **return** $(T_\ell, \mathbf{d})$
13: **else**
  ▷ *Find a subset $S' \subset S$ that is feasible to* PGZ_{T_ℓ}
14: **if** $|T_\ell| \geq \frac{1}{\epsilon}$ **then**
15: $T_\ell \leftarrow$ the set of the first $\frac{1}{\epsilon}$ elements added to T_ℓ
16: $h \leftarrow 1$; $V_1 \leftarrow S \backslash T_\ell$
17: **else**
18: Find a partition $V_1, \ldots, V_h$ over $S \backslash T_\ell$ such that either (i) $\sum_{d \in V_j} d^{\mathrm{R}} \geq \frac{\epsilon}{4} w^{\mathrm{R}}_{T_\ell}$ for all $j = 1, \ldots, h$, or (ii) $\sum_{d \in V_j} d^{\mathrm{I}} \geq \frac{\epsilon}{4} w^{\mathrm{I}}_{T_\ell}$ for all $j = 1, \ldots, h$
19: **end if**
20: Pick $\widehat{k} \in \mathrm{argmin}\{v_k(d_k) \mid d_k \in T_\ell\}$
21: Pick $\widehat{j} \in \mathrm{argmin}\{\sum_{k: d_k \in V_j} v_k(d_k) \mid j = 1, \ldots, h\}$
22: **if** $v_{\widehat{k}}(d_{\widehat{k}}) < \sum_{k: d_k \in V_{\widehat{j}}} v_k(d_k)$ **then**
23: $\widetilde{d}_{\widehat{k}} = 0$
24: **return** $(T_\ell, \widetilde{\mathbf{d}})$
25: **else**
26: $\widetilde{d}_k \leftarrow 0$ for all $k : d_k \in V_{\widehat{j}}$
27: **return** $(T_\ell, \widetilde{\mathbf{d}})$
28: **end if**
29: **end if**

4.3 A Truthful PTAS for multiCKP$[0, \frac{\pi}{2}]$

We now state our main result for this section.

Theorem 4.9. *For any $\epsilon > 0$ there is a $(1 - 3\epsilon)$-socially efficient truthful mechanism for* multiCKP$[0, \frac{\pi}{2}]$. *The running time is $\left|\bigcup_k D_k\right|^{O(\frac{1}{\epsilon^2})}$.*

Proof. It is enough to define a declaration-independent range $\mathcal{S}$ of feasible allocations, such that $\max_{\mathbf{d} \in \mathcal{S}} v(\mathbf{d}) \geq (1 - 3\epsilon) \cdot \textsc{Opt}$, and we can optimize over $\mathcal{S}$ in the stated time. For every set $T \subseteq \mathcal{D}$ of size at most $\frac{1}{\epsilon}$, we solve a slightly modified version of problem Multi-mDKP$\{\sigma_i^T\}$:

- We choose the grid lines from a *fixed set*, where all horizontal and vertical lines are at distances, form the real and imaginary axes, which are integer multiples of $\frac{C}{2^i}$, for some integer $i \in \mathbb{Z}_+$. In particular, instead of defining the grid separation distances to be $(\frac{\epsilon w^{\mathrm{R}}_T}{4}, \frac{\epsilon w^{\mathrm{I}}_T}{4})$, we use instead $(\frac{C}{2^{i(T)}}, \frac{C}{2^{j(T)}})$, where $i(T)$ and $j(T)$ are the smallest integers such that

$$\frac{C}{2^{i(T)}} \leq \frac{\epsilon w^{\mathrm{R}}_T}{4} \text{ and } \frac{C}{2^{j(T)}} \leq \frac{\epsilon w^{\mathrm{I}}_T}{4}.$$

  In this case, we say that the vertical and horizontal lines are *at levels* $i(T)$ and $j(T)$, respectively. Let us denote by $\mathcal{L}_1(i(T))$ the vertical grid lines at level $i(T)$, and by $\mathcal{L}_2(j(T))$ the horizontal grid lines at level $j(T)$.

- We also slightly change the definition of the polygon $\mathcal{P}_T(\epsilon)$ by *expanding* the region $\mathcal{R}_T$ slightly so that its vertical and horizontal boundary lines are from $\mathcal{L}_1(i(T))$ an $\mathcal{L}_2(j(T))$, respectively.

- We solve problem Multi-mDKP$\{\sigma_i^T\}$, imposing all vectors $d \in T$ are in the solution; see Remark 2.

It is straightforward to verify that these changes will only possibly increase the size of $\mathcal{P}_T(\epsilon)$ by a factor of 2, but otherwise, all other claims (in particular, Lemma 4.7) remain valid. As we shall see below, these modifications are only needed to ensure computational efficiency.

For $T \subseteq \mathcal{D}$, let $G(T)$ be the set of vectors in $\mathbb{C}$ defined by the union of (a) the (component-wise) minimal grid points z inside $\mathcal{R}_T$, only considering lines $\mathcal{L}_1(i(T)+1)$ and $\mathcal{L}_2(j(T)+1)$, that have *exclusively* either $i(\{z\}) = i(T)+1$ *or* $j(\{z\}) = j(T)+1$, and (b) the (component-wise) minimal grid points z inside $\mathcal{R}_T$, only considering lines in $\mathcal{L}_1(i(T)+1)$ and $\mathcal{L}_2(j(T)+1)$, that have $i(\{z\}) = i(T)+1$ *and* $j(\{z\}) = j(T)+1$. For $d \in \mathcal{D}$, let us denote by $\mathcal{S}_z$ the range of feasible allocations defined as in (13) with respect to the MULTI-mDKP problem with constraints (15)-(16), when $T = \{z\}$ (and hence, $d_T = z$). Note that $|G(T)| = O(\frac{1}{\epsilon})$. Then we define the range $\mathcal{S}(\mathcal{D})$ as the union:

$$\mathcal{S}(\mathcal{D}) \triangleq \bigcup_{T\subseteq\mathcal{D}:\ |T|\le\frac{1}{\epsilon}} \left(\bigcup_{z\in G(T)} \mathcal{S}_z \right).$$

By Lemmas 4.1 and 4.7, we have $\max_{\mathbf{d}\in\mathcal{S}(\mathcal{D})} v(\mathbf{d}) \ge (1-3\epsilon)\text{OPT}$. It remains to argue that we can efficiently optimize over $\mathcal{S}$. This essentially follows from the following claim.

CLAIM 4.10. *Let $T, T' \subseteq \mathcal{D}$ be such that $d \preceq d'$ for all $d \in T$ and $d' \in T'$. Consider a vector $\kappa \in \mathbb{C}$ such that $d_{T'} + \kappa \in \mathcal{P}_{T'}(\epsilon)$. Then either (i) $d_T + \kappa \in \mathcal{P}_T(\epsilon)$, or (ii) $i(T') \le i(T)+1$ and $j(T') \le j(T)+1$.*

PROOF. Suppose that $d_T + \kappa \notin \mathcal{P}_T(\epsilon)$. Then it also holds that $d_{T'} + \kappa \notin \mathcal{P}_T(\epsilon)$ (since $d_{T'} \succeq d_T$). This implies that both $d_T+\kappa$ and $d_{T'}+\kappa$ lie within the same grid cell in $\mathcal{P}_T(\epsilon)$, and hence $d^{\text{R}}_{T'} - d^{\text{R}}_T \le \frac{C}{2^{i(T)}}$ and $d^{\text{I}}_{T'} - d^{\text{I}}_T \le \frac{C}{2^{j(T)}}$. Now, form $w^{\text{R}}_T = w^{\text{R}}_{T'} + d^{\text{R}}_{T'} - d^{\text{R}}_T$, $\frac{\epsilon w^{\text{R}}_{T'}}{8} < \frac{C}{2^{i(T')}}$, and $\frac{C}{2^{i(T)}} \le \frac{\epsilon w^{\text{R}}_T}{4}$, follows that $i(T') \le i(T)+1$. Similarly, we have $j(T') \le j(T)+1$. □

By this claim, we can solve the optimization problem over $\mathcal{S}$ assuming that $\mathcal{D} = \bigcup_k D_k$, that is, $\max_{\mathbf{d}\in\mathcal{S}(\mathcal{D})} v(\mathbf{d}) = \max_{\mathbf{d}\in\mathcal{S}(\bigcup_k D_k)} v(\mathbf{d})$. One direction "$\ge$" is obvious; so let us show that $\max_{\mathbf{d}\in\mathcal{S}(\mathcal{D})} v(\mathbf{d}) \le \max_{\mathbf{d}\in\mathcal{S}(\bigcup_k D_k)} v(\mathbf{d})$. Suppose that $\mathbf{d}^* = (d^*_1,\dots,d^*_n)$ is an optimal allocation over $\mathcal{S}$, but such that $\mathbf{d}^* \in \mathcal{S}_{z'}$ for some $z' \in G(T')$, $T' \subseteq \mathcal{D}$, and $T' \not\subseteq \bigcup_k D_k$. Define an allocation $\widetilde{d}$ as follows: Let $N = \{k:\ d^*_k \in T'\}$; for each $k \in N$, we choose $\widetilde{d}_k \in D_k$ such that $\widetilde{d}_k \preceq d^*_k$ and $v_k(\widetilde{d}_k) = v_k(d^*_k)$, and we (temporarily) keep $\widetilde{d}_k = d^*_k$ if $k \notin N$. Let us apply the statement of the claim with $T = \{\widetilde{d}_k :\ k \in N\}$ and $\kappa = \sum_{k:k\notin N} d^*_k$. If (i) holds then $d_T + \kappa \in \mathcal{P}_T(\epsilon)$ and therefore we have

$$\max_{\mathbf{d}\in\mathcal{S}(\mathcal{D})} v(\mathbf{d}) \le \max_{\substack{\mathbf{d}\in\mathcal{S}(\{\mathbf{d}:\ d_k\in D_k\ \forall k\in N,\\ d_k\in\mathcal{D}\ \forall k\notin N\})}} v(\mathbf{d}) = \max_{\mathbf{d}\in\mathcal{S}(\bigcup_k D_K)}, \quad (17)$$

where the equality follows from (the proof of) Lemma 4.2. On the other hand, if (ii) holds, then $i(T') \in \{i(T), i(T)+1\}$ and $j(T') \in \{j(T), j(T)+1\}$. In this case, if $i(T') = i(T)$ and $j(T') = j(T)$ then $\mathcal{P}_{T'}(\epsilon) \subseteq \mathcal{P}_T(\epsilon)$ (since $d_T \preceq d_{T'}$); otherwise, there is a point $z \in G(T)$ such that $z \preceq z'$, $i(\{z\}) = i(\{z'\})$ and $j(\{z\}) = j(\{z'\})$. Then $d_{T'} + \kappa \preceq d_{T'} + \kappa \in \mathcal{P}_{\{z\}}(\epsilon)$, and we get again (17). □

5. A TRUTHFUL FPTAS FOR MULTICKP$[0, \pi\text{-}\varepsilon]$

As in [9], the basic idea is to round off the set of possible demands to obtain a range, by which we can optimize over in polynomial time using dynamic programming (to obtain an MIR).

Let $\theta = \max\{\phi - \frac{\pi}{2}, 0\}$, where $\phi \triangleq \max_{d\in\mathcal{D}} \arg(d)$. We assume that $\tan\theta$ is bounded by an *a-priori* known polynomial $P(n) \ge 1$ in n, that is *independent* of the customers valuations. We can upper bound the total projections for any feasible allocation $\mathbf{d} = (d_1,\dots,d_n)$ of demands as follows:

$$\sum_{k\in\mathcal{N}} d^{\text{I}}_k \le C, \quad \sum_{k\in\mathcal{N}_-} -d^{\text{R}}_k \le C\tan\theta, \quad \sum_{k\in\mathcal{N}_+} d^{\text{R}}_k \le C(1+\tan\theta),$$

where $\mathcal{N}_+ \triangleq \{k \in \mathcal{N} \mid d^{\text{R}}_k \ge 0\}$ and $\mathcal{N}_- \triangleq \{k \in \mathcal{N} \mid d^{\text{R}}_k < 0\}$. Define $L \triangleq \frac{\epsilon C}{n(P(n)+1)}$, and for $d \in \mathcal{D}$, define the new rounded demand $\widehat{d}$ as follows:

$$\widehat{d} = \widehat{d}^{\text{R}} + \mathbf{i}\widehat{d}^{\text{I}} \triangleq \begin{cases} \left\lceil \frac{d^{\text{R}}}{L} \right\rceil \cdot L + \mathbf{i} \left\lceil \frac{d^{\text{I}}}{L} \right\rceil \cdot L, & \text{if } d^{\text{R}} \ge 0, \\ \left\lfloor \frac{d^{\text{R}}}{L} \right\rfloor \cdot L + \mathbf{i} \left\lceil \frac{d^{\text{I}}}{L} \right\rceil \cdot L, & \text{otherwise.} \end{cases} \quad (18)$$

Consider an optimal allocation $\mathbf{d}^* = (d^*_1,\dots,d^*_n)$ to MULTICKP$[0, \pi\text{-}\varepsilon]$. Let ξ_+ (and ξ_-), ζ_+ (and ζ_-) be the respective guessed real and imaginary absolute total projections of the rounded demands in $S^*_+ \triangleq \{k : d^{\text{R}}_k \ge 0\}$ (and $S^*_- \triangleq \{k : d^{\text{R}}_k < 0\}$). Then the possible values of $\xi_+, \xi_-, \zeta_+, \zeta_-$ are integral mutiples of L in the following ranges:

$$\xi_+ \in \mathcal{A}_+ \triangleq \left\{0, L, 2L, \dots, \left\lceil \frac{C(1+P(n))}{L} \right\rceil \cdot L \right\},$$

$$\xi_- \in \mathcal{A}_- \triangleq \left\{0, L, 2L, \dots, \left\lceil \frac{C\cdot P(n)}{L} \right\rceil \cdot L \right\},$$

$$\zeta_+, \zeta_- \in \mathcal{B} \triangleq \left\{0, L, 2L, \dots, \left\lceil \frac{C}{L} \right\rceil \cdot L \right\}.$$

Let further $\widehat{\mathcal{D}} \triangleq \{\frac{d}{L} \in \mathcal{D} :\ d^{\text{R}} \in \mathcal{A}_+ \text{ and } d^{\text{I}} \in \mathcal{B}\}$, and note that $|\widehat{\mathcal{D}}| = O(\frac{n^2 P^3(n)}{\epsilon^2})$.

We first present a $(1, 1+3\epsilon)$-approximation algorithm (MULTICKP-BIFPTAS) for MULTICKP$[0, \pi\text{-}\varepsilon]$. Let $\mathcal{N}_+ \triangleq \{k \in \mathcal{N} \mid d^{\text{R}} \ge 0\ \forall d \in D_k\}$ and $\mathcal{N}_- \triangleq \{k \in \mathcal{N} \mid d^{\text{R}} < 0\ \forall d \in D_k\}$ be the subsets of users with demand sets in the first and second quadrants respectively (recall that we restrict users' demand sets to allow such a partition).

The basic idea of Algorithm MULTICKP-BIFPTAS is to enumerate the guessed total projections on real and imaginary axes for S^*_+ and S^*_- respectively. We then solve two separate MULTI-2DKP problems (one for each quadrant) to find subsets of demands that satisfy the individual guessed total projections. But since MULTI-2DKP is generally NP-hard, we need to round the demands to get a problem that can be solved efficiently by dynamic programming. We note that the violation of the optimal solution to the rounded problem w.r.t. to the original problem is small in ϵ.

LEMMA 5.1. *For any optimal allocation $\mathbf{d}^* = (d^*_1,\dots,d^*_n)$ to* MULTICKP $[0, \pi\text{-}\varepsilon]$, *we have $\left|\sum_k \widehat{d}^*_k\right| \le (1+2\epsilon)C$.*

The next step is to solve the each rounded instance exactly. Assume an arbitrary order on $\mathcal{N} = \{1,...,n\}$. We define a 3D table, with each entry $U(k, c_1, c_2)$ being the maximum utility obtained from a subset of users $\{1,2,\dots,k\} \subseteq \mathcal{N}$, each with choosing from $\widehat{\mathcal{D}}$, that can fit exactly (i.e., satisfies the capacity constraint as an equation) within capacity

c_1 on the real axis and c_2 on the imaginary axis. This table can be filled-up by standard dynamic programming; we denote such a program by MULTI-2DKP-EXACT$[\cdot]$. For a user $k \in \mathcal{N}_-$, we redefine the valuation as $\bar{v}_k(d) = v_k(\bar{d})$, where, for $d \in \mathcal{D}$, $\bar{d}^{\mathrm{R}} = -d^{\mathrm{R}}$ and $\bar{d}^{\mathrm{I}} = d^{\mathrm{I}}$. For a set $F \subseteq \mathcal{D}$, we write $\bar{F}$ for the set $\{\bar{d} : d \in F\}$.

Algorithm 3 MULTICKP-BIFPTAS $(\{v_k, D_k\}_{k\in\mathcal{N}}, C, \epsilon)$

Require: Users' multi-minded valuations $\{v_k, D_k\}_{k\in\mathcal{N}}$; capacity C; accuracy parameter ϵ
Ensure: $(1, 1+3\epsilon)$-allocation $(\widetilde{d}_1, \ldots, \widetilde{d}_n)$ to MULTICKP$[0, \pi\text{-}\varepsilon]$
1: $(d_1, \ldots, d_n) \leftarrow (\mathbf{0}, \ldots, \mathbf{0})$
2: $\widehat{\mathcal{D}}_+ \leftarrow \{\frac{d}{L} \in \mathcal{D} : \ d^{\mathrm{R}} \in \mathcal{A}_+ \text{ and } d^{\mathrm{I}} \in \mathcal{B}\}$
3: $\widehat{\mathcal{D}}_- \leftarrow \{\frac{d}{L} \in \mathcal{D} : \ -d^{\mathrm{R}} \in \mathcal{A}_- \text{ and } d^{\mathrm{I}} \in \mathcal{B}\}$
4: **for all** $\xi_+ \in \mathcal{A}_+, \xi_- \in \mathcal{A}_-, \zeta_+, \zeta_- \in \mathcal{B}$ **do**
5:     **if** $(\xi_+ - \xi_-)^2 + (\zeta_+ + \zeta_-)^2 \le (1+2\epsilon)^2 C^2$ **then**
6:         $F_+ \leftarrow$ MULTI-2DKP-EXACT$(\{v_k, D_k\}_{k\in\mathcal{N}_+}, \frac{\xi_+}{L}, \frac{\zeta_+}{L}, \widehat{\mathcal{D}})$
7:         $F_- \leftarrow$ MULTI-2DKP-EXACT$(\{\bar{v}_k, D_k\}_{k\in\mathcal{N}_-}, \frac{\xi_-}{L}, \frac{\zeta_-}{L}, \widehat{\mathcal{D}})$
8:         $(d'_1, \ldots, d'_n) \leftarrow F_+ \cup \overline{F}_-$
9:         **if** $\sum_k v_k(d'_k) > \sum_k v_k(d_k)$ **then**
10:             $(d_1, \ldots, d_n) \leftarrow (d'_1, \ldots, d'_n)$
11:         **end if**
12:     **end if**
13: **end for**
14: **for all** $k \in \mathcal{N}_+$ **do**
15:     Choose $\widetilde{d}_k \in D_k$ s.t. $\widetilde{d}_k \preceq d_k$ and $v_k(d_k) = v_k(\widetilde{d}_k)$
16: **end for**
17: **return** $(\widetilde{d}_1, \ldots, \widetilde{d}_n)$

The following lemma states that the allocation returned by MULTICKP-BIFPTAS does not violate the capacity constraint by more than a factor of $1 + 3\epsilon$.

LEMMA 5.2. *Let* $\widetilde{\mathbf{d}}$ *be the allocation returned by* MULTICKP-BIFPTAS. *Then* $|\sum_k \widetilde{d}_k| \le (1 + 3\epsilon)C$.

THEOREM 5.3. *For any* $\epsilon > 0$, *there is a truthful for* MULTICKP$[0, \pi\text{-}\varepsilon]$, *that returns a* $(1, 1+3\epsilon)$*-approximation. The running time is polynomial in* n *and* $\frac{1}{\epsilon}$.

PROOF. We define a declaration-independent range $\mathcal{S}$ as follows. For $\xi_+ \in \mathcal{A}_+, \xi_- \in \mathcal{A}_-, \zeta_+, \zeta_- \in \mathcal{B}$, define

$$\mathcal{S}_{\xi_+,\xi_+,\zeta-+,\zeta_-} \triangleq \{\mathbf{d} = (d_1, \ldots, d_n) \in \widehat{\mathcal{D}}_+^n : \sum_{k\in\mathcal{N}_+} d_k^{\mathrm{R}} = \xi_+, \ \sum_{k\in\mathcal{N}_+} d_k^{\mathrm{I}} = \zeta_+, \ -\sum_{k\in\mathcal{N}_-} d_k^{\mathrm{R}} = \xi_-, \ \sum_{k\in\mathcal{N}_+} d_k^{\mathrm{R}} = \zeta_-\}.$$

Define further

$$\mathcal{S} \triangleq \bigcup_{(\xi_+ - \xi_-)^2 + (\zeta_+ + \zeta_-)^2 \le (1+2\epsilon)^2 C^2} \mathcal{S}_{\xi_+,\xi_+,\zeta-+,\zeta_-}.$$

Using Algorithm MULTICKP-BIFPTAS, we can optimize over $\mathcal{S}$ in time polynomial in n and $\frac{1}{\epsilon}$. Thus, it remains only to argue that the algorithm returns a $(1, 1+3\epsilon)$-approximation w.r.t. the original range $\mathcal{D}^n$. To see this, let $d_1^*, \ldots, d_n^* \in \mathcal{D}$ be the demands allocated in the optimum solution to MULTI-CKP, and $\widetilde{d}_1, \ldots, \widetilde{d}_n \in \mathcal{D}$ be the demands allocated by MULTICKP-BIFPTAS. Then by Lemma 5.1, the truncated optimal allocation $(\widehat{d}_1^*, \ldots, \widehat{d}_n^*)$ is feasible with respect to a capacity of $(1 + 2\epsilon)C$, and thus its projections will satisfy the condition in Step 5 of Algorithm 3. It follows that $v(\widetilde{\mathbf{d}}) \ge v(\widehat{\mathbf{d}}^*) \ge v(\mathbf{d}^*) = \text{OPT}$, where the second inequality follows from the way we round demands (18) and the monotonicity of the valuations. Finally, the fact that the solution returned byMULTICKP-BIFPTAS violates the capacity constraint by a factor of at most $(1 + 3\epsilon)$ follows readily from Lemma 5.2. $\square$

6. CONCLUSION

In this paper, we provided truthful mechanisms for an important variant of the knapsack problem with complex-valued demands. We gave a truthful PTAS when all demand sets of users lie in the positive quadrant, and a bi-criteria truthful FPTAS when some of the demand sets can lie in the second quadrant. In the full version of the paper, we show that these are essentially the *best possible* results in terms of approximation guarantees, assuming P≠NP.

Acknowledgment

We thank Piotr Krysta for bringing the results from [5, 9] (presented in Section 4.1) to our attention.

7. REFERENCES

[1] *National Electrical Code (NEC) NFPA 70-2005.*

[2] P. Briest, P. Krysta, and B. Vocking. Approximation techniques for utilitarian mechanism design. In *STOC'05*, pages 39–48, 2005.

[3] Chandra Chekuri and Sanjeev Khanna. A ptas for the multiple knapsack problem. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '00, pages 213–222, Philadelphia, PA, USA, 2000. Society for Industrial and Applied Mathematics.

[4] S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *ACM EC*, 2007.

[5] Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. *J. Artif. Intell. Res. (JAIR)*, 37:85–98, 2010.

[6] A. Frieze and M. Clarke. Approximation algorithm for the m-dimensional 0-1 knapsack problem. *European Journal of Operational Research*, 15:100–109, 1984.

[7] J. Grainger and W. Stevenson. *Power System Analysis.* McGraw-Hill, 1994.

[8] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems.* Springer, 2010.

[9] Piotr Krysta, Orestis Telelis, and Carmine Ventre. Mechanisms for multi-unit combinatorial auctions with a few distinct goods. In *AAMAS*, 2013.

[10] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory.* Cambridge University Press, 2007.

[11] Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. *J. Artif. Int. Res.*, 29(1):19–47, May 2007.

[12] Lan Yu and Chi-Kin Chau. Complex-demand knapsack problems and incentives in AC power systems. In *AAMAS*, 2013.