Truthful Mechanisms for Exhibitions

George Christodoulou[1,2], Khaled Elbassioni[1], and Mahmoud Fouz[3]

[1] Max-Planck-Institut für Informatik, Saarbrücken, Germany;
[2] Cluster of Excellence "Multimodal Computing and Interaction", Universität des Saarlandes
[3] FR Informatik, Universität des Saarlandes, Saarbrücken, Germany;
gchristo,elbassio,fmahmoud@mpi-inf.mpg.de

Abstract. We consider the following combinatorial auction: Given a range space $(U, \mathcal{R})$, and m bidders interested in buying only ranges in $\mathcal{R}$, each bidder j declares her bid $b_j : \mathcal{R} \rightarrow \mathbb{R}_+$. We give a deterministic truthful mechanism, when the valuations are single-minded: when $\mathcal{R}$ is a collection of fat objects (respectively, axis-aligned rectangles) in the plane, there is a truthful mechanism with a $1+\epsilon$- (respectively, $\lceil \log n \rceil$)-approximation of the social welfare (where n is an upper bound on the maximum integral coordinate of each rectangle). We also consider the non-single-minded case, and design a randomized truthful-in-expectation mechanism with approximation guarantee $O(1)$ (respectively, $O(\log m)$).

1 Introduction

In a combinatorial auction, there are m bidders competing on a finite set of k items for sale. The preferences of a player over the different subsets of items are expressed via a valuation function, that assigns to every subset of items a non-negative real value. An important, well-studied objective is to allocate the items to the bidders in a way that maximizes the *social welfare*, i.e., the sum of the valuations of the players on the allocated subsets.

We consider auctions where the bundles of items of interest have a geometric interpretation, i.e., they form connected geometric objects. Our work is inspired by the following applications: An owner of a space considers renting her land to exhibitors who will participate in a certain exhibition (say a computer show). The exhibitors bid on certain subsets of the space, and the owner has to decide which parts of the land she should allocate to which bidder, and how she should charge each winner. Typically, the owner may only allow bidders to bid on regions of certain shape (say squares or rectangles), and each bidder has a bid that depends on the location of the region (for instance, central regions or regions close to the entry of the whole exhibition can be more valuable to the bidder). A similar situation arises in advertisements: there is a certain space on a screen that can be used for displaying ads. A number of bidders compete for the total space, and have their individual valuations for each region on the screen[4]. Again, it is natural to expect the regions of interest to be squares or rectangles.

[4] see, for instance, "http://www.milliondollarhomepage.com".

In order to capture the above scenarios, we consider the following combinatorial auction: Let $U \subseteq \mathbb{R}^2$ be a set of points in the plane. Given m customers who are interested in buying subsets of U, each customer declares her bids on certain subsets of U (called ranges). Based on the bids, the auctioneer has to decide a feasible allocation of subsets to customers, and a payment to be charged to every winner (who gets allocated a non-empty subset).

In the above applications, it is natural to consider the situation where the possible subsets are connected regions in the plane, and moreover, those that are axis-aligned, or have some sort of *fatness*, such as squares or discs.

All the participants in an auction are selfish agents whose only goal is to maximize their utility, i.e., they want to obtain the bundle of items that maximizes the value minus the price. Therefore, they will try to manipulate the mechanism by misreporting their true values if this will increase their utility. In order to neutralize the effects of selfishness, a standard desired property of a mechanism that determines the allocation and payment is *truthfulness*, or *incentive-compatibility*. We look for truthful mechanisms, where the best strategy of each bidder is to report his true valuation. At the same time, we are interested in maximizing the *social welfare*, i.e., the sum of the valuations of the winners on their allocated subsets. The celebrated VCG mechanism [7, 9, 18], achieves both goals, but it runs in exponential time for most interesting scenarios. We consider polynomial-time mechanisms that approximate the optimal social welfare. The quality of an allocation is measured by the ratio of its total value to the optimal total value.

An important, well-studied special class of valuation functions, is the class of *single-minded valuations* (SM), where each bidder is interested in obtaining a particular subset of items. Lehmann et al. [14] showed that when the bidders are single-minded, there exists a truthful auction with approximation ratio $\sqrt{k}$, where k is the number of items. They also proved that this is the best ratio possible, even disregarding strategic issues, like truthfulness, unless NP = ZPP.

Our main contribution is the design of truthful mechanisms with approximation guarantees for two natural geometric settings of single-minded bidders; where the regions are either *axis-aligned rectangles* or *fat* objects. Intuitively, fat objects do not contain long and skinny parts; they are a well-known generalization of many common geometric objects like discs and squares. For instance, in the aforementioned motivating example on exhibitions, the regions of interest are usually rectangular with bounded width-to-height ratio; hence, they are fat. Apart from these results for the single-minded case, we also provide randomized truthful-in-expectation mechanisms for the non-single-minded case (non-SM).

The paper is organized as follows. We discuss related work in Section 1.1, and we give an overview of our results along with a comparison to previous work in Section 1.2. In Section 2, we give a precise problem formulation and describe some preliminaries. Finally, in Sections 3 and 4, we describe our results in detail.

1.1 Related Work

The work that is most related to ours is that by Babaioff and Blumrosen [2]. Motivated by similar applications, the authors in [2] study the single-minded version of the above problem. They generalize the greedy mechanism of Lehmann et al. [14] to obtain truthful mechanisms whose approximation guarantees are parameterized by the *aspect ratio* of the regions under interest, defined as the maximum ratio between the diameter and width of any object[5]. Two different informational settings are considered in [2]. In the *Known Single-Minded model (KSM)*, the auctioneer knows the actual range of each bidder but not the true valuation. In the more general *Unknown Single-Minded model (USM)*, both the ranges and valuations are private information of each bidder. The truthful mechanisms obtained in [2] have approximation guarantees $O(R^{4/3})$ in the USM model and $O(R)$ in the KSM model, where R is the maximum aspect ratio of the objects. These approximation guarantees are improved to $O(R)$ in case of arbitrary rectangles in the USM model, and $O(\log R)$ in case of axis-parallel rectangles in the KSM model. The latter result is obtained by using the bid-monotonic algorithms of Khanna et al. [12], which we discuss below.

In [12], the authors consider the *rectangle packing* problem. Here, given a set of m weighted rectangles in the plane, the problem is to find a disjoint collection of at most p rectangles and maximum weight. Note that without the restriction on the cardinality of the set, the problem corresponds to assigning disjoint rectangles to customers such that the social welfare is maximized. The authors of [12] assume that n rectangles are given on an $n \times n$ grid (i.e., they have integral coordinates) and obtain an $O(\log n)$ approximation algorithm that runs in time $O(n^2 p + np \log n)$. In contrast, our approach (which is similar to the one in [1]) for rectangles in the SM case only assumes that the rectangles have integral x-coordinates in $[0, n]$[6]. It achieves an approximation ratio of $\lceil \log n \rceil$, and runs (with a slight modification for the stated variant) in time $O(m \log m + mp \log n)$, where m is the number of rectangles. Note that m could be much smaller than the 'width' n of the plane.

Our approach for the non-SM case is based on rounding the LP-relaxation for the social welfare maximization (SWM) problem and then resorting to the general results of Lavi and Swamy [13]. Motivated by secondary spectrum auctions, and independently from our work, Hoefer et al. [10] considered a more general setting for the non-SM case, in which the feasible allocations are determined by a *conflict graph*, which is assumed to have a small *inductive independence number*[7]. They obtain randomized truthful mechanisms based also on an LP formulation of an extension of the SWM problem, combined with the results in [13]. We note that the intersection graph of a set of fat objects has a small

[5] More precisely, the aspect ratio of an object is the ratio between the maximum distance between any two points in the object and the minimum length of a projection of the object along any direction; equivalently, it is the ratio between the diameters of the minimum enclosing and maximum enclosed disc.

[6] Actually, it is already sufficient to assume that the rectangles lie in $[0, n] \times \mathbb{R}$ and have a minimum width of at least 1.

[7] A graph is said in [10] to have an inductive independence number ρ if there exists an ordering on the vertices s.t., for each vertex v, the subgraph induced on the neighbors of v, that precede v in the order, has independence number at most ρ.

inductive independence number, and hence some of the results in [10] can also be adapted to our setting.

1.2 Results and Techniques

Our main results concern single-minded valuations. We show that there is a truthful mechanism with a $(1+\epsilon)$-approximation of the social welfare, provided that the interesting regions (that the customers bid on) are fat ranges in $[0,1]^d$. This result is best possible since the SWM problem is already NP-hard for the setting of fat ranges [8, 11]. When the interesting regions are axis-aligned rectangles with integral x-coordinates, the approximation ratio will be $\lceil \log n \rceil$, assuming an upper bound of n on the maximum integral x-coordinate for each rectangle. We remark that, when $n = poly(m)$, getting a better bound than $O(\log n)$ will mean to get a better approximation guarantee than $O(\log m)$ for the SWM problem for rectangle ranges, which is a standing open problem[8].

Theorem 1. *Let $\epsilon > 0$ be an arbitrary constant. There is a polytime deterministic truthful mechanism, in the single-minded case, with approximation ratio $1+\epsilon$ for β-fat ranges with $\beta = O(1)$ and running time $m^{O(\epsilon^{-d+1})}$. For axis-aligned rectangles, there is a truthful mechanism, in the single-minded case, with approximation guarantee $\lceil \log n \rceil$ and $O(m \log(mn))$ running time.*

Here, β-fatness is a measure for how fat the object is [6]; for instance, a disc has a constant fatness, while a line segment has unbounded fatness. A precise definition is given in Section 2.

The results of Theorem 1 improve some of the results in [2]. For the case of constant aspect ratio $R = O(1)$, the mechanism of [2] achieves a constant approximation for the KSM information model, while we are able to show a truthful PTAS in the more general case of the USM model. It is important to emphasize that we consider only the USM model here; the mechanism is not aware of neither the true sizes and places of the objects nor their true values. In particular, for the fat ranges case, we assume that we are given a priori a square of size L in which all figures are guaranteed to lie inside. We also assume that the fatness parameter β is a known constant. For axis-parallel rectangles, we assume that all possible rectangles have integral coordinates, ranging from 1 to n along one of the principal directions. In that sense, our USM model is a bit weaker than the USM model in [2], but much stronger than the KSM model, as once we know the exact sizes and positions of the rectangles, we can compute the premises required in our model. It is not hard to see that the definition we use for fatness, given precisely below, is more general than bounding the aspect ration; in other words, $R = O(1)$ implies $O(1)$-fatness. Thus, we get a truthful PTAS when $R = O(1)$. For the case of axis-parallel rectangles, we strengthen the $O(\log R)$ result for the KSM model in [2]. In particular, in the above USM model, we get a truthful mechanism with a $\lceil \log n \rceil$-approximation ratio.

[8] Recently, Chalermsook and Chuzhoy [5] gave an $O(\log \log m)$-approx. algorithm for unit valuations. However, their result does not seem to extend to general valuations.

For the analysis of the algorithms in the single-minded case, we introduce a general framework for mechanism design and give sufficient conditions for the *monotonicity* of mechanisms that follow the framework. Roughly speaking, monotonicity is essentially a sufficient condition for truthfulness and ensures that when a range is shrunk or its value increased, then it remains in the solution if it was there before the change. The framework captures all algorithms that first decompose the instance into several (smaller) instances, solve each instance independently, and then return the best of all these solutions. Both for the fat ranges and the rectangles, the algorithms presented fit the framework and are shown to satisfy the sufficient conditions for monotonicity. The framework might also be of independent interest to show monotonicity of other mechanisms.

For the fat-ranges case, we modify Chan's algorithm [6] for packing d-dimensional fat ranges to get a monotone PTAS. For the rectangles case, we use a natural decomposition technique to partition the set of rectangles into different 'levels'. An instance is then formed by taking all rectangles of one level and, in addition, extensions of all rectangles of higher levels. Each instance is then solved by projecting all its rectangles to one line and solving the corresponding interval packing problem by dynamic programming to optimality. Although in [12], the authors also decompose the problem to instances of the interval packing problem, the decomposition we use is simpler and yields a better running time.

For the case of general valuations, that we refer to as the *non-single-minded* (non-SM) case, the results of Lavi and Swamy [13] allow us to derive randomized truthful-in-expectation mechanisms from LP-rounding algorithms for approximating the optimum social welfare. By developing such rounding algorithms, we obtain truthful-in-expectation mechanisms with the following guarantees.

Theorem 2. *There is a polytime truthful-in-expectation mechanism, in the non-SM case, with approx. ratios* 4β *for* β*-fat ranges, and* $4 \log m$ *for rectangles.*

2 Problem definition and preliminaries

There is an extensive literature on the design of truthful mechanisms for combinatorial auctions. In this section, we will only give the basic definitions needed in this paper. For an excellent introduction, we refer the reader to the book by Nisan et al. [15] and the references therein.

Let $[m] = \{1, \ldots, m\}$. Let $(U, \mathcal{R})$ be a range space, defined by a set U and a collection $\mathcal{R}$ of subsets of U, called *ranges*. Given m bidders, we assume that bidder j declares her bid $b_j(r)$ over every possible range $r \in \mathcal{R}$. The true valuation of the bidder j on range r will be denoted by $v_j(r)$. It is naturally assumed that the empty range $\emptyset \in \mathcal{R}$, and that $v_j(\emptyset) = 0$ for all j. In the single-minded case, each bidder is only interested in a single range $r_j \in \mathcal{R}$, that is, $v_j(r) = v_j(r_j)$ for all $r \supseteq r_j$ and $v_j(r) = 0$ for all $r \not\supseteq r_j$.

In this paper, U will be a set of points in d-dim. Euclidean space, and $\mathcal{R}$ a collection of connected regions. By the *size* of a region, we mean the side length of its smallest enclosing hypercube. More specifically, we will consider rectangles with integral x-coordinates ranging from 1 to n and *fat* objects in $[0, 1]^d$.

There are several definitions of fatness in the literature [3, 6, 17]. We use the following definition by Chan [6]. Recall that a *box* is the generalization of a square to higher dimensions. In what follows we assume that boxes are axis-aligned.

Definition 1 ([6]). *Let $\beta > 0$ be a constant. A collection C of ranges is β-*fat *if for any ℓ and any box B of size ℓ, we can choose β points $H(B)$ s.t. every range intersecting B and has size at least ℓ contains one point in $H(B)$.*

E.g., axis-aligned squares have fatness of $\beta = 4$ since for any box B of size ℓ and any such square S of size at least ℓ that intersects B, S must contain one of the four corners of B. On the other hand, line segments have unbounded fatness.

The social welfare maximization problem (SWM) is to find the optimum *integer* solution of the following linear program:

$$\max \sum_{\substack{j\in[m] \\ r\in\mathcal{R}}} v_j(r)x_{j,r} \tag{P}$$

$$\text{s.t.} \sum_{\substack{j\in[m] \\ r\in\mathcal{R}:u\in r}} x_{j,r} \leq 1 \qquad \text{for all } u \in U \tag{1}$$

$$\sum_{r\in\mathcal{R}} x_{j,r} = 1 \qquad \text{for all } j \in [m] \tag{2}$$

$$x_{j,r} \geq 0 \qquad \text{for all } j \in [m], r \in \mathcal{R}.$$

Informally, we want to find an allocation that maximizes the total valuations (i.e., the social welfare) while making sure that (1) each item is only assigned to one bidder and (2) each bidder is only assigned at most one range.

A *mechanism* takes as an input the set of bids $\{b_j : j \in [m]\}$ and outputs (i) a feasible allocation, that is, a 0/1-vector $\tilde{x}$ (or in other words an assignment of a (possibly empty) range r_j to each bidder $j \in [m]$) satisfying (1) and (2); and (ii) a payment $p : [m] \to \mathbb{R}_+$, that is, an amount $p_j \geq 0$ that is charged to bidder j, for all $j \in [m]$. The mechanism is said to satisfy *individual rationality* if it results in a non-negative utility for each bidder, that is, $v_j(r_j) - p_j \geq 0$, if the mechanism allocates r_j to bidder j, for $j = [m]$. For that, we assume that the bids are also single-minded, i.e., similar to the valuations, for all bidders $j \in [m]$ there exists a range $r'_j \in \mathcal{R}$ such that for all $r \supseteq r'_j$ we have $b_j(r) = b_j(r'_j)$ and for all $r \not\supseteq r'_j$, $b_j(r) = 0$. Thus, we can specify b_j by a range r_j and $b_j(r_j)$ and can write $(r_j, b_j(r_j))$ for a bid. The mechanism is said to be *truthful* if a bidder cannot improve his utility, under the mechanism, by bidding something different from his true valuation, regardless of the other players' bids. Formally, if the mechanism outputs the allocation-payment (r_j, p_j) for bidder $j \in [m]$, given the vector of bids (b_j, b_{-j}), and it outputs $(\hat{r}_j, \hat{p}_j)$ given the vector of bids (v_j, b_{-j}), where b_{-j} denotes the vector of bids of all other bidders $j' \neq j$, then the mechanism will be truthful if it satisfies $v_j(\hat{r}_j) - \hat{p}_j \geq v_j(r_j) - p_j$, for all j and all v_j, b_j, b_{-j}.

For the single-minded case, a sufficient condition for truthfulness is *monotonicity* and *critical payment*. The following formulation is adopted from [15, Lemma 11.9] and slightly adjusted to our notation.

Lemma 1 ([15]). *A mechanism for single-minded bidders in which losers pay 0 is incentive compatible if it satisfies the following two conditions: (i)* Monotonicity: *A bidder j who wins with bid $(r_j, b_j(r_j))$ keeps winning for any $b'_j > b_j$ and for any $r \subset r_j$ (for any fixed settings of the other bids); (ii)* Critical Payment: *A bidder j who wins pays the minimum value needed for winning: the infimum of all values b such that (r_j, b) still wins.*

We note that usually it is easy to define prices that satisfy the critical payment condition. Thus, the main difficulty in obtaining truthful mechanisms lies in ensuring monotonicity.

For randomized mechanisms, the mechanism is said to be *truthful-in-expectation* if $\mathbb{E}[v_j(\hat{r}_j) - \hat{p}_j] \geq \mathbb{E}[v_j(r_j) - p_j]$ for all $j \in [m]$ and all v_j, b_j, b_{-j}, where the expectation is taken over the random choices made by the algorithm.

The approximation guarantee of the mechanism is the ratio between the social welfare given by the mechanism $\sum_{j\in[m]} v_j(r_j)$ and the optimal social welfare, defined as the optimal integral solution of LP (P). Lavi and Swamy [13] showed that an "LP-based" approximation algorithm for SWM can be used to get a truthful-in-expectation mechanism of the same approximation guarantee.

Theorem 3 ([13]). *An LP-based α-approximation algorithm for the SWM problem can be used to obtain a truthful-in-expectation mechanism with approximation guarantee α.*

3 The single-minded case

In this section, we introduce a general framework for mechanism design in the single-minded case and give sufficient conditions for monotonicity of such mechanisms. We will then apply our framework to the two cases where the ranges are fat objects and rectangles.

3.1 A general framework for monotone mechanisms

The SWM problem can be thought of as finding a maximum-weight packing among a set of ranges with given weights $b(1), \ldots, b(m)$. Several existing algorithms for solving the packing problem can be put into the following framework:

(F1) k ordered instances $\mathcal{R}_1, \ldots, \mathcal{R}_k$ are obtained from the original input $\mathcal{R}$;
(F2) an algorithm $\mathcal{A}$ is used to solve each instance $\mathcal{R}_i$ independently, returning a packing $\mathcal{R}'_i$;
(F3) the packing $\mathcal{R}' = \operatorname{argmax}_i\{b(\mathcal{R}'_i)\}$ with maximum weight is returned, where ties are broken according to the order of the instances.

The following lemma describes sufficient conditions for such an algorithm to be monotone. In these conditions, we will consider changing one range or its corresponding bid, while all other ranges and bids are kept the same.

Lemma 2. *An algorithm that satisfies (F1)-(F3) is monotone if it satisfies further the following conditions (with $\mathcal{A}$ denoting the algorithm from (F2)):*

(C1) If a range is shrunk or its bid increased, the order of the instances is unaffected and no new instances are created.
(C2) Algorithm $\mathcal{A}$ is monotone.
(C3) If a range r is shrunk or its bid increased, then for all instances $\mathcal{R}_i$, the total weight of the solution returned by $\mathcal{A}$ on $\mathcal{R}_i$ can not decrease if r is contained in the solution (after the change) and remains the same otherwise.

Remark 1. Note that (C3) is always satisfied if we assume that the algorithm $\mathcal{A}$ is (1) *monotone in its output value*, i.e., if a range is shrunk or its bid increased, $\mathcal{A}$ returns a solution of total weight at least as large as the weight of the solution returned before the change, and (2) for all instances $\mathcal{R}_i$, changing range r can only result in the addition or drop of r from instance $\mathcal{R}_i$, i.e., all other ranges remain in exactly the same instances as before.

In the next subsections, we apply this general framework to the special cases of ranges of bounded fatness and axis-aligned rectangles, to prove Theorem 1.

3.2 The fat ranges case

We describe now how to modify Chan's algorithm [6] for packing m d-dimensional fat ranges to get a monotone PTAS for this case. By scaling, we can assume that all objects lie in $[0,1]^d$. We define the size of a range r to be the size of the smallest bounding box $B(r)$. We first sketch Chan's algorithm, which follows the framework (F1)-(F3). The idea is to choose the instances $\mathcal{R}_j$ such that they can be solved to optimality using dynamic programming. For that, the space is divided recursively by the following procedure: start with a hypercube containing all the ranges, then partition every non-empty hypercube recursively into 2^d equally sized hypercubes. Such a division is represented by means of a 2^d-dimensional tree (for $d = 2$ it is called a *quadtree*). Each instance $\mathcal{R}_j$ is obtained from a (different) shifted version of this basic division by including all ranges into $\mathcal{R}_j$ that are 'large' w.r.t. the smallest cell in which they are contained. Each instance is then solved to optimality and the one of maximum value is returned.

The dynamic program (DP) introduces for each cell $\mathcal{C}$ and disjoint subcollection $\mathcal{B}$ of the ranges crossing the boundary of $\mathcal{C}$ the table entry $pack[\mathcal{C},\mathcal{B}]$, which is defined as the maximum weight of a subcollection $\mathcal{B}'$ of ranges that lie completely inside $\mathcal{C}$, such that $\mathcal{B}\cup\mathcal{B}'$ is a disjoint collection. For a collection $\mathcal{B}'$ of ranges, let $\mathcal{B}'|_{\partial\mathcal{C}}$ be the subsets of ranges from $\mathcal{B}'$ crossing the boundary of $\mathcal{C}$. Let $\mathcal{C}_1,\ldots,\mathcal{C}_{2^d}$ denote the children of cell $\mathcal{C}$ in the tree. The table is filled bottom-up by the easily verifiable formula: $pack[\mathcal{C},\mathcal{B}] = \max_{\mathcal{B}'}\left(\sum_{i=1}^{2^d} pack[\mathcal{C}_i,(\mathcal{B}'\cup\mathcal{B})|_{\partial\mathcal{C}_i}] + b(\mathcal{B}')\right)$, where $b(\mathcal{B}')$ is the sum of bids of all ranges in $\mathcal{B}'$, and where the maximum is over all disjoint subcollections $\mathcal{B}' \subseteq \bigcup_i \mathcal{R}|_{\partial\mathcal{C}_i} \setminus \mathcal{R}|_{\partial\mathcal{C}}$ s.t. $\mathcal{B}'\cup\mathcal{B}$ is a disjoint collection.

The problem with such a straightforward approach is that the size of the table might be too large since we might have to consider an exponential number of disjoint subcollections $\mathcal{B}$ for some cell $\mathcal{C}$. To overcome this problem, we only consider ranges that are large with respect to their (smallest) enclosing cell.

Definition 2 ([6]). *A range of size ℓ is* k-aligned *if it is inside a tree cell of size at most $k\ell$.*

The key observation now is that when all objects are β-fat and k-aligned, the packing problem can be solved *exactly* in polytime using dynamic programming. The reason is that any feasible solution to the packing problem can have at most $K = 2\beta dk^{d-1}$ objects that intersect the boundary of any cell and there are only $O(m)$ many relevant cells in the bottom-up approach. Thus, for every such cell we consider only collections $\mathcal{B}'$ with $|\mathcal{B}'| \leq 2^d K$ ranges. Hence, the table will have at most $m^{O(K)}$ entries. Note that K is a constant if k and d are constants. In summary, we have the following lemma.

Lemma 3 ([6]). *If all ranges in C are β-fat and k-aligned, then the packing problem can be solved in $m^{O(\beta dk^{d-1})}$ time.*

Chan now considers $O(k)$ shifts of the basic division, each of which defines an instance $\mathcal{R}_j$ by removing all ranges that are not k-aligned w.r.t. that shift. Each instance is solved to optimality by the above DP, and the packing of maximum value is returned. Choosing k to be roughly d/ϵ, a $1+\epsilon$-approximation is achieved.

Theorem 4 ([6]). *Given a collection of m $O(1)$-fat objects in $\mathbb{R}^d$, the above algorithm gives a $1+\epsilon$-approximation to the packing problem in $m^{O(\epsilon^{-d+1})}$ time.*

A truthful algorithm. In order to make the algorithm monotone, we show how to ensure that all conditions of Lemma 2 are met. First, we order the instances given by the shifts in an arbitrary, but fixed order. Second, we need to ensure that, for each shifted instance, the (quadtree) partitioning is *independent* of the ranges. This can be achieved by assuming that all ranges in the instance lie in a fixed range, say $[0,1)^d$. So the initial cell will be the unit hypercube. We keep partitioning a cell until it does not contain any range in its interior. Even though the number of cells is not polynomial in m, it is easy to see that one can still implement the DP in polytime (assuming fixed d), by 'zooming' into the relevant cells (those that are intersected by at least one range).

Since the partitioning of the instances is independent from the ranges, (C1) is immediately satisfied. In order to satisfy (C2) and (C3) we do *not* drop any ranges from the instances even if they are not k-aligned, and modify the DP accordingly. In particular, we now have a table entry $pack[\mathcal{C}, \mathcal{B}]$ for every cell $\mathcal{C}$ and every disjoint subcollection $\mathcal{B} \subseteq \mathcal{R}|_{\mathcal{C}}$ of ranges that either cross the boundary of $\mathcal{C}$ or are completely inside $\mathcal{C}$, *and* have cardinality at most K. In contrast to the DP before, the maximum in the above recurrence is now over all subcollections of ranges in $\mathcal{H}(\mathcal{C}, \mathcal{B})$, where $\mathcal{H}(\mathcal{C}, \mathcal{B})$ is the set of disjoint subcollections of ranges $\mathcal{B}' \subseteq \mathcal{R}|_{\mathcal{C}} \setminus \mathcal{R}|_{\partial\mathcal{C}}$ such that $\mathcal{B}' \cup \mathcal{B}$ is a disjoint collection and $|\mathcal{B}'| \leq K$:

$$pack[\mathcal{C}, \mathcal{B}] = \max_{\mathcal{B}' \in \mathcal{H}(\mathcal{C}, \mathcal{B})} \Big(\sum_{i=1}^{2^d} pack[\mathcal{C}_i, (\mathcal{B}' \cup \mathcal{B})|_{\mathcal{C}_i}] + b(\mathcal{B}') \Big). \tag{3}$$

Note that in contrast to the original DP we do not only consider ranges that lie on the boundary of the child cells of $\mathcal{C}$ but all ranges that lie (completely) in $\mathcal{C}$. Despite these modifications, we still have $m^{O(K)}$ relevant table entries.

Clearly, the solution computed by the modified DP is no worse than the solution computed by the original DP (since we just search over a larger space). Still, we have to ensure that the DP is monotone (i.e., satisfies (C2)). For that we specify how to break ties during the computation of $pack[\mathcal{C}, \mathcal{B}]$ in (3). Consider an arbitrary but fixed order on all ranges: $r_1, \ldots, r_m$. Note that during the computation (3), we consider all optimal subcollections $\mathcal{B}' \in \mathcal{H}(\mathcal{C}, \mathcal{B})$. Whenever there is more than one such subcollection we return the subcollection that forms the lexicographically smallest ordered sequence of ranges.

Lemma 4. *Let $\epsilon > 0$. The modified algorithm for fat ranges returns a $1+\epsilon$-approximate, monotone solution and runs in time $m^{O(\epsilon^{-d+1})}$.*

In order to ensure that the algorithm is truthful, we still have to specify the critical payment of each winner. We use a similar payment as in the VCG mechanism (see e.g. [15]). Let W_j be the maximum value returned by the DP over all instances if we remove the object r_j and W'_j the value if in addition to r_j we also remove all objects that intersect r_j. Clearly, $W'_j \leq W_j$. Then define the payment of bidder j to be $p_j = W_j - W'_j$. To see why this definition indeed yields the critical value for bidder j, note that if bidder j bids less, no solution of maximum value over all instances contains r_j, so bidder j does not win. On the other hand, if bidder j bids $p_j + \epsilon$ for any $\epsilon > 0$, every solution of maximum value must contain r_j, so bidder j wins for sure. Finally, individual rationality follows as in the VCG mechanism, since for a winner j, $b_j + W'_j \geq W_j$.

3.3 The rectangles case

We now consider the case when the given objects are axis-aligned rectangles with integral x-coordinates assuming values from 1 to n. We assume that n is a power of 2; otherwise we can extend the interval to the nearest power of 2. In that case, it is easy to verify that we achieve an approximation ratio of $\lceil \log n \rceil$.

Any set of rectangles $r_1, \ldots, r_m$ can be partitioned into at most $k = \log m$ sets, which can be interpreted as lying at different "levels". The rectangles at level $\ell \in [k]$ can be further partitioned into $h = 2^{\ell-1}$ sets, such that all rectangles in one set intersect one vertical line, while every pair of rectangles from two different sets are disjoint (see, e.g., [1]). We will denote by $\mathcal{L}_\ell$ the set of these $2^{\ell-1}$ vertical lines at level ℓ. To make this decomposition independent of the rectangles themselves, which is needed for monotonicity, we give the decomposition explicitly. We define $\mathcal{L}_i$ to be the set of vertical lines $x = n/2^i, x = 3 \cdot n/2^i, \ldots, x = (2^i - 1) \cdot n/2^i$. Clearly, we get $k = \log n$ levels. We assign each rectangle to the smallest level such that one of its lines intersects the rectangle. Note that each rectangle is intersected by exactly one vertical line of its level.

Given a set of rectangles $\{r_1, \ldots, r_m\}$, we iterate the following two steps, for $\ell \in [k]$: First, we project every rectangle whose level is at least ℓ onto the nearest vertical line at level ℓ. This gives a set of intervals on each vertical line in $\mathcal{L}_\ell$. The bid of a bidder j on a given projection is the same as on the original rectangle. Second, we apply a monotone DP for the maximum weight-independent set problem on the set of intervals on each vertical line in $\mathcal{L}_\ell$. The resulting independent set of intervals corresponds to a disjoint set of rectangles $\mathcal{R}_\ell$. At the end we output the set of rectangles $\mathcal{R}_\ell$ that achieves the highest total bid, i.e., maximizes $\sum_{r\in\mathcal{R}_\ell} b_j(r)$. In case of ties, we choose the $\mathcal{R}_\ell$ at the *lowest* level. It is straightforward to verify that this algorithm is a special case of the general framework described in Section 3.1, so monotonicity follows from Lemma 2. Finally each winner is charged the critical payment of the projection that made him a winner.

Lemma 5. *The above procedure for rectangles is monotone, gives an approximation guarantee of* $\log n$ *and has a running time of* $O(m\log(mn))$.

4 The Non-Single-Minded case

In this section, we show that in the non-single-minded case, we get a truthful-in-expectation mechanism with approximation guarantee $O(1)$ (resp., $O(\log m)$) for fat objects (resp., axis-aligned rectangles.) The idea is to use Theorem 3. For that, we show that the integrality gap of the LP (P) is $O(1)$ for fat objects and $O(\log m)$ for rectangles (and there is an algorithm verifying this).

We give a randomized algorithm that with high probability returns an integral solution $\tilde{x}$ such that $b^T\tilde{x} = \Omega(b^T x)$, where b denotes the (column) vector of bids and x denotes the optimal fractional solution This algorithm can be derandomized using the work of [16]. Combining the results yields Theorem 2.

The fat ranges case. Assume all ranges are β-fat. We apply randomized rounding with alteration (see, e.g., [4]). Let $(x_{j,r} : j \in [m], r \in \mathcal{R})$ be any feasible solution for (P). Let $\gamma \in (0,1)$ be a constant to be specified later. We define the rounded solution $\tilde{x}$ by its *winner set* W, obtained by the following procedure. First, for every bidder j, we choose a range $r = r_j$ with probability $\gamma x_{j,r}$ if r is non-empty and with probability $1-\gamma(1-x_{j,r})$ if $r = \emptyset$. Second, let $W = \emptyset$, and $r_1, r_2, \ldots, r_m$ be the ranges selected in Step 1 in *non-increasing* order of size. For $j = 1, \ldots, m$, if r_j does not intersect any range r_i with $i \in W$, add j to W.

By construction, the set $(r_j : j \in W)$ is a valid allocation. It remains to prove the approximation guarantee of $O(1)$.

Lemma 6. $\mathbb{E}[\sum_{j\in W} b_{j,r_j}] \geq \frac{1}{4\beta}\sum_{j\in[m],r\in\mathcal{R}} b_{j,r}x_{j,r}$.

The rectangles case. We use a similar technique as in the previous section. However, since rectangles can have unbounded fatness, we refine the previous technique by using the levelwise decomposition described in Section 3.3. Let $(x_{j,r} : j \in [m], r \in \mathcal{R})$ be any feasible solution for (P), and $\gamma = \frac{1}{2}$. We define the rounded solution $\tilde{x}$ by its winner set W, obtained by the following procedure.

1. For every bidder j, we choose a rectangle $r = r_j$ with probability $\gamma x_{j,r}$ if r is non-empty and with probability $1 - \gamma(1 - x_{j,r})$ if $r = \emptyset$.
2. Let $r_1, r_2, \ldots, r_m$ be the rectangles selected in Step 1, ordered from top to bottom (breaking ties arbitrarily).
3. Consider a levelwise decomposition of these rectangles into $\log m$ levels. Pick a level $i \in \{1, 2, \ldots, \log m\}$ at random. Let S be the set of rectangles in level i.

4. Let $W = \emptyset$. For $j \in S$, if r_j does not intersect any r_i, $i \in W$, add j to W.

Lemma 7. $\mathbb{E}[\sum_{j \in W} b_{j,r_j}] \geq \frac{1}{4 \log m} \sum_{j \in [m], r \in \mathcal{R}} b_{j,r} x_{j,r}$.

Acknowledgments. We thank Rene Sitters for suggesting the application, and Rajiv Raman and Rene Sitters for helpful discussions.

Bibliography

[1] Agarwal, P.K., van Kreveld, M.J., Suri, S.: Label placement by maximum independent set in rectangles. Comput. Geom. 11(3-4), 209–218 (1998)
[2] Babaioff, M., Blumrosen, L.: Computationally-feasible truthful auctions for convex bundles. GEB 63(2), 588–620 (2008)
[3] de Berg, M., Gudmundsson, J., Katz, M., Levcopoulos, C., Overmars, M., van der Stappen, A.F.: TSP with Neighborhoods of varying size. J. of Algorithms 57, 22–36 (2005)
[4] Chakrabarti, A., Chekuri, C., Gupta, A., Kumar, A.: Approximation algorithms for the unsplittable flow problem. Algorithmica 47(1), 53–78 (2007)
[5] Chalermsook, P., Chuzhoy, J.: Maximum independent set of rectangles. In: SODA. pp. 892–901 (2009)
[6] Chan, T.M.: Polynomial-time approximation schemes for packing and piercing fat objects. J. Algorithms 46(2), 178–189 (2003)
[7] Clarke, E.H.: Multipart pricing of public goods. Public Choice pp. 17–33 (1971)
[8] Fowler, R.J., Paterson, M.S., Tanimoto, S.L.: Optimal packing and covering in the plane are NP-complete. IPL 12, 133–137 (1981)
[9] Groves, T.: Incentives in teams. Econometrica pp. 617–631 (1973)
[10] Hoefer, M., Kesselheim, T., Vöcking, B.: Approximation algorithms for secondary spectrum auctions. CoRR abs/1007.5032 (2010)
[11] Imai, H., Asano, T.: Finding the connected components and a maximum clique of an intersection graph of rectangles in the plane. J. Algorithms 4, 310–323 (1983)
[12] Khanna, S., Muthukrishnan, S., Paterson, M.: On approximating rectangle tiling and packing. In: SODA. pp. 384–393 (1998)
[13] Lavi, R., Swamy, C.: Truthful and near-optimal mechanism design via linear programming. In: FOCS. pp. 595–604 (2005)
[14] Lehmann, D., O'Callaghan, L.I., Shoham, Y.: Truth revelation in approximately efficient combinatorial auctions. In: JACM 49(5). pp. 577–602 (2002)
[15] Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press (2007)
[16] Sivakumar, D.: Algorithmic derandomization via complexity theory. In: STOC. pp. 619–626 (2002)
[17] Stappen, A.: Motion Planning amidst Fat Obstacles. Ph.D. dissertation, Utrecht University, the Netherlands (1994)
[18] Vickrey, W.: Counterspeculation, auctions and competitive sealed tenders. Journal of Finance pp. 8–37 (1961)