Mechanisms for Multi-Unit Combinatorial Auctions with a Few Distinct Goods *

Piotr Krysta[1] Orestis Telelis[2] Carmine Ventre[3]

[1]Deparment of Computer Science, The University of Liverpool, U.K.
p.krysta@liverpool.ac.uk

[2]Dept. of Informatics, Athens University of Economics and Business, Greece
telelis@gmail.com

[3]School of Computing, Teesside University, U.K.
c.ventre@tees.ac.uk

Abstract

We design and analyze deterministic truthful approximation mechanisms for multi-unit Combinatorial Auctions with only a *constant* number of distinct goods, each in arbitrary limited supply. Prospective buyers (bidders) have preferences over *multisets* of items, i.e. for more than one unit per distinct good. Our objective is to determine allocations of multisets that maximize the Social Welfare. Despite the recent theoretical advances on the design of truthful Combinatorial Auctions (for several distinct goods) and multi-unit auctions (for a single good), results for the combined setting are much scarser. Our main results are for *multi-minded* and *submodular* bidders. In the first setting each bidder has a positive value for being allocated one multiset from a prespecified *demand set* of alternatives. In the second setting each bidder is associated to a submodular valuation function that defines his value for the multiset he is allocated.

For multi-minded bidders we design a truthful FPTAS that *fully* optimizes the Social Welfare, while violating the supply constraints on goods within factor $(1+\epsilon)$ for any fixed $\epsilon > 0$ (i.e., the approximation applies to the constraints and not to the Social Welfare). This result is best possible, in that full optimization is impossible without violating the supply constraints. It also improves significantly upon a related result of Grandoni *et al.* [SODA 2010]. For submodular bidders we extend a general technique by Dobzinski and Nisan [JAIR, 2010] for multi-unit auctions, to the case of multiple distinct goods. We use this extension to obtain a PTAS that approximates the optimum Social Welfare within factor $(1 + \epsilon)$ for any fixed $\epsilon > 0$, without violating the supply constraints. This result is best possible as well. Our allocation algorithms are *Maximum-in-Range* and yield truthful mechanisms when paired with Vickrey-Clarke-Groves payments.

*This work has been supported by the EPSRC grants EPSRC grant EP/K01000X/1 and EP/M018113/1. The second author is supported by the research project 'DDCOD'(PE6-213), implemented within the framework of the Action "Supporting Postdoctoral Researchers" of the Operational Program "Education and Lifelong Learning" (Action's Beneficiary: General Secretariat for Research and Technology), and is co-financed by the European Social Fund (ESF) and the Greek State.

1 Introduction

In this paper we study the design and analysis of truthful *multi-unit* Combinatorial Auctions, for a constant number of distinct goods, each in limited supply. Arguably, the most widespread modern application of this general setting is the allocation of *radio spectrum licences* [14]; each such license is for the use of a specific frequency band of electromagnetic spectrum, within a certain geographic area. In the design of such "Spectrum Auctions", licenses for the same area are considered as identical units of a single good (the area), while the number of distinct geographic areas is, of course, bounded by a constant.

More formally, we consider the problem of auctioning (allocating) "in one go" multiple units of each out of a constant number of distinct goods, to prospective buyers with *private* multi-demand combinatorial valuation functions, so as to maximize the Social Welfare. A multi-demand buyer in this setting may have distinct positive values for distinct multisets of goods, i.e., for each such multiset they may demand more than one unit per good. Our aim is to devise deterministic truthful auction mechanisms, wherein every bidder finds it to his best interest to reveal his value truthfully for each multiset of items (i.e., truthful report of valuation functions is a *dominant strategy*). Additionally, we are interested in mechanisms that can compute an approximately efficient allocation in polynomial time. This problem generalizes simultaneously *Combinatorial Auctions* of multiple goods and *Multi-Unit Auctions* of a single good to the multi-unit and combinatorial setting.

Mechanism Design for Combinatorial Auctions of multiple heterogeneous goods (each in unitary supply) has received a significant attention in recent years, since the work of [12] (see [13] for the complete version), due to their various applications, especially in online trading systems over the internet. A mechanism elicits bids from interested buyers, so as to determine an assignment of bundles to them and payments in such a way, that it is to each bidder's best interest to reveal his valuation function truthfully to the mechanism. This line of research was initiated by Nisan and Ronen in their seminal paper [17]. The related problem of auctioning multiple – say s – units of a single good to multi-demand bidders has already been considered by Vickrey in his seminal paper [20]. For bidders with submodular private valuation functions, Vickrey gave an extension of his celebrated single-item Second-Price Auction mechanism, that retains truthful revelation of valuation functions as a (weakly) dominant strategy for bidders and fully optimizes the Social Welfare. The only drawback of this mechanism is that it is computationally efficient only for a few (constant number of) units, in that the allocation algorithm must process $\Theta(s)$ bids, whereas because s is an input number, it should process a number of bids bounded by a polynomial of $\log s$. Polynomial-time approximation mechanisms for multi-unit auctions were designed relatively recently [15, 16, 6, 21]. In particular, Vöcking designed and analyzed very recently a *randomized* universally truthful polynomial-time approximation scheme, for bidders with unrestricted valuation functions [21].

Results for the more general setting of multi-unit Combinatorial Auctions are much scarcer [1, 7]. It is exactly this setting we consider here, with a constant number of distinct goods, as in [7]; in particular, for a number of cases of such auctions we analyze *Maximum-in-Range* (MIR) allocation algorithms [18], that can be paired with the Vickrey-Clarke-Groves payment rule, so as to yield truthful mechanisms.

1.1 Contribution

Our main results concern multi-unit Combinatorial Auctions with a constant number of distinct goods for two broad classes of bidders, as specified by their associated valuation functions:

1. Multi-Minded Bidders: in this setting each bidder is associated with a demand set of alternative multisets (the multiple minds). Each bidder's valuation function assigns a (possibly distinct) positive value for every alternative in the demand set (and at least as much for the value of every superset of the alternative) and zero elsewhere.
2. Submodular Bidders: in this setting the value of each bidder for a particular multiset of items is given by a submodular valuation function.

For multi-minded bidders we design and analyze in Section 4 a truthful FPTAS[1], that fully optimizes the Social Welfare in polynomial time, while violating the supply constraints on the goods by a factor at most $(1+\epsilon)$, for any fixed $\epsilon > 0$. The violation of the supply constraints has a practical as well as a theoretical justification. On one hand it is conceivable that, in certain environments, a slight augmentation of supply can be economically viable, for the sake of better solutions (e.g., auctioneers with well supplied stocks can easily handle occurrences of modest overselling). On the other hand, we note that a relaxation of the supply constraints is necessary for obtaining an FPTAS, as the problem is otherwise *strongly* **NP**-hard, for $m \geq 2$ goods (please see the related discussion in Section 4). This result significantly improves on a FPTAS given in [7], which approximates the Social Welfare *and* the supply constraints within factor[2] $(1+\epsilon)$, and *only* when bidders are single-parameter (i.e. associate the same positive value with each multiset from their demand set) and *do not overbid* their demands. Technically, the FPTAS in [7] is based on the design of *monotone* algorithms [10, 2] and, as such, it needs the assumption of no-overbidding on the demands (cf. discussion therein).

In Section 5 we revisit the general technique introduced by Dobzinski and Nisan in [6] for multi-unit auction Mechanism Design and generalize it to multiple distinct goods, each in limited supply. We discuss how this generalization yields a truthful PTAS immediately for multi-minded bidders, that does *not violate* the supply constraints and approximates the Social Welfare within factor $(1+\epsilon)$, for any fixed $\epsilon > 0$. Subsequently, we use the technique to design a truthful PTAS for bidders with submodular valuation functions, assuming that the values (bids) are accessed through value queries by the algorithm. Prior to this result, no time-efficient deterministic truthful mechanism was known for submodular bidders, even for a single good. Interestingly, the direct extension of the technique of [6] for multiple distinct goods does not yield a factor 2 approximation mechanism for general valuation functions accessed by value queries, as was the case for a single good in [6]; we show that an appropriate extension of a more dedicated treatment of this case from [6] yields a 2-approximation (Section 6).

The assumption of a constant number $m = O(1)$ of distinct goods is important, for otherwise our problems become Combinatorial Auctions, thus, hard to approximate within less than $O(\sqrt{m})$ [13]. Regarding the generalization of the Dobzinski-Nisan technique, existence of a FPTAS for multi-minded bidders and one good is excluded by a result from [6], unless

[1] *Fully Polynomial Time Approximation Scheme.*

[2] In the context of Social Welfare *maximization*, by "approximation within factor $\rho \geq 1$" (or, equivalently, "ρ-approximation", for $\rho \geq 1$) we mean recovering at least a fraction ρ^{-1} of the welfare of an optimum allocation. We only switch temporarily to using $\rho < 1$ in Section 5, for technical convenience.

$\mathbf{P} = \mathbf{NP}$. Finally, as shown in [6, 19] regarding general valuation functions, no *deterministic* MIR algorithm achieves better than 2-approximation – even for a single good – with communication complexity $o(s)$, where s is the supply of this good. These lower bounds imply that our results are best possible.

2 Related Work

Mechanism Design for multi-unit auctions was initiated already by the celebrated work of Vickrey [20], where he extended his famous mechanism for the case of multiple units, when bidders have symmetric submodular valuation functions [11]. This mechanism is however not computationally efficient with respect to the number of available units, as we already discussed. It requires that bidders place a *marginal bid* per additional unit they wish to receive and the allocation algorithm processes all these marginal bids. The design of multi-unit mechanisms with polynomially bounded running time in $\log s$, the number of units, was first considered in [16]. In this work, Mu'alem and Nisan designed a truthful 2-approximation mechanism for single-minded bidders, when their single preference on the number of units is known to the mechanism (i.e., is not part of the bidders' private information). The only private datum is the non-zero value of the bidders for their single preference on the number of units that they wish to be allocated. Briest, Krysta and Vöecking gave later a FPTAS in [2] for the more general case where the bidders' single preference is also private information.

Dobzinski and Nisan in [5, 6] analyzed a general scheme for designing MIR polynomial-time truthful approximation mechanisms. This resulted in a PTAS for the case of k-minded bidders, a 2-approximation for general valuation functions that are accessed (by the allocation algorithm) through value queries, and a $\frac{3}{4}$-approximation for symmetric subadditive valuation functions. Dobzinski and Dughmi gave a *truthful in expectation* FPTAS for multi-minded bidders in [4]. Very recently, Vöcking gave a universally truthful randomized PTAS for general valuation functions accessed by value queries [21] (in contrast, all of our mechanisms are deterministic). For the multi-unit combinatorial setting (i.e., with more than one distinct goods) the known results concern mainly bidders that have demands for a single unit from each good (see e.g. [13, 2]). In contrast, we consider a constant number of goods, but multi-demand bidders. Bartal, Gonen and Nisan [1] proved approximation and competitiveness results for multi-unit Combinatorial Auctions with multi-demand bidders, where the bidders' demands on numbers of units are upper and lower bounded. The approximation guarantees depend on these bounds.

The study of a constant number of goods, each in arbitrary limited supply, was initiated by Grandoni *et al.* in [7]. The authors utilized methods from multi-objective optimization (approximate Pareto curves and Langrangian relaxation) to design and analyze truthful polynomial-time approximation schemes for a variety of settings. In particular, they devised truthful FPTASs that approximate both the objective function (Social Welfare or Cost) of multi-capacitated versions of problems within factor $(1 + \epsilon)$, while violating the capacity constraints by a factor $(1 + \epsilon)$ (capacity here corresponds to limited supply of each out of a few distinct goods). Problems considered in [7] include *multi-unit auctions*, *minimum spanning tree*, *shortest path*, *maximum* (perfect) *matching* and *matroid intersection*; for a subclass of these problems a truthful PTAS is also analyzed, that does not violate any of the capacity constraints.

3 Definitions

Let $[m] = \{1, \ldots, m\}$ be a set of m goods, where m is assumed to be a fixed constant. There are $s_\ell \in \mathbb{N}$ units (copies) of good $\ell \in [m]$ available. A *multiset* of goods is denoted by a vector $\mathbf{x} = (x(1), x(2), \ldots, x(m))$, where $x(\ell)$ is the number of units of good $\ell \in [m]$, $\ell = 1, \ldots, m$. The set of all multisets is denoted by $\mathcal{U} = \times_{\ell=1}^{m}\{0, 1, \ldots, s_\ell\}$. Let $\mathcal{N} = [n] = \{1, \ldots, n\}$ be the set of n agents (prospective buyers/bidders). Every bidder $i \in [n]$ has a private valuation function

$$v_i : \mathcal{U} \mapsto \mathbb{R}^+,$$

so that $v_i(\mathbf{x})$ for any $\mathbf{x} \in \mathcal{U}$ denotes the maximum monetary amount that i is willing to pay for $\mathbf{x} \in \mathcal{U}$, referred to as his value for $\mathbf{x}$. The valuation functions are normalized, i.e. $v_i(0, \ldots, 0) = 0$ and assumed to be *monotone non-decreasing*: for any two multisets $\mathbf{x} \leq \mathbf{y}$ where "$\leq$" holds component-wise, we assume $v_i(\mathbf{x}) \leq v_i(\mathbf{y})$.

A mechanism is an *allocation* method $\mathcal{A}$ and a *payment* rule p. The allocation method elicits the bidders' bids $\mathbf{b} = (b_1, \ldots, b_n)$ for their valuation functions and determines an allocation $\mathbf{x}(\mathbf{b}) = \mathcal{A}(\mathbf{b})$, so that $\mathbf{x}_i(\mathbf{b}) \in \mathcal{U}$ is the multiset of goods allocated to bidder i. The *payment rule* determines a vector $\mathbf{p}(\mathbf{b})$, where $p_i(\mathbf{b})$ is the payment of bidder i. Every bidder i bids so as to maximize his quasi-linear utility, defined as:

$$u_i(\mathbf{b}) = v_i(\mathbf{x}_i(\mathbf{b})) - p_i(\mathbf{b}).$$

We study *truthful* mechanisms $(\mathcal{A}, p)$ wherein each bidder i maximizes his utility by reporting his valuation function truthfully, i.e., by bidding $b_i = v_i$, independently of the other bidders' reports, $\mathbf{b}_{-i}$:

Definition 1 *A mechanism $(\mathcal{A}, p)$ is truthful if, for every bidder i and bidding profile $\mathbf{b}_{-i}$, it satisfies $u_i(v_i, \mathbf{b}_{-i}) \geq u_i(v_i', \mathbf{b}_{-i})$, for every v_i'.*

We will drop notation $\mathbf{b}$ and will use only $\mathbf{x}$ for an allocation output by $\mathcal{A}$, since we analyze truthful mechanisms. Our objective function is the *Social Welfare* of the output allocation $\mathbf{x}$, which we aim at maximizing:

$$SW(\mathbf{x}) = \sum_i v_i(\mathbf{x}_i),$$

Our mechanisms use VCG payments and *Maximum-in-Range* (MiR) [18] allocation algorithms:

Definition 2 *[18] An algorithm choosing its output from the set A of all possible allocations is* MiR*, if it fully optimizes the Social Welfare over a subset $R \subseteq A$ of allocations.*

Nisan and Ronen [18] identified MiR allocation algorithms as the sole device that, along with VCG payments, yields truthful mechanisms for Combinatorial Auctions.

4 Multi-Minded Bidders

In this section we consider *multi-minded* bidders; every such bidder $i \in [n]$ is associated with a collection of multisets $\mathcal{D}_i \subseteq \mathcal{U}$, referred to as his *demand-set*. We assume that he values each $\mathbf{d} \in \mathcal{D}_i$, an amount $v_i(\mathbf{d}) > 0$. For every other multiset $\mathbf{d} \in \mathcal{U} \setminus \mathcal{D}_i$ we define:

$$v_i(\mathbf{d}) = \begin{cases} \max_{\mathbf{e} \in \mathcal{D}_i} \Big\{ v_i(\mathbf{e}) \Big| \mathbf{e} \leq \mathbf{d} \Big\} & \text{if such } \mathbf{e} \in \mathcal{D}_i \text{ exists} \\ 0 & \text{otherwise.} \end{cases}$$

Naturally, $v_i(\emptyset) = 0$. Consequently, in this setting, the valuation function of a bidder i can be compactly expressed as the collection $(v_i(\mathbf{d}), \mathbf{d})_{\mathbf{d} \in \mathcal{D}_i}$. As in related literature, we assume therefore that an algorithm expects in input bids of this form, rather than (an oracle representing) the entire valuation function. We say that a bidder i is a *winner* of the auction, if he is assigned exactly one of his alternatives from $\mathcal{D}_i$ (or a superset of one of these alternatives); this corresponds to the `XOR`-bidding language in Combinatorial Auctions [11].

We design a $(1, 1+\varepsilon, 1+\varepsilon, \ldots, 1+\varepsilon)$-approximation FPTAS, that *maximizes the Social Welfare* and may violate the supply constraints on goods by a factor at most $(1+\varepsilon)$, for any fixed $\epsilon > 0$. This will be the allocation algorithm of our mechanism. We will show that it is MIR, thus can be paired with VCG payments to yield a truthful mechanism.

In light of turning it into a truthful mechanism, we use notation of actual valuation functions in its definition. The approach is reminiscent of the one that yields the FPTAS for the well-known one-dimensional knapsack problem. Fix any $\varepsilon > 0$. First, remove all the alternatives $\mathbf{d} \in \mathcal{D}_i$ for any $i \in [n]$ and $\ell = 1, \ldots, m$, such that $d(\ell) > s_\ell$ (if all alternatives of some bidder i are removed, remove i). Henceforth, we use the same notation, $\mathcal{U}$, $[n]$, $\mathcal{D}_i$, etc., for the remaining alternatives and bidders. The demands of the alternatives $\mathbf{d} \in \mathcal{D}_i$ are rounded so that for any good $\ell \in [m]$, we have $d'(\ell) = \lfloor \frac{n \cdot d(\ell)}{\varepsilon s_\ell} \rfloor$, for any agent $i \in [n]$. The new supply for good ℓ becomes now $s'_\ell = \lceil \frac{n}{\varepsilon} \rceil$.

We define the dynamic programming table $\mathcal{V}(i, Y_1, \ldots Y_m)$ for $i = 1, \ldots, n$ and $Y_\ell \in \{0, 1, 2, \ldots, s'_\ell\}$ for any $\ell \in [m]$. The cell $\mathcal{V}(i, Y_1, \ldots, Y_m)$ stores the maximum welfare, i.e., $\sum_j v_j(\mathbf{x}_j)$, of an allocation $\mathbf{x}$, whose rounded version $\mathbf{x}' = (\lfloor \frac{n \cdot \mathbf{x}_j(\ell)}{\varepsilon s_\ell} \rfloor)_{j,\ell}$ uses alternatives of the bidders in $\{1, 2, \ldots, i\}$, and has total demand w.r.t. good $\ell = 1, \ldots, m$ which is precisely Y_ℓ, i.e., $\sum_i x'_i(\ell) = Y_\ell$.

To compute the entries of table $\mathcal{V}$, we observe that, the problem $\mathcal{V}(1, Y_1, \ldots Y_m)$ for any collection of Y_ℓ's such that: $(Y_1, \ldots, Y_m) \in \{0, 1, \ldots, \lceil \frac{n}{\varepsilon} \rceil\}^m$, is easy to solve. For each such entry $\mathcal{V}(1, Y_1, \ldots Y_m)$ we check if bidder 1 has an alternative $\mathbf{d} \in \mathcal{D}_1$ such that $d'(\ell) = Y_\ell$, for all $\ell \in [m]$. If yes, let $\mathbf{d}$ be an alternative of maximum valuation; we assign $\mathcal{V}(1, Y_1, \ldots, Y_m) = v_1(\mathbf{d})$ and build an auxiliary table $A[1, Y_1, \ldots Y_m]$ which we set in this case to $\{(1, \mathbf{d})\}$. Otherwise, if bidder 1 does not have any such alternative, we assign $\mathcal{V}(1, Y_1, \ldots Y_m) = 0$ and $A[1, Y_1, \ldots Y_m] = \{(1, \emptyset)\}$. To define $\mathcal{V}(i+1, Y_1, \ldots, Y_m)$, consider bidder $i+1$ and his alternatives $\mathbf{d} = (d(1), \ldots, d(m)) \in \mathcal{D}_{i+1}$; let now

$$\nu_{i+1} = \max_{\mathbf{d} \in \mathcal{D}_{i+1}} \Big\{ v_{i+1}(\mathbf{d}) + \mathcal{V}\Big(i, Y_1 - d'(1), ..., Y_m - d'(m)\Big) \Big| \mathbf{d}' \leq \mathbf{Y} \Big\} \tag{1}$$

where, for all i, we set $\mathcal{V}(i, Y_1, \ldots Y_m) = -\infty$ and, accordingly, $A[i, Y_1, \ldots Y_m] = \emptyset$, if at least one of the Y_i's is negative. Consequently:

$$\mathcal{V}(i+1, Y_1, \ldots, Y_m) = \max\Big\{\nu_{i+1}, \mathcal{V}(i, Y_1, \ldots Y_m)\Big\}.$$

Accordingly, if $\nu_{i+1} \leq \mathcal{V}(i, Y_1, \ldots Y_m)$, we set:

$$A[i+1, Y_1, \ldots, Y_m] = A[i, Y_1, \ldots, Y_m] \cup \{(i+1, \emptyset)\},$$

otherwise:

$$A[i+1, Y_1, \ldots, Y_m] = A[i, Y_1 - d'(1), \ldots, Y_m - d'(m)] \cup \{(i+1, \mathbf{d})\},$$

where $\mathbf{d}$ is an alternative in $\mathcal{D}_{i+1}$ maximizing (1). Finally, we inspect all the solutions from entries $\mathcal{V}(n, Y_1, \ldots, Y_m)$ for all vectors $(Y_1, \ldots, Y_m) \in \{0, 1, \ldots, \lceil \frac{n}{\varepsilon} \rceil\}^m$, take one which maximizes the Social Welfare and output the solution given by the corresponding entry of the A table.

The size of table $\mathcal{V}$ is $n(\lceil \frac{n}{\varepsilon} \rceil + 1)^m$ and we need time roughly $O(\max_i |\mathcal{D}_i| + m)$ to compute one entry of the table, so the overall time of the algorithm leads to an FPTAS. The optimality with respect to the sum of values is easy to verify; for any feasible solution $\mathbf{x}$ to the original problem, we have for each good $\ell = 1, \ldots, m$: $\sum_i x_i(\ell) \leq s_\ell$, or, equivalently, $\sum_i \frac{x_i(\ell) \cdot n}{\varepsilon \cdot s_\ell} \leq \frac{n}{\varepsilon}$ thus $\sum_i \left\lfloor \frac{x_i(\ell) \cdot n}{\varepsilon \cdot s_\ell} \right\rfloor \leq s'_\ell = \left\lceil \frac{n}{\varepsilon} \right\rceil$. So, the feasibility of $\mathbf{x}$ is preserved in the rounded problem. Because, the dynamic programming algorithm will inspect all feasible solutions to the rounded problem, also an optimum solution to the original problem will be detected.

We argue that the supply constraints s_ℓ, $\ell = 1, \ldots, m$, are violated by at most a factor of $1 + 2\varepsilon$. Fix any good $\ell \in \{1, \ldots, m\}$ and let $\mathbf{x}$ be the output allocation. Now we have $\sum_i \left\lfloor \frac{n \cdot x_i(\ell)}{\varepsilon \cdot s_\ell} \right\rfloor \leq s'_\ell = \lceil \frac{n}{\varepsilon} \rceil$, and because:

$$\begin{aligned}\sum_i \frac{n \cdot x_i(\ell)}{\varepsilon \cdot s_\ell} &\leq \sum_i \left\lfloor \frac{n \cdot x_i(\ell)}{\varepsilon \cdot s_\ell} \right\rfloor + |\{i | \mathbf{x}_i \in \mathcal{D}_i\}| \\ &\leq \left\lceil \frac{n}{\varepsilon} \right\rceil + n \leq \frac{n}{\varepsilon} + 1 + n,\end{aligned}$$

so finally $\sum_i x_i(\ell) \leq (1+2\varepsilon) s_\ell$. Note that the algorithm is *exact*, in that it grants every bidder a multiset from his demand set (or none). Assuming $m = O(1)$ is essential for the result, even in presence of the supply constraints' relaxation. A proof of this claim is given at the end of this section. The truthfulness of the FPTAS, denoted by $\mathcal{A}$ below, follows from the fact that it optimizes over a fixed range of solutions.

Theorem 1 *There exists a truthful FPTAS for the multi-unit combinatorial auction problem with a fixed number of goods, when bidders have private multi-minded valuation functions, defined, for each bidder, over a (private) collection of multisets of goods. The approximation guarantee is $(1, 1+\varepsilon, \ldots, 1+\varepsilon)$, i.e., only the supplies of the goods may be violated.*

Proof. To prove the theorem we show that $\mathcal{A}$ is MiR with range $R = \{\mathbf{x} | \exists \mathbf{b} : \mathcal{A}(\mathbf{b}) = \mathbf{x}\}$. That is, for any allocation $\mathbf{x} \in R$ and bid vector $\mathbf{b}$, we show $SW(\mathcal{A}(\mathbf{b}), \mathbf{b}) \geq SW(\mathbf{x}, \mathbf{b})$, where for a bid vector $\mathbf{b} = \left((b_i(\mathbf{d}), \mathbf{d})_{\mathbf{d} \in \mathcal{D}_i}\right)_{i \in \mathcal{N}}$ and an allocation $\mathbf{x} \in R$, we let $SW(\mathbf{x}, \mathbf{b})$ be the Social Welfare of allocation $\mathbf{x}$, evaluated according to the bid vector $\mathbf{b}$, i.e., $SW(\mathbf{x}, \mathbf{b}) = \sum_i b_i(\mathbf{x})$.

Fix allocation $\mathbf{x}$ and bid vector $\mathbf{b} = \left((b_i(\mathbf{d}), \mathbf{d})_{\mathbf{d} \in \mathcal{D}_i}\right)_{i \in \mathcal{N}}$; by definition of range, there exists a bid vector $\bar{\mathbf{b}}$, with $\bar{\mathbf{b}} = \left((\bar{b}_i(\bar{\mathbf{d}}), \bar{\mathbf{d}})_{\bar{\mathbf{d}} \in \bar{\mathcal{D}}_i}\right)_{i \in \mathcal{N}}$ such that $\mathcal{A}(\bar{\mathbf{b}}) = \mathbf{x}$. Let $x_i(\ell)$, for bidder i and $\ell = 1, \ldots, m$, denote the variable indicating how many copies of item ℓ, the allocation $\mathbf{x}$ grants to bidder i. Note that because $\mathbf{x} = \mathcal{A}(\bar{\mathbf{b}})$ and $\mathcal{A}$ grants only demanded alternatives (by its exactness), there exists a demand $\bar{d}_i \in \bar{\mathcal{D}}_i \cup \{\emptyset\}$ such that, for $\ell = 1, \ldots, m$,

$x_i(\ell) = \bar{d}_i(\ell)$. Since $\mathbf{x}$ is output of $\mathcal{A}$, by definition of $\mathcal{A}$ we have that for any $\ell = 1, \ldots, m$, $\sum_i \left\lfloor \frac{n \cdot \bar{d}_i(\ell)}{\varepsilon \cdot s_\ell} \right\rfloor \leq \left\lceil \frac{n}{\varepsilon} \right\rceil$.

Now let C be the set of bidders such that $\mathbf{b} = (\mathbf{b}_C, \mathbf{b}_{-C})$ and $\bar{\mathbf{b}} = (\bar{\mathbf{b}}_C, \mathbf{b}_{-C})$, that is, $\mathbf{b}$ and $\bar{\mathbf{b}}$ only differ in the bids of bidders in the set C. For all bidders $i \in C$ we assume that their true valuation function is b_i. Any such bidder i evaluates the alternative $\mathbf{x}_i = \bar{\mathbf{d}}_i$ granted to him by allocation $\mathbf{x}$ as some $\mathbf{e}_i \in \mathcal{D}_i \cup \{\emptyset\}$. That is, $v_i(\bar{\mathbf{d}}_i) = v_i(\mathbf{e}_i)$. Assume, for the sake of contradiction, that $SW(\mathbf{x}, \mathbf{b}) > SW(\mathcal{A}(\mathbf{b}), \mathbf{b})$, i.e.:

$$\sum_{i \in C} b_i(\mathbf{e}_i) + \sum_{j \notin C} b_j(\mathbf{x}) > \sum_{i \in C} b_i(\mathcal{A}(\mathbf{b})) + \sum_{j \notin C} b_j(\mathcal{A}(\mathbf{b})). \tag{2}$$

Since $\bar{d}_i(\ell) \geq e_i(\ell)$ for $\ell = 1, \ldots, m$ and $i \in C$, then by setting $\mathbf{e}_i = \bar{\mathbf{d}}_i$ for $i \notin C$, we obtain:

$$\sum_i \left\lfloor \frac{n \cdot e_i(\ell)}{\varepsilon \cdot s_\ell} \right\rfloor \leq \sum_i \left\lfloor \frac{n \cdot \bar{d}_i(\ell)}{\varepsilon \cdot s_\ell} \right\rfloor \leq \left\lceil \frac{n}{\varepsilon} \right\rceil,$$

for $\ell = 1, \ldots, m$. Then the solution which grants to bidder i the alternative $\mathbf{e}_i \in \mathcal{D}_i \cup \{\emptyset\}$ is considered by algorithm $\mathcal{A}$ on input $\mathbf{b}$. This solution has Social Welfare $SW(\mathbf{x}, \mathbf{b})$ and therefore (2) is in contradiction with the definition of $\mathcal{A}$. $\square$

A related result from [2] is a truthful FPTAS for a single good in limited (not violated) supply; this cannot be generalized for our setting of more than one supply constraints.

A note on hardness Note that this problem is strongly **NP**-hard when we do not allow to violate supply constraints and $m \geq 2$ [3]. Also the assumption that m is a fixed constant is necessary. Otherwise the problem is equivalent to multi-unit Combinatorial Auctions and is hard to approximate within $m^{1/2-\epsilon}$ for any $\epsilon > 0$ [13]. This claim is true even if we allow for solutions to violate the capacities. In particular:

Proposition 1 *In a multi-unit combinatorial auction with m distinct goods, it is* **NP***-hard to approximate the Social Welfare within factor better than $m^{-1/2}$, even if we allow a multiplicative $(1+\varepsilon)$-relaxation of the supply constraints, for any $\varepsilon < 1$.*

Proof. The argument is as follows: it is known that it is hard to approximate the maximum independent set problem in a graph $G = (V, E)$ within a factor $m^{1/2-\epsilon}$ for any $\epsilon > 0$, where $|E| = m$ [8]. By using a reduction from [13], we reduce this problem to our problem by having the set of goods $[m] = E$ and the set of single-minded bidders V; each bidder's $u \in V$ set contains all edges adjacent to u in graph G and each bidder's valuation for his set is 1. Now if we allow to violate the capacity (of 1) of each good by a factor of $1+\varepsilon$, where $\varepsilon < 1$, then a feasible solution to the relaxed problem is an independent set in graph G. Thus the relaxed problem is equivalent to the maximum independent set problem in G. $\square$

4.1 Multi-Dimensional Knapsack

We discuss an application of our FPTAS, in relation to the multi-dimensional Knapsack problem [3]. Suppose we are given a multi-dimensional knapsack with a constant number of compartments m and compartment $\ell = 1, 2, \ldots, m$ has capacity s_ℓ. An object $\mathbf{d}_i \in \mathcal{U}$ has dimensions $(d_i(0), d_i(1), \ldots, d_i(m))$ and corresponds to a single agent i, where $d_i(0) = v_i$.

This multiple knapsack problem corresponds to a single-parameter version of the problem we treated above. (It is worth mentioning that for this version with unique valuation for all the demand sets, our FPTAS from Section 4 can be shown to be monotone [10, 2] when one carefully fixes some tie-breaking rule.) Now, we can further generalize this problem by adding an additional constant number of cost functions and budgets, say $d_i(\ell)$, s_ℓ, $\ell = m+1, m+2, \ldots, m'$, for each $\mathbf{d}_i \in \mathcal{U}$, where $m' > m$ is a fixed constant, and we may assume any "capacity" or "covering" constraints in $\{\geq, \leq\}$ for each $\ell = m+1, \ldots, m'$. For such a generalized scenario we can follow the approach similar to our approach in Section 4 and obtain a truthful FPTAS which gives a $(1, 1+\varepsilon, \ldots, 1+\varepsilon)$-approximation (violation of constraints is needed for the reason mentioned above). The issue of monotonicity of the algorithm is similar to the case of Combinatorial Auctions.

5 The Generalized Dobzinski-Nisan Method

We discuss here a direct generalization of a method designed by Dobzinski and Nisan in [6], for truthful single-good multi-unit auction mechanisms. We will use the method's generalization for multiple goods in the next subsection, to obtain a truthful PTAS for bidders with submodular valuation functions (over multisets). Let $\mathcal{A}$ be a polynomial-time MiR allocation algorithm for $t = O(1)$ bidders and s_ℓ units from each good $\ell = 1, \ldots, m$, with complexity $\mathcal{A}(t, \mathbf{s})$, $\mathbf{s} = (s_1, \ldots, s_m)$, and approximation ratio $\alpha \leq 1$. Then, one may use algorithm $\mathcal{A}$ as a routine within the procedure of Figure 1, to obtain a polynomial-time MiR algorithm for n bidders, with approximation ratio $(\alpha - \frac{m}{t+1})$.

The procedure executes algorithm $\mathcal{A}$ on every subset of at most t bidders for any $t = O(1)$ and for every combination of certain pre-specified quantities of the goods. For each output allocation it considers the rest of the bidders and allocates optimally to them an integral number of bundles from each good. The main result shown in [6] for a single good can be also shown for m goods:

Theorem 2 *Let $\mathcal{A}$ be a Maximum-in-Range algorithm, with complexity $\mathcal{A}(t, (s_1, \ldots, s_m))$ for t bidders and at most s_ℓ units from each good $\ell = 1, \ldots, m$. There exists a range of allocations, $\mathcal{R}$, such that the Dobzinski-Nisan Method runs in poly*$(\log s_1, \ldots, \log s_m, n, \mathcal{A}(t, (s_1, \ldots, s_m)))$ *time, for every $t = O(1)$, and outputs an allocation with value at least $(\alpha - \frac{m}{t+1})$ times the optimum Social Welfare.*

The proof is a direct extension of the proof given in [6] for a single good. Consider the MiR algorithm $\mathcal{A}$, to be used within the Dobzinski-Nisan method; it executes in polynomial time for $t = O(1)$ bidders and m distinct goods, each in limited supply s_ℓ, $\ell \in [m]$. Let $\mathcal{R}_{\mathcal{A},t,m}$ denote the *range* of this algorithm. It is straightforward to verify that the method outputs allocations that are "$(\mathcal{R}, t, \chi_1, \ldots, \chi_m)$-round", where $\mathcal{R} = \mathcal{R}_{\mathcal{A},t,m}$, given the following definition (generalized appropriately from [6]), for "round" allocations:

Definition 3 *An allocation is $(\mathcal{R}, t, \chi_1, \ldots, \chi_m)$-round if:*

- *$\mathcal{R}$ is a set of allocations and in each $\mathbf{x} \in \mathcal{R}$ at most t bidders are allocated non-empty bundles. The bidders are allocated together up to $s_\ell - \chi_\ell$ units from each good $\ell = 1, \ldots, m$.*

1. **for** $\ell = 1, \ldots, m$ **do:**
   (a) **define** $u_\ell := (1 + \frac{1}{2n})$
   (b) **define** $L_\ell := \left\{0, 1, \lfloor u_\ell \rfloor, \lfloor u_\ell^2 \rfloor, \ldots, u_\ell^{\lfloor \log_{u_\ell} s_\ell \rfloor}, s_\ell\right\}$
2. **for** every subset $T \subseteq \mathcal{N}$ of bidders, $|T| \leq t$, **do:**
   1. **for** every $(\chi_1, \ldots, \chi_m) \in \left(\times_{\ell=1}^m L_\ell\right)$ **do:**
      1 Run $\mathcal{A}$ with $s_\ell - \chi_\ell$ units from each good $\ell \in [m]$ and bidders in T.
      2 Split the remaining χ_ℓ units from each good $\ell \in [m]$ into $\leq 2n^2$ bundles (per good), each of $\max\left\{\lfloor \frac{\chi_\ell}{2n^2} \rfloor, 1\right\}$ units.
      3 Find the optimal allocation of the equi-sized bundles among bidders $\mathcal{N} \setminus T$.
3. **Return** the best allocation found.

Figure 1: The Dobzinski-Nisan Method for multiple goods.

- *There exists a set T of bidders, $|T| \leq t$, such that they are all allocated according to some allocation in $\mathcal{R}$.*
- *Each bidder $i \in \mathcal{N} \setminus T$ receives an exact multiple of* $\max\left\{\lfloor \frac{\chi_\ell}{2n^2} \rfloor, 1\right\}$ *units from good ℓ and:* $\sum_{i \in \mathcal{N} \setminus T} x_i(\ell) \leq n \cdot \max\left\{\lfloor \frac{\chi_\ell}{2n^2} \rfloor, 1\right\}$ *, for $\ell = 1, \ldots, m$*

Then the range of the method is the subset of all allocations that are $(\mathcal{R}_{\mathcal{A},t,m}, \kappa, \chi_1, \ldots, \chi_m)$-round, such that $(\chi_1, \ldots, \chi_m) \in (\times_{\ell=1}^m L_\ell)$ and $\kappa \leq t$. Call the range of the method $\mathcal{R}_c$. We show that it approximates the socially optimum allocation within factor $(\alpha - \frac{m}{t+1})$.

Lemma 1 *Let $\mathbf{x}^* = (\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ be a socially optimum allocation. There exists an allocation $\mathbf{x} \in \mathcal{R}_c$ with $SW(\mathbf{x}) \geq (\alpha - \frac{m}{t+1}) \cdot SW(\mathbf{x}^*)$.*

Proof. Without loss of generality (because of monotonicity of valuation functions), assume that all units of all goods are allocated in $\mathbf{x}^*$ and that $v_1(\mathbf{x}_1^*) \geq v_2(\mathbf{x}_2^*) \geq \cdots \geq v_n(\mathbf{x}_n^*)$. For every good $\ell = 1, \ldots, m$ choose the largest value $\chi_\ell \in L_\ell$ so that $s_\ell - \chi_\ell \geq \sum_{i=1}^t v_i(\mathbf{x}_i^*)$. When executed on bidders on the subset of bidders $T = \{1, \ldots, t\}$ with $s_\ell - \chi_\ell$ units from good $\ell = 1, \ldots, m$, algorithm $\mathcal{A}$ outputs an allocation $(\mathbf{x}_1, \ldots, \mathbf{x}_t)$ such that $\sum_{i=1}^t v_i(\mathbf{x}_i) \geq \alpha \sum_{i=1}^t v_i(\mathbf{x}_i^*)$.

Now consider for each good $\ell = 1, \ldots, m$ the a bidder $j_\ell \in \{t+1, \ldots, n\}$ with the maximum number of units in $\mathbf{x}^*$ from this good. Define $r_\ell = \sum_{i=t+1}^n x_i(\ell)$. Then $x_{j_\ell}^*(\ell) \geq \frac{r_\ell}{n}$. By definition of r_ℓ and χ_ℓ for each good ℓ, we have $r_\ell \geq \chi_\ell$. Also, because χ_ℓ was chosen to have the largest possible value, it must be $\chi_\ell \geq \frac{r_\ell}{u_\ell} \geq r_\ell - \frac{r_\ell}{2n}$. For every bidder $i \geq t+1$ with $i \neq j_\ell$ for $\ell = 1, \ldots, m$, we round up his allocation with respect to good ℓ to a multiple of $\max\left\{\lfloor \frac{\chi_\ell}{2n^2} \rfloor, 1\right\}$. The extra units for each good ℓ we take from bidders j_ℓ who may not obtain any unit of the good. Observe that we may need to add at most $n \cdot \frac{\chi_\ell}{2n^2} \leq \frac{\chi_\ell}{2n}$ extra units from each good ℓ, that we take from bidder j_ℓ, who has at least $\frac{r_\ell}{n} \geq \frac{\chi_\ell}{n}$ units.

Thus, for all bidders except for j_ℓ, $\ell = 1, \ldots, m$ we increased the units of goods they obtain. Because $j_\ell \geq t+1$ and $v_1(\mathbf{x}_1) \geq \cdots \geq v_n(\mathbf{x}_n)$, we have $v_{j_\ell}(\mathbf{x}_{j_\ell}^*) \leq \frac{1}{t+1} \sum_{i=1}^t v_i(\mathbf{x}_i^*)$

and $v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_i^*)$ for $i \neq j_\ell$, $\ell = 1, \ldots m$. Then:

$$
\begin{aligned}
SW(\mathbf{x}) &= \sum_i v_i(\mathbf{x}_i) \geq \alpha \sum_{i=1}^{t} v_i(\mathbf{x}_i^*) + \sum_{i \geq t+1} v_i(\mathbf{x}_i) \\
&\geq \alpha \sum_{i=1}^{t} v_i(\mathbf{x}_i^*) + \sum_{i \geq t+1} v_i(\mathbf{x}_i^*) - \sum_{\ell=1}^{m} v_i(\mathbf{x}_{j_\ell}^*) \\
&= \left(\alpha - \frac{m}{t+1}\right) \sum_{i=1}^{t} v_i(\mathbf{x}_i^*) + \sum_{i \geq t+1} v_i(\mathbf{x}_i^*) \geq \left(\alpha - \frac{m}{t+1}\right) SW(\mathbf{x}^*)
\end{aligned}
$$

which concludes the proof. $\square$

The lemma completes the proof of Theorem 2.

Let us explain how to find an optimal allocation of single-good bundles (i.e. bundles of identical units) for each good to bidders in $\mathcal{N} \setminus T$, in step **2.1.3** of the algorithm (Figure 1). We use a dynamic programming. By re-indexing the bidders appropriately, assume that $T = \{n-t+1, \ldots, n\}$, thus $\mathcal{N} \setminus T = \{1, \ldots, n-t\}$. For every $i = 1, \ldots, n-t$ and for every $\mathbf{q} = (q_1, \ldots, q_m) \in \left(\times_{i=1}^{m} [2n^2]\right)$, define $\mathcal{V}(i, \mathbf{q}) = \mathcal{V}(i, (q_1, \ldots, q_m))$ to be the maximum value of welfare that can be obtained by allocating at most q_ℓ equi-sized bundles (of units) from each good $\ell = 1, \ldots, m$ to bidders $1, \ldots, i$. Each entry $\mathcal{V}(i, \mathbf{q})$ of the dynamic programming table can be computed using:

$$
\mathcal{V}(i, \mathbf{q}) = \max_{\mathbf{q}' \leq \mathbf{q}} \Big(v_i(\mathbf{q}' \cdot \mathbf{b}) + \mathcal{V}(i-1, \mathbf{q} - \mathbf{q}') \Big),
$$

where $\mathbf{q}' \leq \mathbf{q}$ is taken component-wise; i.e. maximization occurs over all vectors $\mathbf{q}'$ such that $q'(\ell) \leq q(\ell)$ for each $\ell = 1, \ldots, m$. Note that $\mathbf{q}' \cdot \mathbf{b}$ denotes the inner product.

Simple application: k-Minded Bidders The Dobzinski-Nisan method for multiple distinct goods can be applied immediately in the setting of (multi-parameter) k-minded bidders, to yield a PTAS while respecting fully the supply constraints of the goods. For $m = O(1)$ goods and for any constant number of t bidders the optimum assignment can be found exhaustively in polynomial time in $\log s_\ell$, $\ell = 1, \ldots, s$, and m. In particular, there are exactly $O(k^t)$ cases to be examined exhaustively, so that the optimum is found. Plugging this algorithm in the procedure of Figure 1, yields a PTAS that, contrary to the developments of the previous section, approximates the optimum Social Welfare within factor $(1+\epsilon)$ and respects the supply constraints.

5.1 Submodular Valuation Functions

We consider submodular valuation functions over multisets in $\mathcal{U}$, as defined in [9]:

Definition 4 *For any $\ell = 1, \ldots, m$ let $\mathbf{e}_\ell$ be the unary vector with $e_\ell(\ell) = 1$ and $e_\ell(j) = 0$, for $j \neq \ell$. Let $\mathbf{x}$ and $\mathbf{y}$ denote two multisets from $\mathcal{U}$, so that $\mathbf{x} \leq \mathbf{y}$, where "$\leq$" holds component-wise. Then, a non-decreasing function $v : \mathcal{U} \mapsto \mathbb{R}^+$ is submodular if $v(\mathbf{x} + \mathbf{e}_\ell) - v(\mathbf{x}) \geq v(\mathbf{y} + \mathbf{e}_\ell) - v(\mathbf{y})$.*

We assume that these valuation functions, being exponentially large to describe, are accessed by the algorithm through *value queries*; i.e., for any value that the algorithm needs to process as input, it asks for it from the corresponding bidder, for the particular corresponding multiset.

We will design the MiR approximation algorithm $\mathcal{A}$, needed by the method. The range we consider for this setting is an extension of the one considered by Dobzinski and Nisan in [6]. For any $\epsilon > 0$, define $\delta = 1 + \epsilon$; we will be assigning to bidders unit bundles of each good $\ell \in [k]$, that have cardinality equal to an integral power of δ. For every good $\ell \in [k]$, one of the n bidders (possibly a different bidder per good) will always obtain the remaining units of the specific good. We show that optimization over this range provides a good approximation of the unrestricted optimum Social Welfare; also, optimizing over this range yields an FPTAS for a constant number n of bidders. This, used within the generalized Dobzinksi-Nisan method will yield a PTAS for any number of bidders.

Lemma 2 *An optimum assignment within the defined range, recovers at least a factor* $\left(\frac{2-\epsilon}{2+2\epsilon}\right)^m$ *of the socially optimum welfare.*

Proof. Let $\mathbf{x}^* = (\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ denote the welfare maximizing assignment. We will round iteratively – for a particular good $\ell \in [m]$ in each iteration – the assignment of units to each bidder in $\mathbf{x}^*$, to an integral power of δ. Let $\mathbf{x}^{[\ell]}$ be the assignment after rounding with respect to the ℓ-th good. The final assignment $\mathbf{x} \equiv \mathbf{x}^{[m]}$ will approximate the welfare of $\mathbf{x}^{[0]} \equiv \mathbf{x}^*$.

In the beginning of ℓ-th iteration we process the assignment $\mathbf{x}^{[\ell-1]}$, by rounding the assignment of unit bundles of good ℓ. Assume w.l.o.g. that $x_1^{[\ell-1]}(\ell) \geq x_2^{[\ell-1]}(\ell) \geq \cdots \geq x_n^{[\ell-1]}(\ell)$. Also w.l.o.g., we assume that every bidder except for bidder 1 receives an integral power of δ units of good ℓ; bidder 1 receives the remaining units. Let the set of bidders be partitioned as $\mathcal{N} = \mathcal{O} \cup \mathcal{E}$ where $\mathcal{O}$ contains the odd indices of bidders and $\mathcal{E}$ the even ones. We will consider two cases:

$$\sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right) \geq \sum_{i\in\mathcal{E}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right) \qquad \text{and} \qquad \sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right) < \sum_{i\in\mathcal{E}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right). \tag{3}$$

For the first case, for every $i \in \mathcal{O} \setminus \{1\}$ we will round $x_i^{[\ell-1]}(\ell)$ up to the closest integral power of δ, while obtaining the extra units to do so by rounding $x_{i-1}^{[\ell-1]}(\ell)$, $i-1 \in \mathcal{E}$ down to the nearest appropriately chosen integral power of δ. We obtain $x_i^{[\ell]}(\ell) \leq \delta \cdot x_i^{[\ell-1]}(\ell)$ and:

$$\hat{x}_{i-1}^{[\ell-1]}(\ell) = x_{i-1}^{[\ell-1]}(\ell) - (\delta-1)x_i^{[\ell-1]}(\ell) \geq x_{i-1}^{[\ell-1]}(\ell) - (\delta-1)x_{i-1}^{[\ell-1]}(\ell)$$

thus, $\hat{x}_{i-1}^{[\ell-1]}(\ell) \geq (2-\delta)x_{i-1}^{[\ell-1]}(\ell)$. To ensure that for bidder $i-1$ we obtain an integral power of δ, we may need to divide $\hat{x}_{i-1}^{[\ell-1]}(\ell)$ at most by δ, thus: $x_{i-1}^{[\ell]}(\ell) \geq \frac{1}{\delta}\hat{x}_{i-1}^{[\ell-1]}(\ell) = \frac{2-\delta}{\delta}x_{i-1}^{[\ell-1]}(\ell)$.

The welfare of the emerging assignment $\mathbf{x}^{[\ell]}$ is:

$$\begin{aligned}
SW\left(\mathbf{x}^{[\ell]}\right) &= \sum_{i\in\mathcal{N}} v_i\left(\mathbf{x}_i^{[\ell]}\right) = \sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell]}\right) + \sum_{i\in\mathcal{E}} v_i\left(\mathbf{x}_i^{[\ell]}\right)\\
&\geq \sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right) + \frac{2-\delta}{\delta}\sum_{i\in\mathcal{E}} v_i\left(\mathbf{x}_i^{[\ell]}\right)\\
&= \sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right) + \frac{2-\delta}{\delta}\left(SW\left(\mathbf{x}^{[\ell-1]}\right) - \sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right)\right)\\
&= \frac{2\delta-2}{\delta}\sum_{i\in\mathcal{O}} v_i\left(x_i^{[\ell-1]}\right) + \frac{2-\delta}{\delta}SW\left(\mathbf{x}^{[\ell-1]}\right)\\
&\geq \frac{\delta-1}{\delta}SW\left(\mathbf{x}^{[\ell-1]}\right) + \frac{2-\delta}{\delta}SW\left(\mathbf{x}^{[\ell-1]}\right) = \frac{1}{1+\epsilon}SW\left(\mathbf{x}^{[\ell-1]}\right)
\end{aligned}$$

The second line follows by submodularity; for any $\ell \in [m]$, we have $x_{i-1}^{[\ell]}(\ell) \geq \frac{2-\delta}{\delta}x_{i-1}^{[\ell-1]}(\ell)$, so $v_{i-1}\left(\mathbf{x}_{i-1}^{[\ell]}\right) \geq \frac{2-\delta}{\delta}v_i\left(\mathbf{x}_{i-1}^{[\ell-1]}\right)$.

Consider now the second case in (3), where $\sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right) < \sum_{i\in\mathcal{E}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right)$. For $i \in \mathcal{E}\setminus\{2\}$ we round up $x_i^{[\ell-1]}(\ell)$ to the closest integral power of δ; the extra units for this we will obtain from $i-1 \in \mathcal{O}$, by rounding $x_i^{[\ell-1]}(\ell)$ down to an appropriately chosen closest integral power of δ. $x_2^{[\ell-1]}(\ell)$ will be rounded down to closest integral power of δ (contrary to the rest of $x_i^{[\ell-1]}(\ell)$, $i \in \mathcal{E}$), i.e., $x_2^{[\ell]}(\ell) \geq \frac{1}{\delta}x_2^{[\ell-1]}(\ell)$. For $i \in \mathcal{E}\setminus\{2\}$ it will be $x_i^{[\ell]}(\ell) \leq \delta \cdot x_i^{[\ell-1]}(\ell)$ and then we take:

$$x_{i-1}^{[\ell]}(\ell) \geq \frac{1}{\delta}\left(x_{i-1}^{[\ell-1]}(\ell) - (\delta-1)x_i^{[\ell-1]}(\ell)\right) \geq \frac{2-\delta}{\delta}x_{i-1}^{[\ell-1]}(\ell) \tag{4}$$

Then, for the Social Welfare of $\mathbf{x}^{[\ell]}$ we have:

$$\begin{aligned}
SW\left(\mathbf{x}^{[\ell]}\right) &= \sum_{i\in\mathcal{N}} v_i\left(\mathbf{x}_i^{[\ell]}\right) = \sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell]}\right) + \sum_{i\in\mathcal{E}} v_i\left(\mathbf{x}_i^{[\ell]}\right)\\
&\geq \frac{2-\delta}{\delta}\sum_{i\in\mathcal{O}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right) + \frac{1}{\delta}v_2\left(\mathbf{x}_2^{[\ell-1]}\right) + \sum_{i\in\mathcal{E}\setminus\{2\}} v_i\left(\mathbf{x}_i^{[\ell]}\right)\\
&= \frac{2-\delta}{\delta}\left(SW\left(\mathbf{x}^{[\ell-1]}\right) - \sum_{i\in\mathcal{E}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right)\right) + \frac{1}{\delta}v_2\left(\mathbf{x}_2^{[\ell-1]}\right) + \sum_{i\in\mathcal{E}\setminus\{2\}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right)\\
&= \frac{2\delta-2}{\delta}\sum_{i\in\mathcal{E}\setminus\{2\}} v_i\left(\mathbf{x}_i^{[\ell-1]}\right) + \frac{\delta-1}{\delta}v_2\left(\mathbf{x}_2^{[\ell-1]}\right) + \frac{2-\delta}{\delta}SW\left(\mathbf{x}^{[\ell-1]}\right)\\
&\geq \frac{\delta-1}{2\delta}SW\left(\mathbf{x}^{[\ell-1]}\right) + \frac{2-\delta}{\delta}SW\left(\mathbf{x}^{[\ell-1]}\right) = \frac{2-\epsilon}{2+2\epsilon}SW\left(\mathbf{x}^{[\ell-1]}\right)
\end{aligned}$$

The second line of this derivation is again due to submodularity: the factors on the sum over odd-indexed bidders and on $v_2(\mathbf{x}_2^{[\ell-1]})$ follow by (4) and because $x_2^{[\ell]}(\ell) \geq \frac{1}{\delta}x_2^{[\ell-1]}(\ell)$ Thus, for any $\epsilon > 0$, there is an assignment within the described range that approximates the optimum Social Welfare within factor $\left(\frac{2-\epsilon}{2+2\epsilon}\right)^p \cdot \left(\frac{1}{1+\epsilon}\right)^q$, for some integers p, q, such that $p+q=m$ The result follows by $\frac{1}{1+\epsilon} \geq \frac{2-\epsilon}{2+2\epsilon}$. $\square$

We obtain the following (intermediate) result:

Theorem 3 *For multi-unit combinatorial auctions with n submodular bidders, $n = O(1)$, and $m = O(1)$ distinct goods, each good $\ell \in [m]$ available in an arbitrary supply, there exists a truthful deterministic FPTAS for any $\epsilon \leq 1$, that approximates the optimum Social Welfare within factor $(1 + \epsilon)$.*

Proof. For any fixed $\epsilon > 0$ we can search the specified range exhaustively in polynomial time; to find the allocation with maximum Social Welfare, we have to try $O(\log_\delta s_\ell)$ cases for each of $n - 1$ bidders, given a fixed bidder for assigning the remaining units. Thus the time required for trying all possible bundle assignments of a specific good ℓ and for all possible choices of a "remainders" bidder is $O\left(n(\log_\delta s_\ell)^{n-1}\right)$. Because for every fixed allocation of a specific good we need to try all possible allocations for the remaining $m - 1$ goods, the overall complexity is in total $O\left(n^m(\log_\delta \max_\ell s_\ell)^{(n-1)m}\right)$, which is polynomially bounded for constant m and n. Also notice that, for $\epsilon \leq 1$ we obtain a FPTAS, because:

$$\log_\delta \max_\ell s_\ell = (\log_2(1 + \epsilon))^{-1} \cdot (\log_2 \max_\ell s_\ell)$$

and $\log_2^{-1}(1 + \epsilon) \leq \epsilon^{-1}$. □

Using Theorem 3 within the general Dobzinski-Nisan method, we obtain:

Corollary 1 *There exists a truthful PTAS for multi-unit combinatorial auctions with constant number of distinct goods and submodular valuation functions.*

6 General Valuation Functions

Interestingly, the direct generalization of the Dobzinski-Nisan method for a constant number of multiple goods, does not immediately yield, for general valuation functions, a result comparable to the one shown in [6] for a single good; for $m = 1$ a truthful 2-approximation mechanism was obtained (and this factor was shown to be optimal). When $m = 1$, the relevant MIR algorithm $\mathcal{A}$ involved in Theorem 2 solves optimally the case of $t = 1$ bidder, by allocating all units of all goods to him. The monotonicity of the valuation functions guarantees that this allocation is optimal for $t = 1$ bidder. The factor 2 approximation follows. For $m > 1$ goods however, Theorem 2 appears to require a different algorithm $\mathcal{A}$ (for, possibly, $t > 1$ bidders), to yield a comparable result. Instead, 2-approximation for the case of general valuation functions accessed by value queries can be obtained, by simple modification of the direct approach that was given in [6], for general valuation functions. The modification is small, but we include it here for completeness.

We describe from scratch a MIR allocation algorithm. The algorithm splits for every good the number of units into n^2 equi-sized bundles of size $b_\ell = \lfloor \frac{s_\ell}{n^2} \rfloor$; it also creates a single extra bundle containing the remaining units r_ℓ, so that $n^2 \cdot b_\ell + r_\ell = s_\ell$. The algorithm allocates optimally *whole bundles of units* from each good to the n bidders.

First we show that this range approximates by a factor 2 the optimum Social Welfare. Let $\mathbf{x}^* = (\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ denote the socially optimum allocation. We start with $\mathbf{x}^*$ and produce an allocation $\mathbf{x}$ in the range within which the algorithm optimizes, that approximates $SW(\mathbf{x}^*)$ within factor 2. Assume w.l.o.g. that all items are allocated in $\mathbf{x}^*$ (by the monotonicity of valuation functions) and, for each good $\ell = 1, \dots, m$, let $j_\ell = \arg\max_i x_i(\ell)$. Then $x^*_{j_\ell}(\ell) \geq \frac{s_\ell}{n}$. Let $L = \{j_1, \dots, j_\ell\}$. We consider two cases here.

$$\text{Either:} \qquad \sum_{\ell=1}^{m} v_{j_\ell}(\mathbf{x}^*_{j_\ell}) \geq \sum_{i \notin L} v_i(\mathbf{x}^*_i), \qquad \text{or:} \qquad \sum_{\ell=1}^{m} v_{j_\ell}(\mathbf{x}^*_{j_\ell}) < \sum_{i \notin L} v_i(\mathbf{x}^*_i).$$

The summation of all $v_{j_\ell}(\mathbf{x}^*_{j_\ell})$ makes the difference of this proof from the one in [6].

In the first case simply allocating all unit bundles to bidders in L yields a $\frac{1}{2}$ approximation of $SW(\mathbf{x}^*)$; this allocation is examined by the MiR algorithm. In the second case we round up – separately for each good ℓ – the allocation of bidders $i \notin L$ to the nearest multiple of b_ℓ. The units needed for this purpose we find for each good ℓ from the corresponding bidder j_ℓ who may not obtain any unit in $\mathbf{x}$. This is possible because we add at most $n \cdot \frac{s_\ell}{n^2} = \frac{s_\ell}{n} \leq x^*_{j_\ell}(\ell)$ units in total by this rounding. This way we make up an allocation $\mathbf{x}$ that gives all unit bundles for each good to bidders in $\mathcal{N} \setminus L$; this allocation is also examined by the MiR algorithm. Thus, there exists a solution within the range that approximates $SW(\mathbf{x}^*)$ within factor 2.

To complete our analysis, we show how to compute a MiR allocation for the described range, using dynamic programming. Let $\mathbf{r} = (r_1, \dots, r_m)$ denote the vector of amounts that correspond to bundles of "remainders" per good as described above. Given $L \subseteq 2^{\{1,\dots,m\}}$ we denote by $\mathbf{r}[L]$ the projection of $\mathbf{r}$ on indices in L; the remaining coordinates are set to 0. Let $\mathbf{b} = (b_1, \dots, b_m)$. For any subset $L \in 2^{\{1,\dots,m\}}$, define $\mathcal{V}^L(i, \mathbf{q})$, $\mathbf{q} = (q_1, \dots, q_m)$ as the maximum welfare achievable when allocating at most q_ℓ unit-bundles for each good $\ell = 1, \dots, m$ among bidders $1, \dots, i$ *and* the "remainders" bundle for each of the goods $\ell \in L$. We compute each $\mathcal{V}^L(i, \mathbf{q})$ as follows:

$$\mathcal{V}^L(i, \mathbf{q}) = \max_{L' \subseteq L} \max_{q'_1 \leq q_1, \dots, q'_m \leq q_m} \left\{ v_i\left(\mathbf{q} \cdot \mathbf{b} + \mathbf{r}[L']\right) + \mathcal{V}^{L \setminus L'}(i-1, \mathbf{q} - \mathbf{q}') \right\}$$

Because $m = O(1)$, the entries of the dynamic programming table can be computed in polynomial time. Thus:

Theorem 4 *There exists a truthful polynomial-time mechanism for multi-unit Combinatorial Auctions with a constant number of distinct goods and general valuation functions that, using value queries, approximates the welfare of a socially optimum assignment within factor* 2.

7 Conclusions

In this paper we analyzed deterministic mechanisms for multi-unit Combinatorial Auctions with a constant number of distinct goods, each in limited supply. We analyzed in particular Maximum-in-Range allocation algorithms [18] for optimizing the Social Welfare in this multi-unit combinatorial setting that, paired with VCG payments, yield truthful auctions. Our main results include **(i)** a truthful FPTAS for *multi-minded* bidders, that approximates the supply constraints within factor $(1+\epsilon)$ and optimizes the Social Welfare; **(ii)** a deterministic truthful PTAS for submodular bidders, that approximates the Social Welfare within factor $(1+\epsilon)$ without violating the supply constraints. For achieving **(ii)**, we used a direct generalization of a single-good multi-unit allocation method proposed by Dobzinski and Nisan in [6]. Finally, we showed how to treat unrestricted valuation functions in our setting, by adjusting appropriately an analysis from [6]. All of the discussed developments are best possible in terms of time-efficient approximation, as follows by relevant hardness results.

Acknowledgements. The authors would like to thank Fabrizio Grandoni and Stefano Leonardi for very useful discussions in the early stage of this work.

References

[1] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *ACM TARK*, pages 72–87, 2003.

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

[3] C. Chekuri and S. Khanna. A PTAS for the multiple knapsack problem. In *SODA*, pages 213–222, 2000.

[4] S. Dobzinski and S. Dughmi. On the Power of Randomization in Algorithmic Mechanism Design. In *IEEE FOCS*, pages 505–514, 2009.

[5] S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *ACM EC*, pages 346–351, 2007.

[6] S. Dobzinski and N. Nisan. Mechanisms for Multi-Unit Auctions. *Journal of Artificial Intelligence Research*, 37:85–98, 2010.

[7] F. Grandoni, P. Krysta, S. Leonardi, and C. Ventre. Utilitarian Mechanism Design for Multi-Objective Optimization. In *SODA*, pages 573–584, 2010.

[8] J. Hastad. Clique is Hard to Approximate Within $n1 - \epsilon$. In *IEEE FOCS*, pages 627–636, 1996.

[9] M. Kapralov, I. Post, and J. Vondrák. Online and stochastic variants of welfare maximization. Technical report, CoRR abs/1204.1025, 2012.

[10] B. Lehmann, D. J. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. In *ACM EC*, pages 18–28, 2001.

[11] B. Lehmann, D. J. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.

[12] D. J. Lehmann, L. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *ACM EC*, pages 96–102, 1999.

[13] D. J. Lehmann, L. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[14] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[15] A. Mu'alem and N. Nisan. Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. In *AAAI*, pages 379–384, 2002.

[16] A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008.

[17] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *STOC*, pages 129–140, 1999.

[18] N. Nisan and A. Ronen. Computationally Feasible VCG Mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007.

[19] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006.

[20] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, 1961.

[21] B. Vöcking. A universally-truthful approximation scheme for multi-unit auctions. In *SODA*, pages 846–855, 2012.