Computationally-Feasible Truthful Auctions for Convex Bundles

Moshe Babaioff and Liad Blumrosen*

School of Engineering and Computer Science
The Hebrew University of Jerusalem, Jerusalem 91904, Israel

Abstract

In many economic settings, convex figures on the plane are for sale. For example, one might want to sell advertising space on a newspaper page. Selfish agents must be motivated to report their true values for the figures as well as to report the true figures. Moreover, an approximation algorithm should be used for guaranteeing a reasonable solution for the underlying NP-complete problem. We present truthful mechanisms that guarantee a certain fraction of the social welfare, as a function of a measure on the geometric diversity of the shapes. We give the first approximation algorithm for packing arbitrary weighted compact convex figures. We use this algorithm, and variants of existing algorithms, to create polynomial-time truthful mechanisms that approximate the social welfare. We show that each mechanism achieves the best approximation over all the mechanisms of its kind. We also study different models of information and a discrete model, where players bid for sets of predefined building blocks.

1 Introduction

The intersection between Micro-Economic theory and Computer-Science theory raises many new questions. These questions were studied recently by researchers from both disciplines (see, e.g., the surveys in [13, 6]). A leading example for a problem in this intersection is the *Combinatorial Auction* problem. In a combinatorial auction, a finite set of heterogenous items is for sale, and each selfish agent has a valuation for every subset of these items. As the auction designers, we try to find an allocation of the items among the agents that maximizes the "social welfare" (i.e., a set of disjoint packages that maximizes the sum of valuations) or at least to find a good approximation.

In this paper, we study a variant of combinatorial auctions: e.g., a newspaper wants to sell an advertising space on a newspaper page. Agents might have different preferences about their desired space: the size of the ad, its location on the page, whether its

*Email: {mosheb,liad}@cs.huji.ac.il.

The authors are grateful to Noam Nisan for many helpful discussions, and to Ron Lavi for his comments on an earlier draft.

figure is rectangular, square, or elliptic etc. Each agent submits a bid for her favorite figure, and we try to find the allocation that maximizes the social welfare. The underlying packing problem is known to be NP-hard, even for the very simple case of packing 2×2 squares ([15]). Thus, we settle for a computationally-efficient mechanisms that approximate the social welfare.

In our model, the plane ($\Re^2$) is for sale. Let N be a finite set of agents ($|N| = n$). Each agent has a private non-negative valuation $v_i \in \Re^+$ for a single compact convex[1] figure s_i and for any figure that contains it (other figures have a valuation of 0). Every agent submits a bid for her desired figure (e.g., an advertiser might want the lower half of the first page in a newspaper). After receiving all bids, the auctioneer determines a set of winning agents with disjoint figures and the payment that each agent should pay. Note that the agents demand figures in fixed locations in the plane, and the auctioneer cannot translate or rotate them. Bidding for convex figures is common in many real-life scenarios that involve "geometric" bids, e.g., selling real-estate lots, newspaper ads and spectrum licenses in different locations. In most existing real-estate or advertising auctions, agents are forced to bid on predefined disjoint figures. This might result in inefficient allocations that can be avoided by allowing the agents to bid for arbitrary figures (which in turn makes the computational problem harder).

Note that the problem addressed in this paper is more than just finding an algorithm with a good approximation ratio. The agents in our model are selfish, and they may report untruthful information if this is beneficial for them. We want to design *incentive-compatible* mechanisms, in which each agent uses a strategy that is best for her own selfish interest (a *dominant strategy*), and yet, a certain approximation for the social welfare is guaranteed. A general scheme for achieving a welfare-maximizing incentive-compatible mechanism is the family of Vickrey-Clarke-Groves (VCG) mechanisms (see [11] for a review). However, for implementing such mechanisms we must allocate the goods optimally (otherwise it will not be truthful [14]). Thus, since finding the optimal allocation is an *NP*-hard problem, we must find incentive-compatible mechanisms that are non-VCG. Almost all the non-VCG mechanisms currently known are for models where the agents hold a single secret value (*single-parameter* models)[2].

In light of these impossibilities, we assume that each agent is interested in a single figure. Lehmann et al. [10] initiated the study of the *Single-Minded Bidders model* for combinatorial auctions. Our model is unique since the bids have some common geometric properties (e.g., convexity), and also because we actually auction an infinite (even uncountable) number of goods (the points in the plane).

We differentiate between two models of information (similar to the differentiation done in [12, 3]) . In the first model, the auctioneer knows which figure each agent wants, but does not know how much this agent is willing to pay for this figure. This model is the *Known Single-Minded (KSM) model*, and this is indeed a "single-parameter" model. In the second model, called the *Unknown Single-Minded (USM) model*, both the figures and the values are unknown. In the KSM model we should motivate the agents to truthfully declare their true values, where in the USM model the

[1] Actually we prove our results for a more general model that allows non-convex bids as well, with some more general restrictions.

[2] Recent results ([9]) show that in many reasonable settings, essentially no IC mechanisms exists for multi-parameter models, except the family of weighted VCG mechanisms.

agents might submit untruthful bids both for their desired figures and their values.[3]

Another differentiation we make is between a continuous and a discrete model. In the *continuous model*, each agent is interested in an arbitrary compact convex figure in $\Re^2$. For example, if a piece of land is for sale, each agent can draw an arbitrary figure on the map. In the *discrete model*, the plane contains predefined atomic building blocks (or tiles), and each agent is restricted to bids for a set of building blocks which are contained in some convex figure. For example, if we wanted to resell the island of Manhattan, the basic building blocks would be street blocks bounded between two consecutive streets and two consecutive avenues. These blocks are typically convex, though not necessarily rectangular (e.g., because of the diagonal Broadway Avenue).

Related Work:

Our research relates to a sub-field of Micro-Economics called *Mechanism Design*, which studies ways to design mechanisms that encourage agents to behave in a way that results in some desired global properties (see, e.g., [11]). Nisan and Ronen [14] introduced this concept to CS by the name *Algorithmic Mechanism Design*.

Weighted packing of rectangles in the plane was studied in several papers. Hochbaum and Maass [7] proposed a *shifting strategy* for a special case of square packing, and generalizations for arbitrary squares appear in, e.g., [5, 4]. Khanna et al. [8] used similar methods in a model where axis-parallel rectangles lie in a $n \times n$ grid. They presented an algorithm that runs in polynomial time and achieves an $O(log(n))$-approximation for the optimal welfare. However, it is an open question whether a better approximation (and in particular, a constant approximation) exists for this problem.

Our Contribution:

We measure the quality of the approximations achieved by mechanisms in our model according to an *aspect ratio* R, which measures how diverse are the dimensions of the figures demanded by the agents. R is defined as the ratio between the maximal *diameter* of a figure and the minimal *width* of a figure (formally defined in Section 2). For different families of figures, we construct IC mechanisms, either for the KSM or for the USM models. This mechanisms are also *individually-rational*, i.e., agents will not pay more than they value the figure they receive (if any). Therefore, our approximation improves as the dimensions of the figures become closer.[4]

We study three different families of figures: compact convex figures, rectangles, and axis-parallel rectangles. For convex figures in the USM model, we achieve an $O(R^{\frac{4}{3}})$-approximation to the social welfare. If the bids are restricted to rectangles (not necessarily axis-parallel), we achieve a better approximation of $O(R)$.

If the agents bid for axis-parallel rectangles, we can use a slight modification of the algorithm due to Khanna et al. [8] to design an IC mechanism that achieves an approximation ratio of $O(log(R))$ (the best known approximation ratio for this problem).

We also present a novel allocation algorithm that achieves an $O(R)$-approximation for packing arbitrary compact convex figures, and as far as we know this is the first

[3] Note that the USM model does not fit into the "single-parameter" definition, since the agents have both their values and their figures as their secret data.

[4] For instance, when all figures are disks with the same radius up to a constant, our mechanisms achieve a constant approximation.

approximation algorithm for this problem. We use this algorithm for constructing an IC mechanism, with the same approximation ratio, for the KSM model.

The incentive-compatible mechanisms we present for the USM model are based on a family of greedy algorithms presented by Lehmann et al. [10]. For standard combinatorial auctions, Lehmann et al. normalized the values by $|S|^{\alpha}$, where $|S|$ is the number of items in the bundle S and α is some real constant, and then run a simple greedy allocation algorithm on the normalized values. They showed that choosing $\alpha = \frac{1}{2}$ guarantees the best polynomial approximation ratio (unless $NP = ZPP$).

We present mechanisms, called α-*greedy mechanisms*, that normalize the values using the *geometric area* of the figures. That is, for $\alpha \in \Re$ we assign a normalized value of $\frac{v}{q^{\alpha}}$ to a bid with a value v for a figure with a geometric area q. We show that, somewhat surprisingly, for compact convex figures the optimal value for α is $\frac{1}{3}$, resulting an $O(R^{\frac{4}{3}})$-approximation. The difference between the results of Lehmann et al. and ours derives from the different divisibility properties of packages in the two models. In their model, a finite set of goods is traded, and for a package to intersect "many" disjoint packages, its size must be "large". However, in our continuous model, a small package can intersect many disjoint packages.

For our discrete model, we present a mechanism that achieves an $O(R^{\frac{4}{3}})$ approximation. However, if the ratio between the minimal width of a figure and the sizes of the building blocks (we denote by Q) is smaller than the aspect ratio R, we can achieve a better approximation of $O(R \cdot Q^{\alpha^*})$, by running the α-greedy algorithm with $\alpha^* = \frac{log(R)}{2log(R)+log(Q)}$.[5]

The paper's organization: Section 2 describes our model. Section 3 describes our results for the USM model, both for the continuous case and the discrete case. Section 4 presents the results for the KSM model, and Section 5 concludes with a discussion of future work. All proofs are given in the full version of the paper ([2]).

2 Model

Let B denote the family (set) of *bids* (figure-value pairs) of the agents[6], that is $B = \{(s_i, v_i) | i \in N\}$. Let F denote the family of agents figures, that is $F = \{s_i | i \in N\}$. Given a family of bids B, we aim to maximize the *social welfare*, i.e., find a collection of non-conflicting bids (bids for disjoint figures) that maximizes the sum of valuations. For a subset $C \subseteq N$ of agents with disjoint figures, denote the value of C by $V(C) = \sum_{i \in C} v_i$. We denote the set of disjoint figures that achieves the maximal welfare by OPT (to simplify the notation we assume that there are no ties, so there is a single optimal solution), i.e.,

$$V(OPT) = \max_{C \subseteq N | \ \forall i,j \in C \ s_i \cap s_j = \emptyset} V(C)$$

[5] Thus, if one can embed the goods of a traditional combinatorial auction as building-blocks in the plane, such that each agent bids for building-blocks contained in some convex figure, then our approximation scheme improves the approximation ratio achieved in [10].

[6] Since all the mechanisms we consider are truthful, we use the same notation for the secret information and the declared information (bid), except of the IC proofs.

Definition 1 *A mechanism consists of a pair of functions (G, P) where:*

– *G is an allocation scheme (rule) that assigns a figure in $\mathcal{T}$ (where $\mathcal{T}$ is the set of compact convex figures in the plane) to every agent such that the figures are disjoint, i.e. $G(B) \in \mathcal{T}^N$ and for every $i \neq j$ in N, $G_j(B) \cap G_i(B) = \emptyset$ (where we denote the figure received by agent i by $G_i(B)$).*

– *P is a payment scheme, i.e. for any B, $P(B) \in \Re^n$. Denote the payment paid by agent i by $P_i(B)$.*

All allocation rules we present in the paper, allocate to an agent either her requested figure or the empty figure. All the payment rules we consider are *normalized*, that is, a losing agent pays zero. Additionally, each agent pays a non-negative payment. We assume quasi-linear utilities and that the agents have no externalities (the utility for each agent does not depend on the packages received by the other agents), i.e., the utility of each agent i is $u_i(B) = v_i(G_i(B)) - P_i(B)$. The agents are rational, so each agent chooses a bid that maximizes her own utility.

A mechanism is *incentive-compatible (IC)* if declaring their true secret information is a dominant strategy for all the agents. In the KSM model, it means that for any set of values reported by the other agents, each agent cannot achieve a higher utility by reporting an untruthful value, i.e., $\forall i \;\; \forall B_{-i} \;\; \forall v_i' \;\; u_i((s_i, v_i), B_{-i}) \geq u_i((s_i, v_i'), B_{-i})$, where B_{-i} denote the family of all bids except i's bid . In the USM model, IC means that each agent's best strategy is to report *both* her figure and her value truthfully, regardless of the other agents' reports, i.e., $\forall i \;\; \forall B_{-i} \;\; \forall v_i', s_i' \;\; u_i((s_i, v_i), B_{-i}) \geq u_i((s_i', v_i'), B_{-i})$.

An incentive-compatible mechanism is also *individually rational (IR)* if for any agent i, bidding truthfully ensures him a non-negative utility. That is, $\forall i \;\; \forall B_{-i}$ $u_i((s_i, v_i), B_{-i}) \geq 0$.

Geometric definitions:

We state our approximation bounds as functions of few geometric properties of the family of figures the agents bid for. We use standard definitions of *diameter* and *width* of compact figures in $\Re^2$:

The *diameter* d_z of a compact set z is the maximal distance between any two points in the set, i.e. $d_z = \max_{p_1, p_2 \in z} ||p_1 - p_2||_2$ ($||p_1 - p_2||_2$ is the Euclidean distance between p_1 and p_2). The *width* w_z of a compact set z is the minimal distance between the closest pair of parallel lines such that the convex set z lies between them.

Definition 2 *Given a family of figures F in $\Re^2$, the maximal diameter L is the maximal diameter of a figure in F, and the minimal width W is the minimal width of a figure in F. The aspect ratio R is the ratio between the maximal diameter and the minimal width. That is, $L = \max_{z \in F} d_z$, $W = \min_{z \in F} w_z$, $R = \frac{L}{W}$.*

The aspect ratio describes how diverse is the family of figures with respect to the figures' diameter and width.[7] The approximations our mechanisms achieve are asymptotic functions of the aspect ratio R.

[7] For example, if all the figures are disks with the same radius, then $R = 1$. If we have disks of diameter 10 and 5×2 rectangles, then $R = \frac{10}{2} = 5$.

Denote the *geometric area* of a compact figure z by $q(z)$. We assume that the diameter, width and area of any agent's figure are polynomial-time computable. We also assume that given any two agent's figures, we can decide if the two figures are disjoint in polynomial time.[8].

The Discrete Model:
In the discrete model, there is a set of atomic building blocks (we call *tiles*) embedded in the plane. Each agent desires a bundle of tiles that are exactly the ones that are fully contained in some compact convex figure, and she reports this figure and her value for the set of tiles contained in it. We assume that all tiles have similar dimensions (specifically, each tile contains a disk of some positive diameter W_0 and its diameter is at most $2W_0$). Two agents are non conflicting if there is no tile which is *fully* contained in the two figures they report.

For a given family of bids, we define the *width-ratio* $Q = \frac{W}{W_0}$. The width-ratio gives an upper bound on the width of any figure, with respect to the size of the tiles. Clearly, we can assume that $Q \geq 1$.

The Greedy Mechanism:
Lehmann et al. [10] presented the following family of *greedy mechanisms* for combinatorial auctions:

Given a family of bids and some function f on the figures, such that f assigns a positive real value to any non empty figure (w.l.o.g. all figures are non empty), the *greedy allocation algorithm* picks an allocation *ALG*, by the following scheme:
Create a list of the bids sorted from high to low according to their values normalized by f (i.e., $\frac{v_1}{f(s_1)} \geq \frac{v_2}{f(s_2)} \geq \ldots \geq \frac{v_n}{f(s_n)}$).

While the list is not empty, choose a figure s_i for which the normalized value is highest in the remaining list (with a consistent tie breaking). Add i to the allocation ALG and update the list by removing all bids for figures that intersect s_i.

The specific algorithm is determined by the choice of the function (or norm) f. Lehmann et al. suggested using the norm $\frac{v_s}{|s|^\alpha}$ for combinatorial auctions, where $|s|$ is the size of the package s and α is some non-negative constant. We generalize this method for compact figures in $\Re^2$ and define the α-*greedy algorithm* to use the norm $\frac{v_z}{q(z)^\alpha}$, where $q(z)$ is the area of figure z.[9]

Definition 3 *The α-greedy mechanism is a mechanism which uses the α-greedy algorithm as its allocation scheme, where a winning agent i pays according to the following payment scheme:*

Let j be the first agent to win, among all the agents whose figures intersects agent i's figure, when the greedy algorithm runs without i. If such j exists, i pays $\frac{q(s_i)^\alpha \cdot v_j}{q(s_j)^\alpha}$, otherwise i pays 0. Losing agents pay 0.

The properties of the α-greedy mechanisms, proved in [10], also hold in our model:

[8] Note that for polygons the above assumptions hold, and that any compact convex figure can be approximated (as good as one wants) by a polygon.

[9] For example, the 0-greedy algorithm sorts the figures according to their values and the 1-greedy algorithm sorts the figures according to their value per unit of area.

Theorem 4 *(essentially due to [10]) For every α, the α-greedy mechanism is polynomial time, individually rational and incentive compatible for agents bidding for compact figures in the USM model.*

3 The Unknown Single-Minded Model

This section considers the problem of designing a polynomial-time, individually-rational and incentive-compatible mechanisms, which guarantee some fraction of the social efficiency for the USM model. We study three families of figures: convex figures and rectangles in the continuous model, and convex figures in the discrete model. We use the α-greedy mechanism to create mechanisms with the desired properties for the three families. For each family, we find the value of α that optimizes the asymptotic approximation for the social welfare, over all α-greedy mechanisms. For convex figures in the continuous model, we show that $\alpha = \frac{1}{3}$ achieves an $O(R^{\frac{4}{3}})$-approximation for the social welfare. For rectangles in the continuous model, we improve the above result and show that $\alpha = \frac{1}{2}$ achieves an $O(R)$-approximation for the social welfare. Finally, for convex figures in the discrete model, we show that a careful choise of α as a function of R and Q, gives an approximation ratio between $O(R^{\frac{4}{3}})$ and $O(R)$. The proofs of all above results are based on a single general result presented in the full version [2].

3.1 Convex Figures and Rectangles in the Continuous Model

The following lemma presents a lower bound on the approximation ratio that can be achieved by α-greedy mechanisms, by presenting two constructions that are in a sense the hardest inputs for these mechanisms.

Lemma 5 *The α-greedy mechanism for compact convex figures achieves an $\Omega(R^{2(1-\alpha)})$ approximation for any $\alpha < 1$, and an $\Omega(R^{1+\alpha})$-approximation for any $\alpha > 0$. Therefore, the α-greedy mechanism for compact convex figures achieves an $\Omega(R^{\frac{4}{3}})$-approximation.*

Proof sketch: Each of the lower bounds is achieved by a construction that can be built for any R large enough. The left part of Figure 1 illustrates the construction used to prove the $\Omega(R^{2(1-\alpha)})$ bound and the right part is used to prove the $\Omega(R^{1+\alpha})$ bound. In the left example, a large rectangle contains $\Theta(R^2)$ small disjoint squares with a side of length W. On the right example, a small disk intersects $\Theta(R)$ disjoint triangles, with equal area of $\Theta(WL)$. In both constructions, there is one figure z (filled by small vertical lines) that is chosen by the greedy mechanism, while the socially optimal mechanism chooses a family of disjoint figures (small horizontal lines), each intersects z. The value of z is chosen such that its normalized value $\frac{v_z}{q(z)^\alpha}$ is a bit greater than 1, and the rest of the figures have a normalized value of 1. The value for α that minimizes the worst case approximation is therefore $\frac{1}{3}$, yielding an $\Omega(R^{\frac{4}{3}})$ lower bound. □

Next, we show that the $\frac{1}{3}$-greedy mechanism achieves an $O(R^{\frac{4}{3}})$-approximation (which, by the above lemma, is the best over all the α-greedy mechanisms). To prove this result, we use few elementary geometric properties of convex figures. First, for any compact convex figure z, $q(z) = \Theta(d_z w_z)$. Additionally, the perimiter of z (denoted

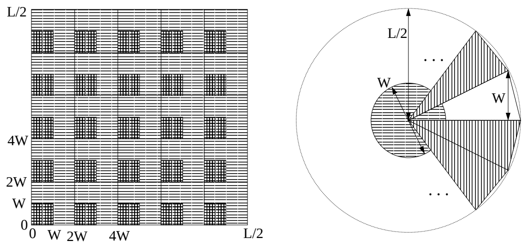

Figure 1: Approximation bounds for the α-greedy mechanism

by p_z) is contiguous, and $p_z = \Theta(d_z)$ (the constants in both cases are independent of the figure z). These properties are sufficient for the approximation to hold.

Theorem 6 *When agents bid for compact convex figures in the plane with an aspect ratio R, the* $\frac{1}{3}$ - greedy mechanism *achieves an* $O(R^{\frac{4}{3}})$*-approximation, and this is the best asymptotic approximation achievable by an* α*-greedy mechanism (for any* α*). This mechanism is individually rational and incentive compatible for the USM model, and it runs in polynomial time.*

Proof sketch: By Theorem 4 the mechanism is IR, IC and it runs in polynomial time. Next, we present the idea behind the proof of the $O(R^{\frac{4}{3}})$-approximation ratio. By the definition of the α-greedy algorithm, any figure is either a winner or intersects a winner. It is the hardest to prove the approximation bound if the set of winners in the optimal solution OPT is disjoint to the set of winners ALG picked by the α-greedy algorithm. We map each agent $x \in OPT$ to a winner $z \in ALG$ that intersects x. We then bound the sum of values of any disjoint set of agents' figures that intersects z, by partitioning them to figures that are contained in z and to figures that are not. We use Hölder inequality and some simple geometric properties of convex figures to show that the upper bounds for $\alpha = \frac{1}{3}$, for both the contained figures and the rest of the intersecting figures, match the lower bounds presented in Lemma 5. □

If agents are only interested in rectangles (not necessarily axis parallel), than we can derive a stronger result of an $O(R)$-approximation for the social welfare. While the construction of a large rectangle containing $\Omega(R^2)$ small rectangles (as presented in Lemma 5) is still possible, the second construction is not. For rectangles, it is impossible for a small rectangle to hit (intersect but not contain) many disjoint rectangles.

Theorem 7 *When agents bid for rectangles in the plane with an aspect ratio R, the* $\frac{1}{2}$ - greedy mechanism *achieves an* $O(R)$*-approximation for the social welfare, and this is the best asymptotic approximation achievable by an* α*-greedy mechanism (for any* α*).*[10] *This mechanism is individually rational and incentive compatible for the USM model, and it runs in polynomial time.*

[10]We actually prove a stronger statement. We show that $\Omega(R)$-approximation is the best over all the greedy mechanisms that sort the bids according to some function of the value and the area only (specifically, this includes the function $\frac{v_z}{q(z)^{\alpha}}$).

3.2 Convex Figures in the Discrete Model

We now turn to look at the discrete model. We first define the mechanism we use for the discrete model, we then present the mechanism properties.[11]

Definition 8 *The Discrete Model Greedy Mechanism is a mechanism that given bids for compact convex figures in* $\Re^2$ *in the discrete model, does the following: If* $Q \geq R$ *then it runs the* $\frac{1}{3}$-greedy mechanism, *and if* $Q \leq R$ *then it runs the* α^* - greedy mechanism *for* $\alpha^* = \frac{log(R)}{2log(R)+log(Q)}$.

Theorem 9 *Consider that the agents bid for compact convex figures in* $\Re^2$ *in the discrete model, with an aspect-ratio R and a width-ratio Q. Then, the* Discrete Model Greedy Mechanism *achieves an* $O(R^{\frac{4}{3}})$*-approximation for the social welfare. Moreover, when* $Q \leq R$ *it achieves a better approximation of* $O(R \cdot Q^{\alpha^*})$*. This mechanism achieves the best asymptotic approximation among all the mechanisms that choose* α *as a function of R and Q, and in particular it is asymptotically better than the* α*-greedy mechanism for any* α*. Additionally, the mechanism is IR and IC for the USM model, and it runs in polynomial time.*

4 The Known Single-Minded Model

In this section, we present polynomial-time mechanisms for different families of figures in the Known Single-Minded (KSM) model (where the auctioneer knows the desired figure of each agent, but does not know her value for the figure). We start by presenting an auction for general compact convex figures. We achieve an $O(R)$-approximation for the social welfare, which is better than the $O(R^{\frac{4}{3}})$-approximation that we proved for the USM model. Next, we present a mechanism (based on an algorithm from [8]), which gives an $O(log(R))$-approximation for axis-parallel rectangles.[12]

4.1 Mechanisms for Convex Figures

Consider the mechanism called the *"Classes-by-Area* 1*-Greedy Mechanism" (CBA-1G mechanism)* presented in Figure 2. This mechanism divides the bids of the agents to classes according to the figures' geometric area, runs a 1-greedy algorithm in each class, and allocates the figures to agents in the class that achieved the highest result. From the algorithmic aspect, this is the first algorithm for packing weighted convex figures that we know of, and it achieves an $O(R)$-approximation. We use this algorithm to construct a polynomial-time and IC mechanism with the same approximation ratio for the social welfare. The payments are exactly the "critical-values" for the agents, i.e., the minimal declaration for which they still win the auction.

[11] Note that an agent can manipulate the values of α by affecting R and Q. Therefore, for incentive compatibility, the mechanism is assumed to know the true values of R and Q.

[12] This approximation is exponentially better than the approximation ratio we achieve for this problem in the USM model and than the ratio we achieve for general convex figures in the KSM model.

The Classes-by-Area 1-Greedy (CBA-1G) Mechanism:

Allocation:

Step 1: Divide the given input to $m = 2log(R)$ classes according to their area. A figure s belongs to class c if $q(s) \in [W^2 \cdot 2^c, W^2 \cdot 2^{c+1})$ (for $c \in \{0, \ldots, m-1\}$).

Step 2: Perform the 1-greedy algorithm per each class. Denote the welfare achieved by class c by V^c.

Step 3: Output the allocation in the class c for which the 1-greedy algorithm achieved the highest welfare, i.e., $c \in argmax_{\tilde{c} \in \{0,\ldots,m-1\}} V^{\tilde{c}}$.

Payments:

Denote the winning class as class 1, and the class with the second-highest welfare as class 2. Let $V^1_{-i} = V^1 - v_i$, and let j be the first figure that intersects figure i and wins, when we run the greedy algorithm where agent i is removed. Let z_i be $\frac{v_j q(i)}{q(j)}$ if such j exists, and 0 otherwise.

A winning agent i pays: $\mathbf{P(i) = max\{V^2 - V^1_{-i}, z_i\}}$, and any losing agent pays 0.

Figure 2: A mechanism for selling arbitrary convex figures. This mechanism is incentive compatible in the KSM model and achieves an $O(R)$-approximation for the social welfare.

The payments in the CBA-1G mechanism are chosen as follows: to win the auction, each agent should be both a winner in her class and her class should beat all other classes. Bidding above the value z_i in the mechanism's description, is a necessary and sufficient condition for agent i to win in her class. However, if agent i bids below $V^2 - V^1_{-i}$ and still wins in her class, her class will definitely lose.

Theorem 10 *When the agents bid for compact convex figures in $\Re^2$ with an aspect ratio R, the* CBA-1G *mechanism achieves an $O(R)$-approximation. This mechanism is IR and IC for the KSM model*[13] *, and runs in polynomial time.*

Proof sketch: We show that the approximation ratio achieved in each class is $O(R_c)$, where R_c is the aspect ratio[14] in class c. Due to a general proposition we prove, the approximation ratio achieved by choosing the best class is $O(\sum_c R_c)$. Finally, we show that $\sum_c R_c = O(R)$, by dividing this sum to two geometric series. For proving IC, we show that the given payments are indeed the critical values for the agents, i.e. the smallest declarations for which they still win. □

4.2 Mechanisms for Axis-Aligned Rectangles

In the full version of this paper ([2]) we present an allocation algorithm, called the *Shifting Algorithm*, which is based on an algorithm by Khanna et al.([8]) with some minor changes. They studied a model where axis-aligned rectangles lie in an $n \times n$ array, and they proved an $O(log(n))$-approximation for the weighted packing problem. This approximation is the best approximation currently known for weighted packing of axis-parallel rectangles. Our algorithm gives an $O(log(R))$-approximation in a slightly more general model where the rectangles can lie in any axis-parallel location in the

[13] We observe that this mechanism is not IC in the USM model.

[14] I.e. the ratio between the maximal diameter and the minimal width of figures in this class.

plane. By carefully defining a payment scheme, we use this allocation rule for designing an IC polynomial-time mechanism achieving an $O(log(R))$-approximation for the social welfare. We call this mechanism the *Shifting Mechanism* and we summarize its properties in the following theorem:[15]

Theorem 11 *When the agents bid for axis-parallel rectangles in $\Re^2$ with an aspect ratio R, the* Shifting Mechanism *achieves an $O(log(R))$-approximation. This mechanism is individually rational and incentive compatible for the KSM model, and runs in polynomial time.*

5 Conclusion and Further Research

In this paper, we study auctions in which the agents bid for convex figures in the plane. We present mechanisms that run in polynomial time, in which the selfish players' best strategy is to send their true private data to the auctioneer. We suggest using the aspect ratio R, which measures how diverse are the dimensions of the figures, for analyzing the economic efficiency of the mechanisms.

In the KSM model, we were able to achieve the best approximation currently known for weighted axis-parallel rectangle packing ($log(R)$) in an IC mechanism. Lehmann et al. [10] showed that the best polynomial-time approximation for combinatorial auctions (for single minded bidders) can be achieved with an IC mechanism. On the other hand, recent results showed settings in which the optimal algorithmic approximation ratio cannot be achieved by IC mechanisms (see, e.g., [1, 9]). Whether such gap exists in our model is an interesting open question:

Open Problem: *Can the best polynomial-time approximation schemes for packing convex figures (general figures, rectangles, or axis-parallel rectangles) be implemented by incentive-compatible mechanisms?*

The mechanism for combinatorial auctions presented in [10] achieves the same approximation both for the USM model and the KSM model. Our results might indicate that, in our model, a gap exists between the approximation achievable in both information models. For general convex figures, the approximation we achieve in the KSM and the USM models are $O(R)$ and $O(R^{\frac{4}{3}})$, respectively. For axis-parallel rectangles, the gap in our results is even exponential.

Open Problem: *In settings where the agents are "single minded", is there a gap between the best approximation achievable in the KSM and in the USM models?*

We present some novel algorithmic results regarding packing of convex figures and arbitrary rectangles. We have not been able to show that these results are tight.

Open Problem: *Is there an $o(R)$-approximation scheme for packing general convex figures, or even for packing rectangles (not necessarily axis-parallel)?*

Our results may also be useful in deriving approximation results for the problem of packing weighted convex bodies in dimensions higher than two[16].

[15] An easy observation is that the Shifting Mechanism is not IC in the USM model. We also note that the Shifting Mechanism achieves an $\Omega(R)$-approximation for general rectangles (not necessarily axis-parallel).

[16] However, the economic interpretation of such auctions is not always clear.

References

[1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *symposium on foundations of computer science*, pages 482–491, 2001.

[2] M. Babaioff and L. Blumrosen. Computationally-feasible truthful auctions for convex bundles. Full version. Available from www.cs.huji.ac.il/˜ liad or ˜ mosheb.

[3] M. Babaioff and W. E. Walsh. Incentive-compatible, budget-balanced, yet highly effi cient auctions for supply chain formation. In *ACM Conference on Electronic Commerce*, pages 64–75, 2003. Extended version to appear in DSS, 2004.

[4] T. M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. In *J. Algorithms*, volume 46, pages 209–218, 2003.

[5] T. Erlebacj, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric graphs. In *SODA 2001*, pages 671–679.

[6] Papadimitriou C. H. Algorithms, games, and the internet. In *proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 749–753, 2001.

[7] D. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and vlsi. In *Journal of the ACM*, pages 130–136, 1985.

[8] S. Khanna, S. Muthukrishnan, and M. Paterson. On approximating rectangle tiling and packing. In *Symposium on Discrete Algorithms*, pages 384–393, 1998.

[9] Ron Lavi, Ahuva Mua'lem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *FOCS 03*, pages 574–583, 2003.

[10] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately effi cient combinatorial auctions. *Journal of the ACM*, 49(5):1–26, 2002.

[11] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[12] Ahuva Mua'lem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *AAAI-02*, 2002.

[13] Noam Nisan. Algorithms for selfi sh agents. In *16th Symposium on Theoretical Aspects of Computer Science*, 1999.

[14] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1/2):166–196, April/May 2001.

[15] M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147, 1998.