Journal of Artificial Intelligence Research 53 (2015) 721-744 Submitted 9/14; published 8/15

Mechanisms for Multi-unit Combinatorial Auctions with a Few Distinct Goods

Piotr Krysta P.KRYSTA@LIVERPOOL.AC.UK
Deparment of Computer Science
University of Liverpool, United Kingdom

Orestis Telelis TELELIS@GMAIL.COM
Department of Digital Systems
University of Piraeus, Greece

Carmine Ventre C.VENTRE@TEES.AC.UK
School of Computing
Teesside University, United Kingdom

Abstract

We design and analyze deterministic truthful approximation mechanisms for multi-unit Combinatorial Auctions involving 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. 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. For submodular bidders, we 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 *Maximal-in-Range* and yield truthful mechanisms, when paired with Vickrey-Clarke-Groves payments.

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* (Milgrom, 2004); 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 settings respectively.

Since the work of Lehmann, O'Callaghan, and Shoham (2002), Mechanism Design for Combinatorial Auctions of multiple heterogeneous goods (each in unitary supply) has received significant attention in recent years (Holzman, Kfir-Dahav, Monderer, & Tennenholtz, 2004; Lehmann, Lehmann, & Nisan, 2006; Dobzinski, Nisan, & Schapira, 2010; Lavi & Swamy, 2011), 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, that introduces algorithmic efficiency considerations in the design of truthful mechanisms, was initialized by the work of Nisan and Ronen (2001).

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 (Vickrey, 1961). 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 in at least as many steps, whereas because s is an input number, it should require a number of steps bounded by a polynomial in $\log s$. Several other drawbacks of the generalized Vickrey-Clarke-Groves (truthful) auction mechanism have been identified by Ausubel and Milgrom (2010). Polynomial-time approximation mechanisms for multi-unit auctions were designed relatively recently (Mu'alem & Nisan, 2002, 2008; Dobzinski & Nisan, 2010; Vöcking, 2012; Nisan, 2014). In particular, Nisan (2014) devised a deterministic, polynomial time auction mechanism, for the multi-unit setting with submodular bidders first considered by Vickrey (1961). Vöcking designed and analyzed recently a *randomized* universally truthful polynomial-time approximation scheme, for bidders with unrestricted valuation functions (Vöcking, 2012).

Results for the more general setting of multi-unit Combinatorial Auctions are relatively scarcer (Bartal, Gonen, & Nisan, 2003; Grandoni, Krysta, Leonardi, & Ventre, 2014; Lavi & Swamy, 2011). It is exactly this setting we consider here, with a constant number of distinct goods, similarly to the setting considered by Grandoni et al. (2014); in particular, for a number of cases of such auctions we analyze *Maximal-in-Range* (MiR) allocation algorithms (Nisan & Ronen, 2007), 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 upon an Fptas by Grandoni et al. (2014), which approximates the Social Welfare *and* the supply constraints within factor[2] $(1+\epsilon)$, *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 of Grandoni et al. (2014) is based on the design of *monotone* algorithms (Lehmann et al., 2002; Briest, Krysta, & Vöcking, 2011) and it requires a no-overbidding assumption on the demands (cf. discussion therein).

In Section 5 we revisit the general technique introduced by Dobzinski and Nisan (2010), for multi-unit auction Mechanism Design, and generalize it for the setting of 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 when $m \geq 2$.[3] Although the technique of Dobzinski and Nisan facilitated the development of a factor 2 approximation mechanism for bidders with general valuation functions in the single-good multi-unit setting, its direct extension for the setting of multiple distinct goods does not appear to work (for general valuation functions). We

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 switch temporarily to using $\rho \leq 1$ in Section 5, for technical convenience.
3. Nisan (2014) devised an optimal polynomial time such auction for $m = 1$, i.e., a single good.

show, however, that an appropriate extension of a more dedicated treatment of this case by Dobzinski and Nisan yields a constant $(m+1)$-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 in polynomial time within less than $O(\sqrt{m})$ (Lehmann et al., 2002) for multi-minded bidders and within less than $\frac{e}{e-1}$ for submodular bidders (Khot, Lipton, Markakis, & Mehta, 2008; Mirrokni, Schapira, & Vondrák, 2008). Moreover, recent results of Daniely, Schapira, and Shahaf (2015) imply that, for unrestricted m, our techniques cannot yield *truthful* polynomial-time mechanisms with approximation factors less than $O(m)$ and $O(\sqrt{m})$, respectively. Regarding the generalization of the Dobzinski-Nisan technique, existence of an Fptas for multi-minded bidders and a single good is excluded, unless $\mathbf{P} = \mathbf{NP}$ (Dobzinski & Nisan, 2010). These lower bounds imply that our results are best possible. Finally, as shown by Nisan and Segal (2006) and Dobzinski and Nisan (2010), regarding general valuation functions, no *deterministic* MiR algorithm achieves better than 2-approximation for a single good – with communication complexity $o(s)$, where s is the supply of this good. Closing the gap between this lower bound and our upper bound of $(m+1)$ for a constant number m of multiple distinct goods, remains an open problem.

2. Related Work

Mechanism Design for multi-unit auctions was initiated already by the celebrated work of Vickrey (1961), where he extended his famous mechanism for the case of multiple units, when bidders have symmetric submodular valuation functions (Lehmann et al., 2006). 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. Very recently, Nisan (2014) exhibited a polynomial-time truthful mechanism for this case. The design of multi-unit mechanisms with polynomially bounded running time in $\log s$, s denoting the number of units, was first considered by Mu'alem and Nisan (2008). In this work, Mu'alem and Nisan designed and analyzed a truthful polynomial-time 2-approximation mechanism for a multi-unit combinatorial setting, involving multiple distinct goods, each in limited supply, and single-minded bidders. Subsequently, Archer, Papadimitriou, Talwar, and Tardos (2003) improved upon this approximation ratio for a similar setting, but the developed mechanism was based on randomized rounding and was truthful only *in expectation*. More recently, Briest et al. (2011) designed and analysed an Fptas, for single-minded bidders in the multi-unit combinatorial setting.

Dobzinski and Nisan (2010) analyzed a general scheme for designing MiR polynomial-time truthful approximation mechanisms, for single-good multi-unit auctions. This resulted in a Ptas for the case of multi-minded bidders, a 2-approximation for general valuation functions that are accessed (by the allocation algorithm) through value queries, and a $\frac{4}{3}$-approximation for symmetric subadditive valuation functions. Moreover, the authors applied their scheme to a class of *piecewise linear* (multi-unit) valuation functions over the number of units of a single good, to obtain a truthful Ptas mechanism. A special case of this class of valuation functions had been earlier studied by Kothari, Parkes, and Suri (2005); the authors designed an Fptas mechanism that was, however, only *approximately* *truthful*. Dobzinski and Dughmi (2013) gave a randomized *truthful in expectation* Fptas for multi-minded bidders. Relatively recently, Vöcking (2012) gave a universally truthful randomized Ptas for general valuation functions accessed by value queries (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 (Lehmann et al., 2002; Briest et al., 2011; Blumrosen & Nisan, 2007). In contrast, we consider a constant number of goods, but multi-demand bidders. Bartal et al. (2003) proved approximation and competitiveness results for truthful multi-unit Combinatorial Auctions with multi-demand bidders, where the bidders' demands on numbers of units are upper and lower bounded. The derived approximation guarantees depend on these bounds. Lavi and Swamy (2011) improved upon these approximation guarantees, by devising randomized mechanisms that are truthful in expectation.

The study of a constant number of goods, each in arbitrary limited supply, was initiated by Grandoni et al. (2014). 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 FptasEs 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 by Grandoni et al. 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, which does not violate any of the capacity constraints.

Of particular interest to the practice of Combinatorial Auctions (also in the multi-unit case, where each good is available in a limited supply of identical copies), is the efficient and (near-)optimal resolution of the *Winner Determination* problem (Lehmann, Müller, & Sandholm, 2010). Given as input a set of bids in a prespecified format (formally termed *language*), for the items on sale, the Winner Determination problem prescribes the determination of a feasible allocation of the items to the bidders, so that the sum of their bids corresponding to their received allocation is maximized. Thus, the Winner Determination problem implicitly prescribes determination of winning bidders and their receiving allocation, so that the revenue collected by the corresponding bids is maximized. Notice that, in comparison to our work, truthful report of the bidders' valuation functions is not a concern in this setting. A significant volume of research has concerned the study of approximation algorithms and derivation of hardness results (see, e.g., Lehmann et al., 2010), as much as the development of global optimization techniques (Sandholm, 2010). Kelly (2004) studies Multi-unit Combinatorial Auctions with only a few distinct goods, for determining the allocation of computational resources. In particular, he devises an optimal algorithm for the Winner Determination problem, in a low-dimensional setting as ours.

Finally let us mention that the work of Bikhchandani, de Vries, Schummer, and Vohra (2011) investigates multi-unit Combinatorial Auction premises very similar to ours, without a restriction on the number of distinct goods. Instead, the authors devise an *ascending price* auction for selling subsets of goods that constitute *bases of a matroid*, or *polymatroid*, in case of multi-unit demand by the bidders and limited supply for each distinct good. Their auction is truthful and runs in polynomial or pseudo–polynomial time, respectively. It accesses the bidders' combinatorial valuation functions through *Demand Queries*; the bidders are presented with prices on the goods and announce the subset they are willing to pay for. In comparison, all of our mechanisms use *Value Queries*, where the mechanism asks for the value of each bidder for a particular set of items. Value queries are a weaker *"device"* in that they can be simulated by (but cannot generally simulate) demand queries (Blumrosen & Nisan, 2007).

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 $[n] = \{1, \ldots, n\}$ 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 $v_i(\mathbf{x}) \leq v_i(\mathbf{y})$. That is, in auction theory terms, we assume *free disposal* (i.e., enlarging the set or increasing the number of items in an allocation never decreases the value incurred to any bidder).

A *mechanism* consists of an *allocation* method (algorithm), $\mathcal{A}$, and a *payment* rule, $\mathbf{p}$. The allocation method $\mathcal{A}$ elicits bids $\mathbf{b} = (\,\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n\,)$ from the bidders that, presumably, describe their valuation functions and outputs an allocation $\mathcal{A}(\mathbf{b}) = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$, where $\mathbf{x}_i \in \mathcal{U}$ is the multiset of goods allocated to bidder i. For the purposes of our discussion in this section, we deliberately ignore the fact that the bidders' valuation functions may not have a succinct representation that will facilitate their efficient communication to the allocation algorithm; recall that the bidders' valuation functions are – generally – defined over $\mathcal{U} = \times_{\ell=1}^{m}\{0, 1, \ldots, s_\ell\}$. When they do not have a succinct representation indeed, the allocation algorithms discussed in this paper access the bidders' valuation functions iteratively, through *polynomially many Value Queries*; that is, the algorithm in each iteration asks every bidder for a bid on a specific multiset of items.

The *payment rule* determines a vector $\mathbf{p}(\mathbf{b}) = (\,p_1(\mathbf{b}), p_2(\mathbf{b}), \ldots, p_n(\mathbf{b})\,)$, 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:

$$u_i(\mathbf{b}) = v_i(\,\mathcal{A}(\mathbf{b})\,) - p_i(\mathbf{b}),$$

where, by an assumption of *no externalities*, i.e., that the value of any bidder for $\mathcal{A}(\mathbf{b})$ depends *only* on his own individual allocation and not at all on the others', we obtain $v_i(\,\mathcal{A}(\mathbf{b})\,) = v_i(\mathbf{x}_i)$.

We study *truthful* mechanisms $(\mathcal{A}, \mathbf{p})$ 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, $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$:

Definition 1 *A mechanism* $(\mathcal{A}, \mathbf{p})$ *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*.*

Under this definition, the profile $\mathbf{b} = \mathbf{v}$ is a *dominant strategy* equilibrium. Our objective is to design and analyze *truthful* mechanisms, $(\mathcal{A}, \mathbf{p})$ that render truthful reporting of the bidders' valuation a dominant strategy equilibrium, wherein, the *Social Welfare* of the resulting allocation, $SW(\mathcal{A}(\mathbf{b})) = SW(\mathcal{A}(\mathbf{v}))$ is (approximately) optimized. The social welfare of an allocation, $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ is defined as:

$$SW(\mathbf{X}) = \sum_{i=1}^{n} v_i(\mathbf{x}_i),$$

In the sequel we will use simply $\mathbf{X}$, for an allocation output by $\mathcal{A}$, without a specific reference to $\mathbf{b}$, since we analyze truthful mechanisms, that dictate $\mathbf{b} = \mathbf{v}$.

The only well understood general method for the design of truthful mechanisms is the Vickrey-Clarke-Groves (VCG) auction mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973), a generalization of Vickrey's Single-Item 2nd Price and Multi-unit Auctions (Vickrey, 1961). Deployment of the VCG auction, however, requires utilization of an allocation algorithm, $\mathcal{A}$, that outputs a welfare-maximizing allocation for the underlying setting; it rarely constitutes a computationally efficient alternative for combinatorial settings, as the underlying optimization problem is **NP**-hard.

As the problems that we consider in our work are indeed **NP**-hard, our mechanisms use *Maximal-in-Range* (MIR) allocation algorithms (Nisan & Ronen, 2007), that maximize the social welfare only approximately.

Definition 2 *(Nisan & Ronen, 2007) 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.*

Note that the subset R, also called a *range*, is defined independently from the bidders' declarations. Nisan and Ronen (2007) identified MIR allocation algorithms as the sole device that, along with VCG payments, yields truthful mechanisms for Combinatorial Auctions. In particular, given any MIR allocation algorithm, $\mathcal{A}$, using this algorithm for computing the output allocation *and* for computing payments in the manner of the VCG payments scheme, suffices to obtain a truthful mechanism. In particular, given any such MIR allocation algorithm, the payment for each bidder i is computed as follows:

$$p_i(\mathbf{b}) = \sum_{i' \neq i} v_{i'}(\mathcal{A}(\mathbf{b}_{-i})) - \sum_{i' \neq i} v_{i'}(\mathcal{A}(\mathbf{b}))$$

Notice how this payment scheme coincides with the VCG payment scheme, once we use the optimal allocation algorithm in place of $\mathcal{A}$. The starting point for the work of Nisan and Ronen (2007) is the pair of observations that: *(i)* the VCG mechanism itself *requires* full optimization of the social welfare of the underlying setting, which is an **NP**-hard problem in most interesting settings *(ii)* VCG-based mechanisms (wherein a polynomial-time allocation algorithm outputs welfare-suboptimal allocations) are *not* necessarily truthful.

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 each $i \in [n]$ values each multiset $\mathbf{d} = (d(1), \ldots, d(m)) \in \mathcal{D}_i$ by an amount $v_i(\mathbf{d}) > 0$. For every other multiset $\mathbf{e} \in \mathcal{U} \setminus \mathcal{D}_i$ we define:

$$v_i(\mathbf{e}) = \begin{cases} \max\limits_{\mathbf{d} \in \mathcal{D}_i} \left\{ v_i(\mathbf{d}) \,\middle|\, \mathbf{d} \leq \mathbf{e} \right\} & \text{if such } \mathbf{d} \in \mathcal{D}_i \text{ exists} \\ 0 & \text{otherwise.} \end{cases}$$

Naturally, $v_i(\mathbf{0}) = 0$, where $\mathbf{0} = (0, \ldots, 0)$. 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 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 (Lehmann et al., 2006).

We design a 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. After analyzing its performance with respect to the welfare optimality of the allocation that it outputs and the bounded violation of the supply constraints, we will prove that it is a MiR algorithm, thus can be paired with VCG payments, to yield a truthful mechanism. At a high level, the algorithm is reminiscent of the one that yields the Fptas for the well-known one-dimensional knapsack problem (see e.g., Vazirani, 2003, ch. 8). It proceeds as follows. For any chosen fixed $\varepsilon > 0$, it first discards any alternatives of bidders from their demand sets, that cannot be satisfied, given the supply constraints. These alternatives are multisets that already exceed the supply of at least one good. Subsequently, the quantities of goods in the multisets remaining within the bidders' demand sets are appropriately rounded; the supply is adjusted as well. We thus obtain a *rounded* instance. Then, we search for a welfare maximizing allocation of the rounded instance, by usage of dynamic programming. This allocation is shown to be optimal for the initial instance, as well, and feasible, modulo a violation of the initial supply constraints within a factor of at most $(1 + \epsilon)$. In light of turning this algorithm into a truthful mechanism, we use notation of actual valuation functions in its definition and analysis below.

Fix any constant $\varepsilon > 0$. First, for any $i \in [n]$, remove all the alternatives $\mathbf{d} \in \mathcal{D}_i$ such that $d(\ell) > s_\ell$ for *any* $\ell = 1, \ldots, m$ (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 $\mathbf{d} \in \mathcal{D}_i$ of *each* $i \in [n]$ are rounded as follows. For every $i \in [n]$ and for every $\mathbf{d} \in \mathcal{D}_i$, we produce a multiset $\mathbf{d}' = (d'(1), \ldots, d'(m))$ so that, for each distinct good $\ell \in [m]$, we have $d'(\ell) = \lfloor \frac{n \cdot d(\ell)}{\varepsilon s_\ell} \rfloor$. Then we adapt the supply of each good appropriately, to $s'_\ell = \lceil \frac{n}{\varepsilon} \rceil$. Given this *rounded* version of the problem instance, we will use dynamic programming to produce an allocation for it, which will immediately translate into an allocation for the original problem instance, that is welfare-optimal and violates the (original) supply constraints by a factor at most $(1 + \epsilon)$. For the purposes of the description that follows, we denote by $\mathbf{d}'$ the rounded version of a demand $\mathbf{d}$.

We define the dynamic programming table $\mathcal{V}(i, Y_1, \dots Y_m)$ for $i = 1, \dots, n$ and $Y_\ell \in \{0, 1, 2, \dots, s'_\ell\}$ for any $\ell \in [m]$. The cell $\mathcal{V}(i, Y_1, \dots, Y_m)$ stores the maximum welfare of an allocation $\mathbf{X}$, i.e., $\sum_j v_j(\mathbf{x}_j)$, whose rounded version $\mathbf{X}' = (\lfloor \frac{n \cdot \mathbf{x}_j(\ell)}{\varepsilon s_\ell} \rfloor)_{j,\ell}$ uses only multisets that are in the demand sets of the bidders in $\{1, 2, \dots, i\}$, and has total demand w.r.t. good $\ell = 1, \dots, m$ 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 $\mathcal{V}(1, Y_1, \dots Y_m)$ for any collection of Y_ℓ's such that: $(Y_1, \dots, Y_m) \in \{0, 1, \dots, \lceil \frac{n}{\varepsilon} \rceil\}^m$, is easy to solve. For each such entry $\mathcal{V}(1, Y_1, \dots Y_m)$ 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 $\mathcal{V}(1, Y_1, \dots, Y_m) = v_1(\mathbf{d})$ and build an auxiliary table $A[1, Y_1, \dots Y_m]$ which we set in this case to $\{(1, \mathbf{d})\}$. Otherwise, if bidder 1 does not have any such alternative, we assign $\mathcal{V}(1, Y_1, \dots Y_m) = 0$ and $A[1, Y_1, \dots Y_m] = \{(1, \mathbf{0})\}$. To define $\mathcal{V}(i+1, Y_1, \dots, Y_m)$, consider bidder $i+1$ and his alternatives $\mathbf{d} = (d(1), \dots, d(m)) \in \mathcal{D}_{i+1}$; let now

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

where, for all i, we define $\mathcal{V}(i, Y_1, \dots Y_m) = -\infty$ and, accordingly, $A[i, Y_1, \dots Y_m] = \{\, (i, \mathbf{0}) \,\}$, if there is no demand $\mathbf{d} \in \mathcal{D}_i$ satisfying $\mathbf{d'} \le \mathbf{Y}$. Consequently:

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

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

$$A[i+1, Y_1, \dots, Y_m] = A[i, Y_1, \dots, Y_m] \cup \{(i+1, \mathbf{0})\},$$

otherwise:

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

where $\mathbf{d}$ is an alternative in $\mathcal{D}_{i+1}$ maximizing (1). Finally, we inspect all the solutions from entries $\mathcal{V}(n, Y_1, \dots, Y_m)$ for all vectors $(Y_1, \dots, Y_m) \in \{0, 1, \dots, \lceil \frac{n}{\varepsilon} \rceil\}^m$, 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 the bidders' values is easy to verify. Let $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ denote any feasible allocation to the original problem instance. For every good, $\ell = 1, \dots, m$, we have: $\sum_i x_i(\ell) \le s_\ell$, or, equivalently, $\sum_i \frac{x_i(\ell) \cdot n}{\varepsilon \cdot s_\ell} \le \frac{n}{\varepsilon}$, thus $\sum_i \left\lfloor \frac{x_i(\ell) \cdot n}{\varepsilon \cdot s_\ell} \right\rfloor \le \lceil \frac{n}{\varepsilon} \rceil = s'_\ell$. That is, $\mathbf{X}$ is also feasible for the rounded problem instance. Because the dynamic programming algorithm will inspect all feasible solutions to the rounded problem instance and output the one with the largest welfare for it, an optimum solution to the original problem instance will be inspected as well.

We argue that the supply constraints s_ℓ, $\ell = 1, \dots, m$, are violated by at most a factor of $1 + 2\varepsilon$. Fix any good $\ell \in \{1, \dots, m\}$ and let $\mathbf{X}$ be the output allocation, *with respect* *to the original problem instance.* Because $\mathbf{X}$ was chosen by the algorithm by means of dynamic programming search over the rounded problem instance, it is feasible for the rounded problem instance. Thus, 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, since:

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

we obtain: $\sum_i x_i(\ell) \leq (1 + 2\varepsilon) s_\ell$.

Example Let us illustrate the algorithm's functionality through a very simple example. Consider $n = 3$ bidders and $m = 2$ distinct goods. Let the supplies of goods be $s_1 = s_2 = 4$. The bidders' values and demand sets are defined as follows:

Bidder	Valuation Function		Demand Set
1	$v_1((3,4)) = 1$	$v_1((4,3)) = 2$	$\mathcal{D}_1 = \{(3,4), (4,3)\}$
2	$v_2((3,3)) = 3$		$\mathcal{D}_2 = \{(3,3)\}$
3	$v_3((2,3)) = 4$	$v_3((3,2)) = 5$	$\mathcal{D}_3 = \{(2,3), (3,2)\}$

In this example it is evident that all feasible allocations involve assignment of a demand to a single bidder, given that the supplies of both goods are 4. Thus, the optimal allocation $\mathbf{X}^*$ is $(\mathbf{0}, \mathbf{0}, (3,2))$. Consider the *rounded problem* for $\epsilon = 2$. The *"rounded"* supply for each of the two goods is $\lceil n/\epsilon \rceil = \lceil 4/2 \rceil = 2$. The *"rounded"* demands of the bidders are as follows:

Bidder	1		2	3	
Demands	$(3,4)$	$(4,3)$	$(3,3)$	$(2,3)$	$(3,2)$
Rounded Demands	$(1,1)$	$(1,1)$	$(1,1)$	$(0,1)$	$(1,0)$

Observe that both demands of bidder 1 are rounded to $(1,1)$. This does not pose any problem, as the algorithm processes the original demands, and only uses their rounded versions to validate feasibility of the allocation it builds with respect to the rounded supply. In this example, because the rounded supply of each good is 2, the algorithm will output the allocation $\mathbf{X} = (\mathbf{0}, (3,3), (3,2))$, which has welfare 8 and is superoptimal for the inital instance. Although the rounded versions of the allocated demands do not violate the rounded supplies of the goods (equal to 2), they do violate the original supplies of 4, by less than a factor of $1 + \epsilon = 3$ (particularly in this example, by no more than a factor of 1.5).

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. For a fixed $\epsilon > 0$, the* Fptas *fully optimizes the social welfare, while violating the supplies of goods within factor at most $(1 + \epsilon)$.*

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 $\mathbf{b} = \Big((b_i(\mathbf{d}), \mathbf{d})_{\mathbf{d} \in \mathcal{D}_i}\Big)_{i \in \mathcal{N}}$ 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 $\mathbf{b} = \Big((b_i(\mathbf{d}), \mathbf{d})_{\mathbf{d} \in \mathcal{D}_i}\Big)_{i \in \mathcal{N}}$; by definition of the range, there exists a bid vector $\bar{\mathbf{b}}$, with $\bar{\mathbf{b}} = \Big((\bar{b}_i(\bar{\mathbf{d}}), \bar{\mathbf{d}})_{\bar{\mathbf{d}} \in \bar{\mathcal{D}}_i}\Big)_{i \in \mathcal{N}}$ such that $\mathcal{A}(\bar{\mathbf{b}}) = \mathbf{X}$. Recall that $x_i(\ell)$, for bidder i and $\ell = 1, \ldots, m$, is 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 $\bar{d}_i \in \bar{\mathcal{D}}_i \cup \{\mathbf{0}\}$ 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 $\mathbf{e}_i \in \mathcal{D}_i \cup \{\mathbf{0}\}$. 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 $\mathcal{A}$. □

A related result of Briest et al. (2011) 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.

4.1 A Note on Hardness

Note that this problem is strongly **NP**-hard, when we do not allow to violate supply constraints and $m \geq 2$ (Chekuri & Khanna, 2005). It is well known that if a problem is strongly **NP**-hard, there does not exist any FPTAS for this problem, unless **P**=**NP** (see, e.g., Vazirani, 2003). 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 in polynomial time within $m^{1/2-\epsilon}$, for any $\epsilon > 0$ (Lehmann et al., 2002). This claim is true, even if we allow for solutions to violate the supplies. 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 $|E| = m$ (Håstad, 1996). By using a reduction of Lehmann et al. (2002), 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 the graph G and each bidder's valuation for his set is 1. Now if we allow to violate the supply 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. $\square$

4.2 Multi-dimensional Knapsack

We discuss an application of our Fptas, in relation to the *Multi-dimensional Knapsack Problem* (MdKP) (Chekuri & Khanna, 2005). Suppose we are given a MdKP instance, with a constant number of distinct compartments, $m = O(1)$, and each compartment, $\ell = 1, 2, \ldots, m$, has capacity s_ℓ. The problem asks to fit in the knapsack a subset out of a universe, $\mathcal{U}$, of n given m-dimensional objects, so that the sum of the collected objects' sizes in each dimension, ℓ, does not exceed s_ℓ and the total value of all collected objects is maximized. Each object, $i = 1, \ldots, n$ of the MdKP instance can be represented by a vector $\mathbf{d}_i$, which represents its m dimensions, $(d_i(1), \ldots, d_i(m))$ and its value, v_i, so that $\langle v_i, \mathbf{d}_i \rangle \in ((\mathbb{R}^+ \cup \{0\}) \times \mathcal{U})$. Then, each object corresponds to a single bidder i, from the setting that we analyzed in Section 4, with valuation function $v_i(\mathbf{d}) \equiv v_i$, for every $\mathbf{d} \geq \mathbf{d}_i$, $v_i(\mathbf{d}) = 0$, for every $\mathbf{d}$ such that $d(\ell) < d_i(\ell)$, for some $\ell = 1, \ldots, m$. Notice that the bidder is *single-parameter*, in that his valuation function takes on a single non-zero value for every $\mathbf{d} \geq \mathbf{d}_i$ and his demand set, $\mathcal{D}_i$, is singleton, i.e., it contains a single multiset, $\mathbf{d}_i$. Thus, the MdKP corresponds to a single-parameter version of the problem we treated above.

We can apply our Fptas to the MdKP, because the algorithm is *exact*, as mentioned previously, in that it allocates every bidder (read as: "fits in the knapsack") either an exact alternative from his demand set, $\mathcal{D}_i$, or none. It is worth mentioning that for this single-parameter version, our Fptas from Section 4 can be shown to be *monotone* (Lehmann et al., 2002; Briest et al., 2011), when one carefully fixes a tie-breaking rule. A monotone allocation algorithm ensures that: if a (single-parameter) bidder i is allocated his single demand $\mathbf{d}_i$ when he declares truthfully $\langle v_i, \mathbf{d}_i \rangle$, he also receives his (declared) demand $\mathbf{d}'_i$, when he declares $\langle v'_i, \mathbf{d}'_i \rangle$, with $v'_i \geq v_i$ and $\mathbf{d}'_i \leq \mathbf{d}_i$ (i.e., intuitively, asks for less items while offering more money). An exact and monotone allocation algorithm can yield a truthful mechanism for this single-parameter setting, with the incorporation of *critical value* payments – see the work of Lehmann et al. (2002) for details.

Let us note that, we can generalize the MdKP further, in the following manner. Instead of having only *packing constraints* (of the form $\leq$) on the dimensions of the knapsack, we can handle any mix of *packing* and *covering* constraints (i.e., of any of the forms $\{\geq, \leq\}$), as long as there is only a constant number of dimensions, $m = O(1)$, and one *covering* or *packing* constraint per dimension. For such a generalized scenario we can follow an approach similar to our approach in Section 4 and obtain a truthful Fptas which fully optimizes the total value of fitted items and violates each of the constraints by a factor at most $(1+\varepsilon)$. Violation of the constraints is needed for the reason mentioned above, in our note on computational hardness, in the end of Section 4.

5. The Generalized Dobzinski-Nisan Method

We discuss here a direct generalization of a method designed by Dobzinski and Nisan (2010), 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 time complexity $\mathcal{T}_\mathcal{A}(t, \mathbf{s})$, $\mathbf{s} = (s_1, \ldots, s_m)$, and approximation ratio $\alpha \leq 1$. Then, algorithm $\mathcal{A}$ can be used as a routine within the procedure of Figure 1, to obtain a polynomial-time MiR algorithm for n bidders, with approximation ratio $(\alpha - \frac{m}{t+1})$.

Given $t = O(1)$, the procedure executes algorithm $\mathcal{A}$ on every subset of at most t bidders 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 (multi-unit) bundles from each good. The main result shown by Dobzinski and Nisan (2010) for a single good can be also proved for m goods:

Theorem 2 *Let $\mathcal{A}$ be a Maximal-in-Range algorithm with complexity $\mathcal{T}_\mathcal{A}(t, (s_1, \ldots, s_m))$, for t bidders and at most s_ℓ units from each good $\ell = 1, \ldots, m$. The Dobzinski-Nisan Method is* MiR *and runs in time polynomial in $\log s_1, \ldots, \log s_m$, n, $\mathcal{T}_\mathcal{A}(t, (s_1, \ldots, s_m))$, for every $t = O(1)$. Moreover, it outputs an allocation with value at least a fraction $(\alpha - \frac{m}{t+1})$ of the optimum Social Welfare.*

The proof is a direct extension of the proof given by Dobzinski and Nisan (2010) 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 $m = O(1)$ distinct goods, each in limited supply s_ℓ, $\ell \in [m]$. Let $\mathcal{R}_\mathcal{A}$ denote the *range* of this algorithm. It can be verified that the method outputs allocations that are "$(\mathcal{R}, t, \chi_1, \ldots, \chi_m)$-round", given the following definition for such *"round"* allocations (Dobzinski & Nisan, 2010):

Definition 3 *For some $t = O(1)$, 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 $\ell = 1, \ldots, m$.*
- *There exists a set T of $|T| \leq t$ bidders, such that they are all allocated according to some allocation in $\mathcal{R}$.*
- *Each bidder $i \in [n] \setminus T$ receives an exact multiple of $\max\left\{\lfloor \frac{\chi_\ell}{2n^2} \rfloor, 1\right\}$ units from good ℓ and: $\sum_{i \in [n] \setminus T} x_i(\ell) \leq n \cdot \max\left\{\lfloor \frac{\chi_\ell}{2n^2} \rfloor, 1\right\}$, for $\ell = 1, \ldots, m$*

1. **for** $\ell = 1, \ldots, m$ **do:**
   (a) **define** $u_\ell := (1 + \frac{1}{2n})$
   (b) **define** $L_\ell := \left\{0, 1, \lfloor u_\ell \rfloor, \lfloor u_\ell^2 \rfloor, \ldots, \left\lfloor u_\ell^{\lfloor \log_{u_\ell} s_\ell \rfloor} \right\rfloor, s_\ell \right\}$
2. **for** every subset $T \subseteq [n]$ 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 (if $\chi_\ell > 0$) 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 $[n] \setminus T$.
3. **Return** the best allocation found.

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

In this definition, $\mathcal{R}$ corresponds to the range of $\mathcal{A}$, parameterized by the subset of bidders T, i.e., $\mathcal{R} = \mathcal{R}_{\mathcal{A}}(T)$, on which it is executed. Then, for some $t = O(1)$, the range of the method is the subset of *all* allocations that are $(\mathcal{R}_{\mathcal{A}}(T), \tau, \chi_1, \ldots, \chi_m)$-round, so that: $(\chi_1, \ldots, \chi_m) \in (\times_{\ell=1}^m L_\ell)$, where L_ℓ is defined as in step **1.(b)** of the method in Figure 1, and $T \subseteq [n]$, with $\tau = |T| \leq t$. Formally, the method's range $\mathcal{R}_{\texttt{DN}}$ is the subset of allocations:

$$\mathcal{R}_{\texttt{DN}} = \left\{ \mathbf{X} \,\middle|\, \mathbf{X} \text{ is } (\mathcal{R}_{\mathcal{A}}(T), \tau, \chi_1, \ldots, \chi_m)\text{-round, for } (\chi_\ell)_\ell \in (\times_{\ell=1}^m L_\ell) \text{ and } \tau = |T| \leq t \right\}$$

Example – Part (I) Before continuing to analyze the method's range, let us exemplify the concept of $(\mathcal{R}, t, \chi_1, \ldots, \chi_m)$-round allocations. We will consider a small instance of multi-minded bidders, similar to the considered in the previous section. As we will argue later, the Dobzinski-Nisan method can yield a truthful Ptas (that respects the supply constraints on goods), for multi-minded bidders. Assume $m = 2$ distinct goods, and $n = 5$ bidders. We assume supplies $s_1 = 200 = s_2$ for the goods. The bidders' demands are as follows:

Bidder	Demand Set
1	$\mathcal{D}_1 = \{(75, 51), (49, 73)\}$
2	$\mathcal{D}_2 = \{(51, 27), (25, 49)\}$
3	$\mathcal{D}_3 = \{(48, 1)\}$
4	$\mathcal{D}_4 = \{(1, 1)\}$
5	$\mathcal{D}_5 = \{(1, 48)\}$

Let us exhibit a *"round"* allocation for this instance, according to Definition 3. For $t = 2$ and $\chi_1 = \chi_2 = 100$, consider first the allocation:

$$\mathbf{X} = (\,(75, 51),\ (25, 49),\ (48, 2),\ (2, 2),\ \mathbf{0}\,),$$

where $\mathbf{x}_1 = (75, 51)$, $\mathbf{x}_2 = (25, 49)$, $\mathbf{x}_3 = (48, 2)$, $\mathbf{x}_4 = (2, 2)$, $\mathbf{x}_5 = \mathbf{0}$. This allocation is $(\mathcal{R}, 2, 100, 100)$-round, according to Definition 3, where $\mathcal{R}$ denotes the subset of allocations with at most 2 bidders receiving non-empty multisets and the remaining ones receiving appropriate multi-unit bundles per good. Indeed, we can set $T = \{1, 2\}$ (for the corresponding subset of at most 2 bidders); each of bidders 1 and 2 obtains one of his demands. The total number of units allocated to these two bidders per good is exactly $100 = s_\ell - \chi_\ell$. For the remaining $100 = \chi_\ell$ units from each good, we make 50 bundles of $\chi_\ell/(2n^2) = 2$ units per bundle. Bidder 3 receives 24 such 2-units bundles from good 1 and one 2-units bundle from good 2. Bidder 4 obtains one 2-units bundle from each good. Finally, bidder 5 receives an empty allocation. Notice that $\mathbf{x}_3$ and $\mathbf{x}_4$ essentially satisfy the unique demands $(48, 1)$ and $(1, 1)$ of bidders 3 and 4 respectively.

Another $(\mathcal{R}, 2, 100, 100)$-round allocation (according to Definition 3) is:

$$\mathbf{X}' = (\,(75, 51),\ (52, 28),\ (48, 2),\ \mathbf{0},\ \mathbf{0}\,),$$

where the required subset T of bidders is $T = \{1\}$. Bidder 1 obtains one of his demands; bidders 2 and 3 receive 2-units bundles from each good; and bidders 4 and 5 receive empty allocations (i.e., zero 2-units bundles from each good). Notice that $\mathbf{X}'$ is also $(\mathcal{R}, 1, 100, 100)$-round (i.e., when we set $t = 1$). Now let us choose algorithm $\mathcal{A}$ for the Dobzinski-Nisan method, to be an exhaustive search procedure, that optimizes the welfare (thus, has approximation ratio $\alpha = 1$). Notice that the range $\mathcal{R}_{\mathcal{A}}(T)$ of $\mathcal{A}$, for the chosen values of χ_1, χ_2, trivially contains the allocation that bidders 1 and 2 receive in $\mathbf{X}$ (when $T = \{1, 2\}$) and the allocation of bidder 1 under $\mathbf{X}'$ (when $T = \{1, 2\}$ or $T = \{1\}$); this is because $\mathcal{A}$ optimizes over all feasible allocations up to supplies $200 - \chi_1 = 100$ and $200 - \chi_2 = 100$ for each of the two choices of T. Thus, the allocations $\mathbf{X}$ and $\mathbf{X}'$ also belong in the range $R_{\mathtt{DN}}$, as defined above.

We show that optimization over $\mathcal{R}_{\mathtt{DN}}$ approximates the socially optimal allocation within factor $(\alpha - \frac{m}{t+1})$.

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

Proof. In the proof we make use of notation L_ℓ and u_ℓ, as defined in Figure 1. 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 $s_\ell - \chi_\ell \geq \sum_{i=1}^t x_i^*(\ell)$. When executed 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 $\ell = 1, \ldots, m$ 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 in $L_\ell = \{\,0,\, 1,\, \lfloor u_\ell \rfloor,\, \lfloor u_\ell^2 \rfloor,\, \ldots,\, s_\ell\}$ satisfying $s_\ell \geq \chi_\ell + \sum_{i=1}^t x_i^*(\ell)$, 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. □

The lemma completes the proof of Theorem 2.

Example – Part (II) We revisit the example discussed right before the statement and proof of Lemma 1, in order to exemplify the approximation implied by the Lemma. To this end, we assign values to the bidders' demands as described in the following table, where $v > 0$ is a very small positive number and $V >> v$ is a very large one. As before, $s_1 = s_2 = 200$.

Bidder	Valuation Function	Demand Set
1	$v_1(\,(75,51)\,) = v \quad v_1(\,(49,73)\,) = V$	$\mathcal{D}_1 = \{\,(75,51),\,(49,73)\,\}$
2	$v_2(\,(51,27)\,) = V \quad v_2(\,(25,49)\,) = v$	$\mathcal{D}_2 = \{\,(51,27),\,(25,49)\,\}$
3	$v_3(\,(48,1)\,) = V$	$\mathcal{D}_3 = \{\,(48,1)\,\}$
4	$v_4(\,(1,1)\,) = v$	$\mathcal{D}_4 = \{\,(1,1)\,\}$
5	$v_5(\,(1,48)\,) = V$	$\mathcal{D}_5 = \{\,(1,48)\,\}$

The following socially optimal allocation $\mathbf{X}^*$ for this instance has welfare $4V + v$:

$$\mathbf{X}^* = (\,(49,73),\,(51,27),\,(48,1),\,(1,1),\,(1,48)\,)$$

By choosing $t = 2$ and $T = \{1,2\}$, we can exhibit the welfare-approximate allocation implied by Lemma 1, as follows. The maximum value possible for each of χ_1 and χ_2 satisfying $s_\ell - \chi_\ell \geq x^*_1(\ell) + x^*_2(\ell)$ is $100 = \chi_1 = \chi_2$. From the remaining bidders, bidder 3 has the maximum number of units from good 1 in $\mathbf{X}^*$ and bidder 5 has the maximum number of units from good 2 in $\mathbf{X}^*$. Thus, we have $j_1 = 3$ and $j_2 = 5$. Then, we round up the allocations of bidders 3 and 4 w.r.t. good 2, to one 2-units bundle (for each of them), by taking two units from bidder $j_2 = 5$. Accordingly, we round up the allocations of bidders 4 and 5 w.r.t. good 1, by taking two units from bidder $j_1 = 3$. The resulting allocation is:

$$\mathbf{X} = (\,(49,73),\,(51,27),\,(46,2),\,(2,2),\,(2,46)\,)$$

and has welfare $2V + v$, which approaches half of the optimal welfare (as v becomes vanishingly small). Lemma 1 for this example guarantees at least 1/3 of the optimal welfare, if the algorithm $\mathcal{A}$ used within the Dobzinski-Nisan method is a welfare-optimizing exhaustive search procedure. On the other hand notice that, for this particular example and $t = 2$, $\chi_1 = \chi_2 = 100$, the allocation $\mathbf{Y} = (\,(49, 73),\ (51, 27),\ (48, 2),\ \mathbf{0},\ (2, 48)\,)$ has almost optimal welfare, $4V$, and is *"round"* according to Definition 3. Thus, the Dobzinski-Nisan method will examine $\mathbf{Y}$ and it will return an allocation at least as good.

Let us explain how to find an optimal allocation of multi-unit bundles of goods (i.e., bundles of identical units) to bidders in $[n] \setminus T$, in step **2.1.3** of the algorithm (Figure 1). We use dynamic programming. By re-indexing the bidders appropriately, assume that $T = \{n-t+1, \ldots, n\}$, thus $[n] \setminus T = \{1, \ldots, n-t\}$. 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:

$$\mathcal{V}(i, \mathbf{q}) = \max_{\mathbf{q}' \leq \mathbf{q}} \Big(v_i(q'_1 \cdot b_1, \ldots, q'_m \cdot b_m) + \mathcal{V}(i-1, \mathbf{q} - \mathbf{q}') \Big),$$

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 $\ell = 1, \ldots, m$.

5.1 Simple Application: Multi-minded Bidders

The generalized Dobzinski-Nisan method for multiple distinct goods can be applied immediately in the setting of multi-minded bidders, to yield a Ptas that respects 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, if every bidder's demand sets contains at most k demands, 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, complementarily to the developments of the previous section, approximates the optimum Social Welfare within factor $(1 + \epsilon)$ and respects the supply constraints.

5.2 Submodular Valuation Functions

We consider submodular valuation functions over multisets in $\mathcal{U}$, as defined by Kapralov, Post, and Vondrák (2013):

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., that the algorithm asks the bidders for their value, for each particular multiset that it needs to process.

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

(2010). For any $\epsilon > 0$, define $\delta = 1 + \epsilon$; we will be assigning to bidders multi-unit bundles of each good $\ell \in [m]$, that have cardinality equal to an integral power of δ. For every good $\ell \in [m]$, 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 a 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 optimal 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 multi-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 $[n] = \mathcal{O} \cup \mathcal{E}$ 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:

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

$$\geq \frac{\delta-1}{\delta} SW\left(\mathbf{X}^{[\ell-1]}\right) + \frac{2-\delta}{\delta} SW\left(\mathbf{X}^{[\ell-1]}\right) = \frac{1}{1+\epsilon} SW\left(\mathbf{X}^{[\ell-1]}\right)$$

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 $v_{i-1}\left(\mathbf{x}_{i-1}^{[\ell]}\right) \geq \frac{2-\delta}{\delta} v_i\left(\mathbf{x}_{i-1}^{[\ell-1]}\right)$. For the last inequality, recall that we are examining the left-hand side case of (3), thus, we use the that: $\sum_{i\in\mathcal{O}} v_i(\, x_i^{[\ell-1]}\,) \geq \frac{1}{2} SW(\mathbf{X}^{[\ell-1]})$.

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 $x_{i-1}^{[\ell-1]}(\ell)$ 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:

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

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

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

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 $x_2^{[\ell]}(\ell) \geq \frac{1}{\delta} x_2^{[\ell-1]}(\ell)$. For the last inequality, we used the fact that we are examining the right-hand side case of (3); then, $\sum_{i\in\mathcal{E}} v_i(\,\mathbf{x}_i^{[\ell-1]}\,) \geq \frac{1}{2} SW(\mathbf{X}^{[\ell-1]})$.

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 $\frac{1}{1+\epsilon} \geq \frac{2-\epsilon}{2+2\epsilon}$. □

We obtain the following (intermediate) result:

Theorem 3 *For multi-unit combinatorial auctions with $n = O(1)$ submodular bidders, and $m = O(1)$ distinct goods, each good $\ell \in [m]$ available in an arbitrary supply, there exists a truthful deterministic* Fptas *that, for any $\epsilon \leq 1$, 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 by Dobzinski and Nisan (2010) 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 (constant approximation) result. Instead, a constant $(m+1)$-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 by Dobzinski and Nisan, for general valuation functions.

We describe from scratch an 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 (per distinct good, ℓ), 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 $(m+1)$ the optimum Social Welfare. Let $\mathbf{X}^* = (\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ denote the socially optimal allocation. Beginning with $\mathbf{X}^*$, we produce an allocation in the range within which the algorithm optimizes, that approximates $SW(\mathbf{X}^*)$ within factor $(m+1)$. Assume w.l.o.g. that all items are allocated in $\mathbf{X}^*$ (by the monotonicity of valuation functions) and, for each good $\ell = 1, \ldots, m$, let

$j_\ell = \arg\max_i x_i^*(\ell)$. Then $x_{j_\ell}^*(\ell) \geq \frac{s_\ell}{n}$. Define $\mathcal{L} = \{j_1, \ldots, j_m\}$. We consider two cases here.

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

For the first case, let us denote by $\mathbf{Y}^\ell = (\mathbf{y}_1^\ell, \mathbf{y}_2^\ell, \ldots, \mathbf{y}_n^\ell)$ – for each $\ell = 1, \ldots, m$ – the allocation which assigns *all bundles* of *all goods* to bidder $j_\ell \in \mathcal{L}$ (thus, $\mathbf{y}_i^\ell = (0, \ldots, 0)$, for every $i \neq j_\ell$). Of these m allocations, consider $\mathbf{Y} = \arg\max_{\mathbf{Y}^\ell} v_{j_\ell}(\mathbf{y}_{j_\ell}^\ell)$. Then, $SW(\mathbf{Y}) \geq \frac{1}{m}\sum_{\ell=1}^m v_{j_\ell}(\mathbf{x}_{j_\ell}^*)$, thus, also: $SW(\mathbf{Y}) \geq \sum_{i \notin \mathcal{L}} v_i(\mathbf{x}_i^*)$. Putting these inequalities together yields $SW(\mathbf{Y}) \geq \frac{1}{m+1} SW(\mathbf{X}^*)$. Notice that the allocation $\mathbf{Y}$ will be examined by the MiR algorithm. For the second case we build an allocation $\mathbf{X}$, by rounding up – separately for each good ℓ – the (optimal) allocation of bidders $i \notin \mathcal{L}$ to the nearest multiple of b_ℓ. The units needed for this purpose we find – for each good ℓ – from the corresponding bidder $j_\ell \in \mathcal{L}$, 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 multi-unit bundles of each good to bidders in $[n] \setminus \mathcal{L}$ and satisfies $SW(\mathbf{X}) \geq \sum_{i \notin \mathcal{L}} v_i(\mathbf{x}_i^*)$, thus, also: $SW(\mathbf{X}) > \frac{1}{m}\sum_{\ell=1}^m v_{j_\ell}(\mathbf{x}_{j_\ell}^*)$. Then, we deduce $SW(\mathbf{X}) \geq \frac{1}{m+1} SW(\mathbf{X}^*)$. Notice that the allocation $\mathbf{X}$ is also examined by the MiR algorithm. Thus, there exists a solution within the range, that approximates $SW(\mathbf{X}^*)$ within constant factor, at most $(m+1)$.

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

$$\mathcal{V}^L(i, \mathbf{q}) = \max_{L' \subseteq L} \max_{q'_1 \leq q_1, \ldots, q'_m \leq q_m} \left\{ v_i\left((q'_1 \cdot b_1, \ldots, q'_m \cdot b_m) + \mathbf{r}[L'] \right) + \mathcal{V}^{L \setminus L'}(i-1, \mathbf{q} - \mathbf{q}') \right\}$$

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, m, and general valuation functions that, using value queries, approximates the welfare of a socially optimal assignment within constant factor, $(m+1)$.*

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 Maximal-in-Range allocation algorithms (Nisan & Ronen, 2007), 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 (2010). All of the discussed developments are best possible in terms of time-efficient approximation, as follows by relevant hardness results. Finally, we showed how to treat general (unrestricted) valuation functions in our setting, by appropriately adjusting an analysis by Dobzinski and Nisan (2010). Closing the gap between a communication complexity lower bound of 2 (for a single good) of Dobzinski and Nisan and our $(m+1)$-approximation result for $m = O(1)$ goods, requires further understanding of the communication complexity of our more general setting.

Acknowledgments

We thank three anonymous reviewers for helping us significantly in improving the presentation of this work. We also thank Jinshan Zhang, for pointing out a technical inconsistency in an earlier version of this paper, and Fabrizio Grandoni and Stefano Leonardi, for very useful discussions in the early stage of this work.

Piotr Krysta acknowledges support by the EPSRC grant EP/K01000X/1.

Carmine Ventre acknowledges support by the EPSRC grant EP/M018113/1.

Orestis Telelis acknowledges support 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 Union (European Social Fund – ESF) and the Greek State.

References

Archer, A., Papadimitriou, C. H., Talwar, K., & Tardos, E. (2003). An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. *Internet Mathematics*, *1*(2), 129–150.

Ausubel, L. M., & Milgrom, P. (2010). The Lovely but Lonely Vickrey Auction. In Cramton, P., Shoham, Y., Smith, V. L., & Steinberg, R. (Eds.), *Combinatorial Auctions*, pp. 17–40. MIT Press.

Bartal, Y., Gonen, R., & Nisan, N. (2003). Incentive Compatible Multi Unit Combinatorial Auctions. In Halpern, J. Y., & Tennenholtz, M. (Eds.), *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2003)*, pp. 72–87. ACM.

Bikhchandani, S., de Vries, S., Schummer, J., & Vohra, R. (2011). An Ascending Vickrey Auction for Selling Bases of a Matroid. *Operations Research*, *59*(2), 400–413.

Blumrosen, L., & Nisan, N. (2007). Combinatorial Auctions. In Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (Eds.), *Algorithmic Game Theory*, pp. 267–299. Cambridge University Press.

Briest, P., Krysta, P., & Vöcking, B. (2011). Approximation Techniques for Utilitarian Mechanism Design. *SIAM Journal on Computing*, *40*(6), 1587–1622.

Chekuri, C., & Khanna, S. (2005). A Polynomial Time Approximation Scheme for the Multiple Knapsack Problem. *SIAM Journal on Computing*, *35*(3), 713–728.

Clarke, E. (1971). Multipart Pricing of Public Goods. *Public Choice*, *11*(1), 17–33.

Daniely, A., Schapira, M., & Shahaf, G. (2015). Inapproximability of Truthful Mechanisms via Generalizations of the VC Dimension. In Servedio, R. A., & Rubinfeld, R. (Eds.), *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pp. 401–408. ACM.

Dobzinski, S., & Dughmi, S. (2013). On the Power of Randomization in Algorithmic Mechanism Design. *SIAM Journal on Computing*, *42*(6), 2287–2304.

Dobzinski, S., & Nisan, N. (2010). Mechanisms for Multi-Unit Auctions. *Journal of Artificial Intelligence Research*, *37*, 85–98.

Dobzinski, S., Nisan, N., & Schapira, M. (2010). Approximation Algorithms for Combinatorial Auctions with Complement-Free Bidders. *Mathematics of Operations Research*, *35*(1), 1–13.

Grandoni, F., Krysta, P., Leonardi, S., & Ventre, C. (2014). Utilitarian Mechanism Design for Multi-Objective Optimization. *SIAM Journal on Computing*, *43*(4), 1263–1290.

Groves, T. (1973). Incentives in Teams. *Econometrica*, *41*(4), 617–631.

Håstad, J. (1996). Clique is Hard to Approximate Within $n^{1-\epsilon}$. In *37th Annual Symposium on Foundations of Computer Science (FOCS'96)*, pp. 627–636. IEEE Computer Society.

Holzman, R., Kfir-Dahav, N. E., Monderer, D., & Tennenholtz, M. (2004). Bundling Equilibrium in Combinatorial Auctions. *Games and Economic Behavior*, *47*(1), 104–123.

Kapralov, M., Post, I., & Vondrák, J. (2013). Online Submodular Welfare Maximization: Greedy is Optimal. In Khanna, S. (Ed.), *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'13)*, pp. 1216–1225. SIAM.

Kelly, T. (2004). Generalized Knapsack Solvers for Multi-Unit Combinatorial Auctions: Analysis and Application to Computational Resource Allocation. In Faratin, P., & Rodríguez-Aguilar, J. A. (Eds.), *Agent-Mediated Electronic Commerce VI, Theories for and Engineering of Distributed Mechanisms and Systems (AAMAS 2004 Workshop, AMEC 2004)*, Vol. 3435 of *LNCS*.

Khot, S., Lipton, R. J., Markakis, E., & Mehta, A. (2008). Inapproximability Results for Combinatorial Auctions with Submodular Utility Functions. *Algorithmica*, *52*(1), 3–18.

Kothari, A., Parkes, D. C., & Suri, S. (2005). Approximately-Strategyproof and Tractable Multiunit Auctions. *Decision Support Systems*, *39*(1), 105–121.

Lavi, R., & Swamy, C. (2011). Truthful and Near-Optimal Mechanism Design via Linear Programming. *Journal of the ACM*, *58*(6), 25.

Lehmann, B., Lehmann, D. J., & Nisan, N. (2006). Combinatorial Auctions with Decreasing Marginal Utilities. *Games and Economic Behavior*, *55*(2), 270–296.

Lehmann, D., Müller, R., & Sandholm, T. (2010). The Winner Determination Problem. In Cramton, P., Shoham, Y., Smith, V. L., & Steinberg, R. (Eds.), *Combinatorial Auctions*. MIT Press.

Lehmann, D. J., O'Callaghan, L., & Shoham, Y. (2002). Truth Revelation in Approximately Efficient Combinatorial Auctions. *Journal of the ACM*, *49*(5), 577–602.

Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.

Mirrokni, V. S., Schapira, M., & Vondrák, J. (2008). Tight Information-Theoretic lower bounds for Welfare Maximization in Combinatorial Auctions. In *Proceedings 9th ACM Conference on Electronic Commerce (ACM EC)*, pp. 70–77.

Mu'alem, A., & Nisan, N. (2002). Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. In Dechter, R., & Sutton, R. S. (Eds.), *Proceedings of the Eighteenth National Conference on Artificial Intelligence and Fourteenth Conference on Innovative Applications of Artificial Intelligence (AAAI/IAAI 2002)*, pp. 379–384. AAAI Press / MIT Press.

Mu'alem, A., & Nisan, N. (2008). Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. *Games and Economic Behavior*, *64*(2), 612–631.

Nisan, N. (2014). Algorithmic Mechanism Design, Through the Lens of Multi-Unit Auctions. In Aumann, R., & Hart, S. (Eds.), *Handbook of Game Theory*, Vol. IV. Elsevier North-Holland.

Nisan, N., & Ronen, A. (2007). Computationally Feasible VCG Mechanisms. *Journal of Artificial Intelligence Research*, *29*, 19–47.

Nisan, N., & Segal, I. (2006). The Communication Requirements of Efficient Allocations and Supporting Prices. *Journal of Economic Theory*, *129*(1), 192–224.

Nisan, N., & Ronen, A. (2001). Algorithmic Mechanism Design. *Games and Economic Behavior*, *35*(1-2), 166–196.

Sandholm, T. (2010). Optimal Winner Determination Algorithms. In Cramton, P., Shoham, Y., Smith, V. L., & Steinberg, R. (Eds.), *Combinatorial Auctions*. MIT Press.

Vazirani, V. V. (2003). *Approximation Algorithms*. Springer-Verlag.

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

Vöcking, B. (2012). A Universally Truthful Approximation Scheme for Multi-Unit Auctions. In Rabani, Y. (Ed.), *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'12)*, pp. 846–855. SIAM.