Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Near-Optimal Approximation Mechanisms for Multi-Unit Combinatorial Auctions * †

Piotr Krysta
Department of Computer Science
University of Liverpool, U.K.
p.krysta@liverpool.ac.uk

Orestis Telelis
Department of Digital Systems
University of Piraeus, Greece
telelis@gmail.com

Carmine Ventre
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 involving 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, that are expressed through their private valuation functions. Our objective is to determine allocations of multisets that maximize the Social Welfare approximately. Despite the recent theoretical advances on the design of truthful combinatorial auctions (for multiple distinct goods in unit supply) and multi-unit auctions (for multiple units of a single good), results for the combined setting are much scarcer. We elaborate on the main developments of [Krysta *et al.*, 2013], concerning bidders with *multi-minded* and *submodular* valuation functions, with an emphasis on the presentation of the relevant algorithmic techniques.

1 Introduction

We present truthful mechanisms for *multi-unit* combinatorial auctions, involving constant number of distinct goods, each in limited supply. A 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 "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 may have distinct positive values for distinct *multisets* of goods, that specify his demand for (potentially) more than one unit per good. We discuss 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.

Mechanism Design for combinatorial auctions of multiple heterogeneous goods (each in unit supply) has received significant attention in recent years, since the work of [Lehmann *et al.*, 1999; 2002], due to their various applications, especially in online trading systems over the Internet. Research in *Algorithmic* Mechanism Design was initialized by [Nisan and Ronen, 1999]. A mechanism elicits bids from buyers, so as to determine an assignment of bundles to them and payments in such a manner, so that it is to each buyer's best interest to reveal his valuation function *truthfully* to the mechanism. The related problem of auctioning multiple units of *a single good* to multi-demand bidders was considered already by [Vickrey, 1961], through a (multi-unit) extension of his celebrated single-item Second-Price mechanism. This extension, however, is not polynomial-time for an arbitrary number of units. Polynomial-time approximation mechanisms for multi-unit auctions were designed relatively recently [Mu'alem and Nisan, 2002; 2008; Dobzinski and Nisan, 2010; Vöcking, 2012; Nisan, 2014]. In particular, [Nisan, 2014] devised a welfare-optimal deterministic polynomial-time auction mechanism, for the multi-unit setting with submodular bidders first considered by [Vickrey, 1961]. [Vöcking, 2012] designed and analyzed a *randomized* universally truthful polynomial-time approximation scheme, for bidders with unrestricted valuation functions.

Results for the more general setting of multi-unit combinatorial auctions are much scarcer [Bartal *et al.*, 2003; Grandoni *et al.*, 2014]. Here, we consider exactly this setting, with a constant number of distinct goods, as in [Grandoni *et al.*, 2014]; for a number of cases of such auctions we analyze *Maximum-in-Range* (MIR) allocation algorithms [Nisan and Ronen, 2007] (see definition 2 in Section 2), that can be paired with the Vickrey-Clarke-Groves payment rule, so as to yield truthful mechanisms.

*This paper was invited for submission to the Best Papers From Sister Conferences Track, based on a paper that appeared in the 12th International Conference AAMAS (2013).

†Work supported by EPSRC grants EP/K01000X/1 and EP/M018113/1.

1.1 Overview

We discuss truthful mechanisms for multi-unit combinatorial auctions with a constant number of goods developed recently in [Krysta *et al.*, 2013], for two broad classes of the bidders' valuation functions.

Multi-Minded Bidders Each bidder has positive value for every distinct multiset belonging to a specific *demand set* of alternatives associated with him (and at least as much for the value of every superset of this multiset). His value is zero elsewhere. For this class of valuation functions we describe in Section 3 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 $\epsilon > 0$. 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 (see the related discussion in Section 3). 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). This result improves upon an FPTAS by [Grandoni *et al.*, 2014], which approximates the social welfare *and* the supplies 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.

Submodular Bidders The value of each bidder for a multiset of items is given by a submodular valuation function. Thus, each bidder's marginal value for each additional item allocated to him (a unit of any good) is non-increasing. For this setting we develop a PTAS [1], that approximates the optimum social welfare within factor $(1+\epsilon)$, for any fixed $\epsilon > 0$, without violating the supply constraints. To this end, we revisit a technique introduced by [Dobzinski and Nisan, 2010], for multi-unit auction Mechanism Design, and generalize it for multiple distinct goods, each in limited supply. We describe how this generalization works and also how to use it to obtain a truthful PTAS for multi-minded bidders, that does *not violate* the supply constraints. This latter result is best possible; a hardness result from [Dobzinski and Nisan, 2010] rules out the possibility of a truthful FPTAS via their technique, for multi-minded bidders and a single good, unless $\mathbf{P} = \mathbf{NP}$.

The assumption of $m = O(1)$ distinct goods is important, for otherwise our two problem cases become hard to approximate in polynomial time, within factor less than $O(\sqrt{m})$ [Lehmann *et al.*, 2002] and $\frac{e}{e-1}$ [Khot *et al.*, 2008; Mirrokni *et al.*, 2008] respectively. Also, our techniques cannot yield *truthful* polynomial-time mechanisms with approximation factors less than $O(m)$ and $O(\sqrt{m})$ respectively [Daniely *et al.*, 2014]. We omit from our presentation a truthful constant-approximation mechanism from [Krysta *et al.*, 2013], for bidders with unrestricted valuation functions.

[1](F)PTAS stands for *(Fully) Polynomial Time Approximation Scheme*, see chapter 8 in [Vazirani, 2003] for a formal definition.

[2]In the context of social welfare *maximization*, by "approximation within factor $\rho \geq 1$" we mean recovering at least a fraction ρ^{-1} of the welfare of an optimum allocation.

1.2 Related Work

Mechanism Design for multi-unit auctions was initiated already by the celebrated work of [Vickrey, 1961], where the well-known single-item Second-Price auction is extended for the case of multiple units and bidders with *symmetric* submodular valuation functions [Lehmann *et al.*, 2006]. This mechanism is however not computationally efficient with respect to the number of available units. A polynomial-time truthful mechanism for this case was analyzed in [Nisan, 2014]. 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]. The authors analyzed a truthful 2-approximation mechanism for a multi-unit setting with multiple distinct goods (in limited supply), and *single-minded* bidders, each valuing positively a particular subset of goods (and its supersets). [Archer *et al.*, 2003] improved on this approximation ratio for a similar setting, but their mechanism was based on randomized rounding and was truthful only *in expectation*. Furthermore, an FPTAS was designed and analyzed in [Briest *et al.*, 2005] for single-minded bidders, in the multi-unit combinatorial setting.

[Dobzinski and Nisan, 2007; 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 k-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. Moreoever, 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. For a special case of this class, [Kothari *et al.*, 2005] had designed earlier an FPTAS mechanism that was, however, only *approximately truthful*. [Dobzinski and Dughmi, 2009] gave a *truthful in expectation* FPTAS for multi-minded bidders. A universally truthful randomized PTAS for general valuation functions accessed by value queries was developed in [Vöcking, 2012] (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 demanding at most one unit per good (e.g. [Lehmann *et al.*, 2002; Briest *et al.*, 2005]). In contrast, we consider a constant number of goods, but multi-demand bidders. [Bartal *et al.*, 2003] proved approximation 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 mechanisms for a constant number of distinct goods, each in limited supply, was initialized by [Grandoni *et al.*, 2014]. The authors utilized methods from multi-objective optimization, to obtain truthful polynomial-time approximation schemes for problems including: *multi-unit auctions*, *minimum spanning tree*, *shortest path*, *maximum* (perfect) *matching* and *matroid intersection*. They devised truthful FPTASes that approximate the objective function (social welfare or cost) of multi-capacitated versions of these problems within factor $(1+\epsilon)$, while exceeding the capacities by a factor $(1+\epsilon)$ (capacities here correspond to limited supplies).

2 Preliminaries

We consider a set $[m]$, of $m = O(1)$ goods, $[m] = \{1, \ldots, m\}$. Each good $\ell \in [m]$ is available in limited supply (number of units) $s_\ell \in \mathbb{N}$. 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 bidders (prospective buyers). 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 *monotone non-decreasing*: for any two multisets $\mathbf{x} \leq \mathbf{y}$ – where "$\leq$" holds componentwise – we have $v_i(\mathbf{x}) \leq v_i(\mathbf{y})$.

A *mechanism* consists of an *allocation* method (algorithm), $\mathcal{A}$, and a *payment* rule, $\mathbf{p}$. The allocation algorithm, $\mathcal{A}$, elicits from the bidders bids $\mathbf{b} = (\, b_1, b_2, \ldots, b_n \,)$ that, presumably, describe their valuation functions, and determines 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 current 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 that we discuss in our work 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:

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

where the second equality stems from a standard assumption of *no externalities*, i.e., that the value of any bidder for $\mathcal{A}(\mathbf{b})$ depends *only* on his own individual allocation.

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 outcome $\mathbf{b} = \mathbf{v}$ is a *dominant strategy* equilibrium. We are interested in designing and analyzing *truthful* mechanisms $(\mathcal{A}, \mathbf{p})$, that render truthful reporting of the bidders' valuation functions 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 Second-Price and multi-unit auctions [Vickrey, 1961]. Deployment of the VCG auction, however, requires allocation algorithms $\mathcal{A}$ that output welfare-maximizing allocations, 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 are indeed **NP**-hard, our mechanisms use *Maximum-in-Range* (MIR) [Nisan and Ronen, 2007] allocation algorithms, that maximize the social welfare *approximately*.

Definition 2 *[Nisan and 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 [Nisan and Ronen, 2007] identified MIR allocation algorithms as the sole device that, along with VCG payments, yields truthful mechanisms for combinatorial auctions. That is, any MIR allocation algorithm, $\mathcal{A}$, can be "turned" into a truthful mechanism, if the payment for each bidder i is computed according to the VCG payments scheme, as follows:

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

This payment scheme coincides with the VCG payment scheme, if $\mathcal{A}$ is replaced by the *optimal* allocation algorithm.

3 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 $\mathbf{e} \in \mathcal{U} \setminus \mathcal{D}_i$ we define:

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

Naturally, $v_i(\mathbf{0}) = 0$. In this setting, the valuation function of a bidder i can be expressed compactly, 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. A bidder i is a *winner* of the auction, if he is assigned exactly one – or a superset of one – of his alternatives from $\mathcal{D}_i$.

The allocation algorithm of our mechanism will be an FPTAS, that *maximizes the social welfare* and may violate the supply constraints on goods by a factor at most $(1 + \varepsilon)$, for any fixed $\epsilon > 0$. The algorithm is reminiscent of the FPTAS for the well-known one-dimensional knapsack problem (see e.g., chapter 8 in [Vazirani, 2003]). For any chosen fixed $\varepsilon > 0$, first it discards any alternative from the bidders' demand sets, that violates any of the supply constraints. Subsequently, the quantities of goods in the remaining multisets are *rounded* appropriately; the supplies are also adjusted. Thus we 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 violates the (initial) supplies by a factor at most $(1+\epsilon)$. In light of turning this algorithm into a truthful mechanism, we use notation of actual valuation functions in its description below.

Rounding 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 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; this will directly translate into an allocation for the original problem instance, that is welfare-optimal and violates the (original) supply constraints by a factor at most $(1+\epsilon)$. In the sequel, for any multiset $\mathbf{d}$ in the demand set $\mathcal{D}_i$ of some bidder i, we use $\mathbf{d}'$ to refer to the *rounded* multiset as described above.

Dynamic Programming We use two tables, $\mathcal{V}$ and A, for storing respectively social welfare *values* and *allocations*. Each cell of each table is indexed by the tuple $(i, q_1, \ldots, q_m)$, for $i = 1, \ldots, n$, where $q_\ell \in \{\, 0, 1, 2, \ldots, \lceil n/\epsilon \rceil\}$, for $\ell = 1, \ldots, m$. Notice that there are $n(1 + \lceil n/\epsilon \rceil)^m$ such cells in each table, which is a polynomially bounded number for $m = O(1)$ and any fixed $\epsilon > 0$. To simplify notation, in the sequel we use $\mathbf{q}$ for $(q_1, \ldots, q_m)$. $\mathcal{V}(i, \mathbf{q})$ stores the maximum welfare of an allocation $\mathbf{X}$, i.e., $\sum_j v_j(\mathbf{x}_j)$, whose rounded version $\mathbf{X}' = (\lfloor \frac{n \cdot \mathbf{x}_j(\ell)}{\varepsilon s_\ell} \rfloor)_{j,\ell}$:

1. uses only multisets that are in the demand sets of the bidders in $\{1, 2, \ldots, i\}$,
2. has total demand w.r.t. good $\ell = 1, \ldots, m$ which is precisely y_ℓ, i.e., $\sum_i x'_i(\ell) = q_\ell$.

The corresponding cell $A[i, \mathbf{q}]$ holds this allocation $\mathbf{X}$.

Let us explain how we compute the cells $\mathcal{V}(1, \mathbf{q})$ and $A[1, \mathbf{q}]$ first, for every $\mathbf{q} \in \{0, 1, \ldots, \lceil n/\epsilon \rceil\}^m$. For each vector $\mathbf{q}$: we identify a multiset $\mathbf{d} \in \mathcal{D}_1$ with $\mathbf{d}' = \mathbf{q}$, that maximizes the value of bidder 1, *if such a multiset exists*. If it *does* exist, we store the value in $\mathcal{V}(1, \mathbf{q})$ and a partial allocation $\{\,(1, \mathbf{d})\,\}$ in $A[1, \mathbf{q}]$. Otherwise, we set $\mathcal{V}(1, \mathbf{q}) = 0$ and $A[1, \mathbf{q}] = \{\,(1, \emptyset)\,\}$. This procedure is described on the left of Figure 1. On the right of the figure, a generalization of this computation is described for bidder $i+1$, once the values $\mathcal{V}(i, \mathbf{q})$ and $A[i, \mathbf{q}]$ are known for all $\mathbf{q} \in \{0, 1, \ldots, \lceil n/\epsilon \rceil\}^m$.

The size of each table 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. For its optimality, we argue that the algorithm optimizes the welfare over a superset of feasible solutions $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ *for the initial problem instance;* indeed, for every good $\ell = 1, \ldots, m$ 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 $\sum_i \left\lfloor \frac{x_i(\ell) \cdot n}{\varepsilon \cdot s_\ell} \right\rfloor \le \lceil \frac{n}{\varepsilon} \rceil = s'_\ell$. Thus, $\mathbf{X}$ is *also feasible for the rounded problem instance* and inspected by the dynamic programming algorithm.

For the approximate violation of the supply constraints, let $\mathbf{X}$ be an output allocation, *for the original problem instance*. For any good $\ell = 1, \ldots, m$ we have: $\sum_i \frac{n \cdot x_i(\ell)}{\varepsilon \cdot s_\ell} \le \sum_i \left\lfloor \frac{n \cdot x_i(\ell)}{\varepsilon \cdot s_\ell} \right\rfloor + n$. The rounded version of $\mathbf{X}$ is feasible for the rounded instance, thus, $\sum_i \left\lfloor \frac{n \cdot x_i(\ell)}{\varepsilon \cdot s_\ell} \right\rfloor \le \lceil \frac{n}{\varepsilon} \rceil \le \frac{n}{\varepsilon} + 1$. Then, $\sum_i \frac{n \cdot x_i(\ell)}{\varepsilon \cdot s_\ell} \le \frac{n}{\varepsilon} + 1 + n$, which yields $\sum_i x_i(\ell) \le (1+2\varepsilon) s_\ell$.

Note that the algorithm is *exact*, in that it grants every bidder a multiset from his demand set (or none). This feature is essential in proving that the algorithm is MIR and, thus, that it yields a truthful mechanism:

Theorem 1 *Multi-unit combinatorial auctions with multi-minded bidders and a constant number of distinct goods admit a truthful* FPTAS*, that optimizes the social welfare fully, while violating the supplies of goods at most by a factor* $(1+\epsilon)$*, for any fixed* $\epsilon > 0$.

Computational Hardness Note that this problem is strongly **NP**-hard, when we do not allow violation of supply constraints and $m \ge 2$ [Chekuri and Khanna, 2000]. It is well known that if a problem is strongly **NP**-hard, it does not admit an FPTAS, unless **P=NP**, see, e.g., [Vazirani, 2003]. Also the assumption that m is a fixed constant is necessary. Otherwise the problem is hard to approximate in polynomial time within $m^{1/2-\epsilon}$, for any $\epsilon > 0$ [Lehmann *et al.*, 2002]. By applying a simple reduction from [Lehmann *et al.*, 2002], from the MAXIMUM INDEPENDENT SET PROBLEM [Hastad, 1996], we can establish the same approximation hardness, even if we allow supply violation, when m is non-constant.

An Application: Multi-Dimensional Knapsack Our FPTAS can be applied to the *Multi-dimensional Knapsack Problem* [Chekuri and Khanna, 2000] (MDKP), where the knapsack has a constant number $m = O(1)$ of distinct compartments, each of capacity s_ℓ, $\ell \in [m]$. 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 $i \in [n]$ is represented by a vector $\mathbf{d}_i = (d_i(1), \ldots, d_i(m))$ and has a value v_i. Thus, each such object corresponds exactly to a *single-parameter* bidder with valuation function $v_i(\mathbf{d}) \equiv v_i$, if $\mathbf{d} \ge \mathbf{d}_i$, and $v_i(\mathbf{d}) = 0$ otherwise. Our FPTAS is applicable to the MDKP because, as mentioned, it is *exact* in that it allocates every bidder (i.e., "fits in the knapsack") either an exact alternative from his demand set, $\mathcal{D}_i$, or none. We can generalize the MDKP setting further, by handling any mix of *packing* and *covering* constraints (i.e., of any of the forms $\{\ge, \le\}$), for $m = O(1)$ dimensions and one *covering* or *packing* constraint per dimension.

for every $\mathbf{q} \in \{\, 0,\, 1,\, \ldots,\, \lceil n/\epsilon \rceil \,\}^m$ **do:**

1. $\mathcal{D}'_1 \longleftarrow \{\, \mathbf{d} \in \mathcal{D}_1 \,|\, \mathbf{d}' = \mathbf{q} \,\}$

2. if $\mathcal{D}'_1 \neq \emptyset$ **then**

1. $\mathbf{d} \longleftarrow \arg\max_{\mathbf{a} \in \mathcal{D}'_1} v_1(\mathbf{a})$

2. $\mathcal{V}(1,\, \mathbf{q}) \longleftarrow v_1(\mathbf{d})$

3. $A[1,\, \mathbf{q}] \longleftarrow \{\, (1, \mathbf{d}) \,\}$

3. else

1. $\mathcal{V}(1,\, \mathbf{q}) \longleftarrow 0.$

2. $A[1,\, \mathbf{q}] \longleftarrow \{\, (1, \mathbf{0}) \,\}$

for every $\mathbf{q} \in \{\, 0,\, 1,\, \ldots,\, \lceil n/\epsilon \rceil \,\}^m$ **do:**

1. $\mathcal{D}'_{i+1} \longleftarrow \{\, \mathbf{d} \in \mathcal{D}_{i+1} \,|\, \mathbf{d}' \leq \mathbf{q} \,\}$

2. if $\mathcal{D}'_{i+1} \neq \emptyset$ **then**

1. $\mathbf{d} \longleftarrow \arg \max_{\mathbf{a} \in \mathcal{D}'_{i+1}} \left[\, v_{i+1}(\mathbf{a}) + \mathcal{V}(i,\, \mathbf{q} - \mathbf{a}') \,\right]$

2. $\nu_{i+1} \longleftarrow v_{i+1}(\mathbf{d}) + \mathcal{V}(i,\, \mathbf{q} - \mathbf{d}')$

3. if $\nu_{i+1} > \mathcal{V}(i,\, \mathbf{q})$ **then**

1. $\mathcal{V}(i+1,\, \mathbf{q}) \longleftarrow \nu_{i+1}$

2. $A[i+1,\, \mathbf{q}] \longleftarrow A[i,\, \mathbf{q} - \mathbf{d}'] \cup \{\, (i+1, \mathbf{d}) \,\}$

4. else

1. $\mathcal{V}(i+1,\, \mathbf{q}) \longleftarrow \mathcal{V}(i,\, \mathbf{q})$

2. $A[i+1,\, \mathbf{q}] \longleftarrow A[i,\, \mathbf{q}] \cup \{\, (i+1, \mathbf{0}) \,\}$

3. else

1. $\mathcal{V}(i+1,\, \mathbf{q}) \longleftarrow \mathcal{V}(i,\, \mathbf{q})$

2. $A[i+1,\, \mathbf{q}] \longleftarrow A[i,\, \mathbf{q}] \cup \{\, (i+1, \mathbf{0}) \,\}$

Figure 1: Computation of the Dynamic Programming tables' entries, for bidder 1 (left) and for bidder $i+1$ (right).

4 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. The method's generalization for multiple goods yields a truthful PTAS for bidders with submodular valuation functions over multisets. The method *requires* as subroutine a MIR α-approximation algorithm $\mathcal{A}$, that, for an instance of the problem involving a *constant number* $t = O(1)$ of bidders and $m = O(1)$ goods, approximates the optimum welfare within factor α. Moreover, $\mathcal{A}$ is MIR with range $\mathcal{R}_{\mathcal{A}}$. The method executes iteratively this algorithm $\mathcal{A}$ for appropriately chosen combinations of a subset of *at most* t bidders and a subset of at most m distinct goods. In each iteration, it completes appropriately the partial allocation output by $\mathcal{A}$ and, in the very end, returns the best found complete allocation for the whole instance of n bidders and $m = O(1)$ goods. If $\mathcal{A}$ is MIR, the method can be shown to be MIR as well; moreover, it recovers a fraction at least $(\alpha^{-1} - \frac{m}{t+1})$ of the optimum welfare of the whole instance.

Let us explain first how the Dobzinski-Nisan method completes a partial allocation output by the presumed algorithm $\mathcal{A}$ described above. Subsequently, we will describe such an algorithm $\mathcal{A}$ for the case of bidders with submodular valuation functions. The method is described in detail in Figure 2. It determines first a set P_ℓ of possible supplies for each good $\ell = 1,\, \ldots,\, m$; these are powers of $(1 + 1/(2n))$ no larger than s_ℓ – and rounded downwards to the largest integer – inclusively of 0, as described in lines **1.1** and **1.2** of Figure 2. Subsequently, for every possible subset T of at most $t = O(1)$ bidders and for every combination of supplies of goods $(\chi_1,\, \ldots,\, \chi_m) \in (\, \times_{\ell=1}^m P_\ell \,)$, the method:

1. first executes the algorithm $\mathcal{A}$ for the bidders in T and for all m goods, each good with the "remaining" supply $s_\ell - \chi_\ell$, to obtain a partial allocation.
2. Then, it splits the other part of each good's supply, χ_ℓ, into at most $2n^2$ multi-unit bundles, of $\beta_\ell = \max\{\lfloor \chi_\ell/(2n^2) \rfloor, 1\}$ units per such bundle. It finds the optimum allocation of these bundles to the remaining $[n] \setminus T$ bidders.

The (optimal) allocation of the splitted multi-unit bundles of goods to bidders in $[n] \setminus T$ is carried out by 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 multi-unit bundles 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:

$$\mathcal{V}(i, \mathbf{q}) = \max_{\mathbf{q}' \leq \mathbf{q}} \Big(v_i(\, q'_1 \cdot \beta_1,\, \ldots,\, q'_m \cdot \beta_m\,) + \mathcal{V}(i-1, \mathbf{q} - \mathbf{q}') \Big),$$

where $\mathbf{q}' \leq \mathbf{q}$ is taken component-wise; i.e., maximization occurs over all $\mathbf{q}'$ with $q'(\ell) \leq q(\ell)$ for each $\ell = 1, \ldots, m$.

The following result quantifies the performance of the Dobzinski-Nisan method for multiple goods; its proof is a direct extension of a related result from [Dobzinski and Nisan, 2010], for the case of a single good.

Theorem 2 *Let $\mathcal{A}$ be a* MIR *algorithm, with time complexity $\mathcal{T}_{\mathcal{A}}(t, (s_1, \ldots, s_m))$ for t bidders and at most s_ℓ units from each good $\ell = 1, \ldots, m$. There exists a range of allocations, $\mathcal{R}$, such that the Dobzinski-Nisan Method 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 $t = O(1)$, and outputs an allocation with value at least a fraction $(\alpha^{-1} - \frac{m}{t+1})$ of the optimum social welfare.*

The Method's Range The (generalized) Dobzinski-Nisan method optimizes over a range of possible allocations, that can be described independently of the bidders' declarations. This range depends on the range of the algorithm $\mathcal{A}$ used within the method. Following [Dobzinski and Nisan, 2010] we describe the method's range as a set of $(\mathcal{R}_{\mathcal{A}}, \tau, \chi_1,\, \ldots,\, \chi_m)$-*round* allocations. Let us first describe the notion of a round allocation, for any $\tau \leq t$, for some fixed $t = O(1)$. An allocation $\mathbf{X}$ is $(\mathcal{R}_{\mathcal{A}}, \tau, \chi_1, \ldots, \chi_m)$-*round* if:

1. There exists a subset $T \subseteq [n]$ of $\tau = |T|$ bidders, such that the (partial) allocation $\mathbf{X}[T]$ is optimum in the range

1. **for** $\ell = 1, \ldots, m$ **do:**
   1. **define** $u_\ell := (1 + \frac{1}{2n})$
   2. **define** $P_\ell := \left\{0, 1, \lfloor u_\ell \rfloor, \lfloor u_\ell^2 \rfloor, \ldots, \left\lfloor u_\ell^{\lfloor \log_{u_\ell} s_\ell \rfloor} \right\rfloor, s_\ell \right\}$
2. **for** every subset $T \subseteq [n]$ of bidders, $|T| \leq t$, **do:**
   1. **for** every $(\chi_1, \ldots, \chi_m) \in \left(\times_{\ell=1}^m P_\ell \right)$ **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 multi-unit bundles among bidders $[n] \setminus T$.
3. **return** the best allocation found.

Figure 2: The Dobzinski-Nisan Method for multiple goods.

$\mathcal{R}_\mathcal{A}$ of algorithm $\mathcal{A}$, for the sub-instance involving bidders in T and $s_\ell - \chi_\ell$ units from each good $\ell = 1, \ldots, m$.

2. Each bidder $i \in [n] \setminus T$ obtains 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 $\ell = 1, \ldots, m$.

Then, for any fixed $t = O(1)$ and for a MIR algorithm $\mathcal{A}$ with range $\mathcal{R}_\mathcal{A}$ as discussed above, the range of the generalized Dobzinski-Nisan method contains *all* the possible $(\mathcal{R}_\mathcal{A}, \tau, \chi_1, \ldots, \chi_m)$-*round* allocations, for any $\tau \leq t$ and $(\chi_1, \ldots, \chi_m) \in (\times_{\ell=1}^m P_\ell)$, where P_ℓ is defined as in step **1.2** of the method, in Figure 2. Quite elaborate as this description may be, it prescribes exactly the set of allocations that the Dobzinski-Nisan method optimizes upon. By the results of [Nisan and Ronen, 2007], we can obtain a truthful mechanism simply by pairing it with VCG payments.

A PTAS for Multi-Minded Bidders The generalized Dobzinski-Nisan method for multiple distinct goods can be applied directly in the setting of (multi-parameter) multi-minded bidders, that we studied in Section 3, to give us a PTAS that respects the supply constraints of the goods. A result of [Dobzinski and Nisan, 2010] rules out the possibility of FPTAS, when the supply constraints are fully respected, thus, the PTAS is best possible in this setting. Let $k = \max_i |\mathcal{D}_i|$ denote the maximum size of any bidder's demand set. For $m = O(1)$ goods and for any number $t = O(1)$ bidders, the *optimum assignment* can be found exhaustively in time polynomial in $\log s_\ell$, $\ell = 1, \ldots, s$, and m: there are exactly $O(k^t)$ cases to be examined exhaustively, so that the optimum is found. Then, by fixing any $\epsilon > 0$, we can use this (trivially MIR) algorithm $\mathcal{A}$ within the procedure of Figure 2, for $t = m \cdot \frac{1+\epsilon}{\epsilon} - 1$ and $\alpha = 1$. By Theorem 2, we thus obtain a PTAS that, complementarily to the developments of Section 3, approximates the optimum social welfare within factor $(1 + \epsilon)$ and respects the supply constraints.

A PTAS for Submodular Bidders We consider bidders with submodular valuation functions over multisets in $\mathcal{U}$:

Definition 3 *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 $v(\mathbf{x} + \mathbf{e}_\ell) - v(\mathbf{x}) \geq v(\mathbf{y} + \mathbf{e}_\ell) - v(\mathbf{y})$.*

These valuation functions, being exponentially large to describe compactly, are accessed by the algorithm through *value queries*; i.e., the algorithm asks the bidders for their value, for each particular multiset that it needs to process.

The MIR algorithm $\mathcal{A}$ needed by the Dobzinski-Nisan method optimizes over a range of allocations, defined as follows. For any $\epsilon > 0$, define $\delta = 1 + \epsilon$; the range of $\mathcal{A}$ contains all allocations that assign bidders multi-unit bundles of each good $\ell \in [m]$, where each bundle has cardinality equal to an integral power of δ. The specified range can be examined exhaustively in polynomial time; to find a welfare-maximizing allocation, we can try $O(\log_\delta s_\ell)$ cases per good, for each of $t - 1$ bidders, given a fixed bidder for assigning the remaining units. Trying all possible multi-unit bundle assignments of a specific good ℓ – and for all possible choices of a "remainders" bidder – is $O\left(t(\log_\delta s_\ell)^{t-1}\right)$. Because for every allocation of a specific good we need to examine all possible allocations for the remaining $m - 1$ goods, the overall complexity is $O\left(t^m (\log_\delta \max_\ell s_\ell)^{(t-1)m}\right)$, which is polynomially bounded for $t = O(1)$ and $m = (1)$. We show in [Krysta *et al.*, 2013] that – for a constant number of bidders and goods – this algorithm $\mathcal{A}$ is in fact an FPTAS, w.r.t. welfare; it yields a truthful PTAS when used within the Dobzinski-Nisan method.

Theorem 3 *Multi-unit combinatorial auctions with a constant number of distinct goods and bidders with submodular valuation functions admit a truthful* PTAS.

5 Conclusions

We elaborated on MIR polynomial-time deterministic mechanisms, for approximate social welfare maximization in multi-unit combinatorial auctions. The 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 **(ii)** we used a generalization of a single-good multi-unit allocation method proposed in [Dobzinski and Nisan, 2010]. Our developments are best possible in terms of time-efficient approximation, as follows by relevant hardness results.

References

[Archer *et al.*, 2003] A. Archer, C. H. Papadimitriou, K. Talwar, and É. Tardos. An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. *Internet Mathematics*, 1(2):129–150, 2003.

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

[Briest *et al.*, 2005] P. Briest, P. Krysta, and B. Vöcking. Approximation Techniques for Utilitarian Mechanism Design. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 39–48, 2005.

[Chekuri and Khanna, 2000] C. Chekuri and S. Khanna. A PTAS for the multiple knapsack problem. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'00)*, pages 213–222. SIAM, 2000.

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

[Daniely *et al.*, 2014] A. Daniely, M. Schapira, and G. Shahaf. Inapproximability of Truthful Mechanisms via Generalizations of the VC Dimension. *CoRR*, abs/1412.6265, 2014. To appear in the Proceedings of ACM STOC 2015.

[Dobzinski and Dughmi, 2009] S. Dobzinski and S. Dughmi. On the Power of Randomization in Algorithmic Mechanism Design. In *50th Annual IEEE Symposium on Foundations of Computer Science (FOCS'09)*, pages 505–514. IEEE Computer Society, 2009.

[Dobzinski and Nisan, 2007] S. Dobzinski and N. Nisan. Mechanisms for Multi-Unit Auctions. In J. K. MacKie-Mason, D. C. Parkes, and P. Resnick, editors, *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007)*, pages 346–351. ACM, 2007.

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

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

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

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

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

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

[Krysta *et al.*, 2013] P. Krysta, O. Telelis, and C. Ventre. Mechanisms for Multi-Unit Combinatorial Auctions with a Few Distinct Goods. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 691–698, 2013. *Full version available at:* http://goo.gl/tjtdYf.

[Lehmann *et al.*, 1999] D. J. Lehmann, L. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proceedings of the 1st ACM Conference on Electronic Commerce (EC'99)*, pages 96–102. ACM, 1999.

[Lehmann *et al.*, 2002] 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.

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

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

[Mirrokni *et al.*, 2008] V. S. Mirrokni, M. Schapira, and J. Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings 9th ACM Conference on Electronic Commerce (ACM EC)*, pages 70–77, 2008.

[Mu'alem and Nisan, 2002] A. Mu'alem and N. Nisan. Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI)*, pages 379–384, 2002.

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

[Nisan and Ronen, 1999] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC'99)*, pages 129–140. ACM, 1999.

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

[Nisan, 2014] N. Nisan. Algorithmic Mechanism Design, Trough the Lens of Multi-Unit Auctions. In R. Aumann and S. Hart, editors, *Handbook of Game Theory*, volume IV. Elsevier North-Holland, 2014.

[Vazirani, 2003] V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2003.

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

[Vöcking, 2012] B. Vöcking. A Universally Truthful Approximation Scheme for Multi-Unit Auctions. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 846–855. SIAM, 2012.