Auctioning a Cake: Truthful Auctions of Heterogeneous Divisible Goods

Yonatan Aumann
Department of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
aumann@cs.biu.ac.il

Yair Dombb
Department of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
yairdombb@gmail.com

Avinatan Hassidim
Department of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
avinatanh@gmail.com

ABSTRACT

We consider the problem of auctioning a one-dimensional continuously-divisible heterogeneous good (a.k.a. "the cake") among multiple agents. Applications include auctioning of time intervals, e.g. auctioning time for usage of a shared device, auctioning TV commercial slots, and more. Different agents may have different valuations for the different possible intervals, and the goal is to maximize the aggregate utility. Agents are self-interested and may misrepresent their true valuation functions, if this benefits them. Thus, we seek auctions that are *truthful*. Considering the case that each agent may obtain a single interval, the challenge is twofold, as we need to determine both where to slice the interval, and who gets which slice. The associated computational problem is NP-hard even under very restrictive assumptions. We consider two settings: discrete and continuous. In the discrete setting we are given a *sequence* of m indivisible elements $(e_1, \ldots, e_m)$, and the auction must allocate each agent a *consecutive subsequence* of the elements. For this setting we provide a truthful auctioning mechanism that approximates the optimal welfare to within a $\log m$ factor. The mechanism works for arbitrary monotone valuations. In the continuous setting we are given a continuous, infinitely divisible interval, and the auction must allocate each agent a sub-interval. The agents' valuations are non-atomic measures on the interval. For this setting we provide a truthful auctioning mechanism that approximates the optimal welfare to within a $O(\log n)$ factor (where n is the number of agents). Additionally, we provide a truthful 2-approximation mechanism for the case that all slices must be of some fixed size.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; G.2.1 [**Combinatorics**]: Combinatorial algorithms

General Terms

Algorithms, Economics

Keywords

Cake Cutting; Auctions; Resource Allocation

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

1. INTRODUCTION

Consider a common resource, e.g. high-end video editor, for which several agents require access. Different agents may attribute different utilities to using the equipment, and this utility may further vary between different time intervals (nights, mornings, Sundays). When and for how long should each agent get to use the common resource?

Similarly, consider a setting where commercials can be placed alongside the screening of some movie. Several advertisers wish to place their commercials, and are willing to pay for doing so. Naturally, different advertisers may have different preferences as to when, and for how long, they get to air their commercials (possibly depending on the exact content of the movie scenes), thus attributing different values to the different possible time slots. How should one decide who gets which time slot and for what price? What is the optimal way to determine the time slots?

In these cases, as in many others, a natural solution is to auction the common good. Auctioning, if done right, allows for the resource to go to the agent(s) that benefits from it the most, and for the best price. However, in the above examples the resource in consideration is not a single item, but rather "time": a *continuously divisible* and *heterogeneous* good. There are many ways to slice-out time (in fact, infinitely many), and each such partition may offer different utilities to the different agents. The question thus arises: *How does one auction time?* In this work we take the first steps in addressing this fundamental question. We focus on the traditional goal of maximizing utilitarian welfare – maximizing the total utility produced by the auction. Additionally, we require that the auction be *truthful*, as agents may misrepresent their valuation functions. We focus throughout on the case where each agent must get a single, contiguous, time interval.

It is interesting to compare this problem to the general problem of combinatorial auctions [9]. Auctioning time is similar to combinatorial auctions in that many items are auctioned at once, and participants place different values on *bundles* of items. However, time also contains a strong "geometric" element, absent in the general combinatorial setting; with time, the core good is an interval, not a set, and the auction needs to decide how to *slice* the good among the agents. Thus, on the one hand, there is the additional challenge of determining where to place the cuts. On the other hand, the set of *permissible allocations* is restricted to those that allocate a contiguous interval to each agent.

We note that the same framework also applies to auctioning other one dimensional continuously divisible goods, e.g. beach-front land. Following the rich literature on fair and efficient division of an interval (see e.g. [26] for a survey), in the remainder of the paper we refer to the auctioned interval as "the cake".

Results.

It was recently shown [2] that it is computationally hard to find the socially optimal allocation in the setting we consider here, even if the valuations are fully known, and even when imposing severe restrictions on the valuation functions. Hence, we seek approximations.

We consider two setting: discrete and continuous. In the discrete setting we are given a *sequence* of m indivisible elements $(e_1, \ldots, e_m)$, and the auction must allocate each agent a *consecutive subsequence* of these elements. For this setting we provide a truthful auctioning mechanism, which approximates the optimal welfare to within a $\log m$ factor. This mechanism works for arbitrary monotone valuations.

The continuous setting is the one presented above: we are given a continuous, infinitely divisible interval, and the auction must allocate each agent a sub-interval. The agents' valuations are non-atomic measures on the interval (in particular, they are additive). Note that in this case there are infinitely many possible sub-intervals, so a major challenge is to choose the right cuts to consider, while maintaining truthfulness. We provide a universally-truthful auctioning mechanism for the continuous setting that approximates the optimal welfare to within a $O(\log n)$ factor (where n is the number of agents).

We also consider a special case of the continuous setting, where there is some fixed interval length ℓ such that any allotted interval must be of exactly this length ℓ. This requirement may arise in many real world settings where the length of the allotted intervals is determined by external considerations, e.g. presentations at a conference. The problem remains NP-hard even under this restriction. Here we provide a truthful mechanism for this case that is guaranteed to output an allocation with welfare at least 1/2 of the optimum.

Communication. A major difficulty in combinatorial auctions is the communication requirements. The representation of individual valuations functions may be prohibitively large. Our mechanisms place very low communication requirements: each agent needs to provide its valuation for only $2 \cdot m$ intervals, in the discrete setting, and $O(n^2)$ intervals in the continuous setting.

Related Work.

The use of auctions for allocating goods (typically discrete and indivisible) has been mathematically studied for decades. In the past 15 years, these questions have also been intensively studied from a computational point of view, mainly under the title of Algorithmic Mechanism Design [23] (see also [25]). It turns out that the problem of finding an optimal allocation is computationally intractable in many settings. Therefore, many works have considered restricted cases, in which some assumptions are made on the bidders' possible valuations and/or on the set of permissible allocations [28, 29, 17].

There is a large body of work concerning auctions of *homogeneous* divisible goods [12, 30, 18, 16, 15]. In this setting, agents may have different valuations for different *amounts* of the good, and the auction needs to determine what amount each agent gets. The results from these works, however, do not carry over to our *heterogeneous* setting, where the valuations are not only a function of the amount, but also of the exact location of the interval within the good.

The problem of dividing a heterogeneous, continuously divisible good among different players is commonly referred to as "cake cutting". Cake cutting problems have been studied for over half a century, with the main focus traditionally being on *fairness* (see [4, 27] for surveys of this literature). In recent years the cake cutting setting has received increasing attention from the AI and multi-agent community, as a model for resource allocation among agents, see [26].

Aumann et al. [2] consider the problem of maximizing welfare in cake cutting with contiguous pieces, showing that the problem is NP-hard even for valuations that have a rather simple structure, and providing a constant-factor approximation algorithm; however, their algorithm does not extend to a truthful allocation mechanism. Bei et al. [3] study the problem of finding a contiguous pieces partition that maximizes welfare while maintaining proportionality. They show that the problem is NP-hard even for piecewise-constant valuation functions, and provide a PTAS for linear functions.

Truthful mechanisms for the cake cutting setting were considered by several previous works. Chen et al. [6] and Mossel and Tamuz [20] both consider the problem of devising truthful mechanisms aimed at producing fair allocations, providing polynomial time truthful mechanisms for various fairness criteria, and various classes of valuation functions. The setting of these works differs from ours in several ways: their goal is fairness while ours is aggregate welfare, their mechanism are without payments while ours may use payments, and they allow non-contiguous pieces while we focus on contiguous pieces.

Truthful cake cutting mechanisms without payments and non-contiguous pieces were also considered by Maya and Nisan [19]. They consider the case of two players with piecewise uniform valuations and show a tight bound of ≈ 0.93 on the utilitarian welfare approximation. Guo and Conitzer [13] and Han et al. [14] consider a similar setting, but where one needs to divide a *set* of divisible goods, each of which is assumed to be homogeneous.

Rothkopf et al. [28] consider the discrete setting considered here. However, unlike in our setting, they allow bidders to receive multiple subsequences, and assume that the values of such subsequences can be (linearly) added, allowing them to obtain an optimal allocation algorithm for this case.

Welfare maximization has been also studied in many other settings of multiagent resource allocations; for a survey of this literature see, e.g. [7].

Finally, it is worth noting that, in practice, auctions have been successfully applied for dividing one-dimensional goods such as TV airtime [22] and spectrum [8].

2. PRELIMINARIES AND NOTATION

We consider two setting: discrete and continuous. We now formally define each.

The Discrete Setting.

In the discrete setting there is a sequence of m items, $E = (e_1, \ldots, e_m)$ to be divided among a set of n bidders/agents, denoted by the integers $[n] = \{1, \ldots, n\}$. Each bidder has a *valuation function*, v_i attributing a non-negative value to each possible subsequence of E. We assume nothing on the valuation functions v_i other than being monotone, i.e. that if I and I' are subsequences of E, with $I \subseteq I'$ (as sets), then $v_i(I) \leq v_i(I')$, for all i. A mechanism allocates non-overlapping subsequences of E to some or all of the n bidders, such that each bidder gets at most one subsequence. Thus, each bidder i is allocated a subsequence I_i (which may be empty).

Note that we do not require the allocation to allot the entire resource; this assumption is known as *free disposal*.

The Continuous Setting.

In the continuous setting, the good at hand is the interval $[0, 1]$, to be divided among the n bidders/agents. Each bidder's *valuation function*, v_i, is a non-atomic measure on $[0, 1]$ (in particular, v_i is additive). The mechanism allocates non-overlapping sub-intervals of $[0, 1]$ (to which we sometimes refer as "plots") to the n bidders, such that each bidder gets at most one subinterval.

A useful consequence of the non-atomicity of the v_i's is that we can ignore the endpoints of intervals, as for any $a < b$ and any $i \in [n]$ it implies that $v_i([a, b]) = v_i((a, b))$. We will therefore consider only allocations that give bidders *open* intervals, thus completely avoiding the issue of deciding who gets the boundary point between two adjacent intervals.

Social Welfare.

In this work we devise mechanisms aiming to maximize the utilitarian social welfare; that is, mechanisms that for any set of valuations $\{v_i(\cdot)\}_{i\in[n]}$ try to find an allocation $\{I_i\}_{i\in[n]}$ with the sum $\sum_{i\in[n]} v_i(I_i)$ as high as possible. However, as we state in Theorem 1 below, finding an allocation that maximizes the utilitarian welfare is computationally intractable. We therefore focus on finding mechanisms that *approximate* the welfare. We say that a mechanism $\mathcal{M}$ achieves an α-approximation of the optimal welfare if it outputs an allocation $\{I_i^{\mathcal{M}}\}_{i\in[n]}$ such that for any possible allocation $\{I_i\}_{i\in[n]}$ it holds that

$$\alpha \cdot \sum_{i\in[n]} v_i(I_i^{\mathcal{M}}) \geq \sum_{i\in[n]} v_i(I_i) \ .$$

Communication.

In general, a bidder's valuation function is her own private data, unknown to the mechanism. This poses two challenges for the mechanism designer. The first is finding a way for the bidders to inform the mechanism about their valuations using a reasonable amount of communication.[1] The second is dealing with the problem that bidders may opt to misreport their valuations in order to increase their gains.

The mechanisms presented in this work all operate with very limited communication. More precisely, in the discrete setting, each bidder must only communicate the values of $2m$ specific subsequences. In continuous setting, each bidder needs to communicate $O(n^2)$ such valuations, and $O(n)$ cut points. For continuous setting with a fixed interval size ℓ, each bidder must communicate $O(1/\ell)$ values. Thus, the amount of communication required by our mechanisms is very modest, and is completely independent of the representation size of the bidders' actual valuations.

[1] This issue has also been studied in the case of combinatorial auctions; see [21, 9, 25].

Truthfulness.

For dealing with the bidders' possible misrepresentation of their true valuations, our mechanisms make use of payments. Formally, let $\mathcal{M}$ be an allocation mechanism that on every vector $v = \{v_1, \ldots, v_n\}$ of valuations produces an allocation giving each bidder i a bundle $I_i(v)$ and charges her a price $P_i(v)$. We say that the mechanism $\mathcal{M}$ is *truthful* if for each bidder i, every vector v of valuations, and every alternative valuation function $v'_i(\cdot)$ of i, it holds that

$$v_i(I_i(v)) - P_i(v) \geq v_i(I_i(v')) - P_i(v')$$

where v' is the vector of (reported) valuations in which v_i is replaced with v'_i. In other words, we require that truthtelling is a *dominant strategy* for each bidder.[2]

VCG Payments and MIR Mechanisms.

A general way for achieving truthfulness in allocations is to use the Vickrey–Clarke–Groves (VCG) scheme (see, e.g. [25] Chapter 9). Fix some vector of valuations $v = (v_1, v_2, \ldots, v_n)$, and denote by A^* the optimal allocation for these valuations, giving each bidder i a plot $A^*(i)$. In addition, for every i denote by A^*_{-i} the optimal allocation that can be achieved when bidder i gets nothing. A VCG mechanism returns the optimal allocation A^*, and charges each bidder i the price

$$P_i = \sum_{k\neq i} v_k(A^*_{-i}(k)) - \sum_{k\neq i} v_k(A^*(k)) \ .$$

Intuitively, bidder i is charged for the loss of value its presence causes to the other bidders; thus, the bidders' individual incentives become "aligned" with the global goal of welfare maximization. We note the besides being truthful, VCG mechanisms (with the above payments) are also *individually rational* in that the utility of any bidder is never negative (and so a bidder can only gain from participating in the auction).

A drawback of the VCG scheme is that its truthfulness critically hinges on that the returned allocation, as well as the allocations used to compute the payments, are all *optimal*. In our setting (and in many other cases), computing optimal allocations is infeasible. However, it is sometimes possible to get good results using the *Maximal-in-Range (MIR)* approach (see, e.g. [24, 10]), based on the following idea. Let $\mathcal{A}$ be a sub-range of all possible allocations such that searching for the optimal allocation in $\mathcal{A}$ is computationally tractable; then a VCG mechanism restricted to the sub-range $\mathcal{A}$ is truthful, individually rational, and computationally efficient. The price of restricting to the range $\mathcal{A}$ is, of course, the degradation in welfare resulting from the fact that the optimal allocation in $\mathcal{A}$ is typically inferior to the "globally-optimal" allocation. However, if this degradation can be bounded (as is the case in our setting), we obtain a truthful mechanism that approximates the optimal welfare. However, in our case, we cannot simply apply this approach, due to the geometric restrictions of the problem; namely, in the continuous setting the problem is not only how to allocate a given set of slices, but also where to slice. Thus, some more work is required before using the MIR technique.

[2] Here we also make the standard assumption that the bidders' utilities are quasi-linear; i.e., that being given an interval I_i and charged a price P_i, bidder i's utility is $v_i(I_i) - P_i$.

3. ALLOCATIONS WITH EQUALLY-SIZED PLOTS

We begin with a relatively restrictive case. Namely, we assume that each agent must be given a plot of size (exactly) ℓ, where ℓ is some given parameter. The mechanism we provide for this case gives a 2-approximation. Before going into the details, however, we state the following theorem, showing that even for such restricted allocations and valuations, finding the optimal allocation is computationally intractable.

THEOREM 1 (FOLLOWING FROM [2]). *Finding an allocation that maximizes the utilitarian welfare is NP-hard, even when given a fixed size ℓ such that the all the pieces given by the allocation must be of size exactly ℓ. This holds in the discrete setting as well (in which each piece is compsed of the same number of items).*

The mechanism we provide here finds an allocation that is guaranteed to have at least $1/2$ of the optimal welfare.

Both this mechanism and the one for the more general case (in which intervals of different sizes may be allocated) use a matching technique. We thus begin with a few definitions that will be useful for both algorithms.

Partitions.

Let $\ell < 1$, we call the *ℓ-partition* of $[0,1]$ the set of intervals

$$\mathcal{P}_\ell^0 = \Big\{ \big(0,\ell\big)\,,\ \big(\ell,2\ell\big)\,,\ \ldots\,,\ \big((\lfloor\frac{1}{\ell}\rfloor-1)\cdot\ell\,,\ \lfloor\frac{1}{\ell}\rfloor\cdot\ell\big) \Big\}\ ;$$

in other words, it is a sequence of ℓ-length consecutive subintervals of $[0,1]$ starting at 0 (hence the superscript). We can similarly define a sequence of such consecutive subintervals beginning at some point $0 < \delta < \ell$; we accordingly call such a sequence a *δ-shifted ℓ-partition* and denote it by

$$\mathcal{P}_\ell^\delta = \Big\{ \big(\delta,\ell+\delta\big)\,,\ \big(\ell+\delta,2\ell+\delta\big)\,,\ \ldots \Big\}$$

where the number of intervals in $\mathcal{P}_\ell^\delta$ may be either $\lfloor\frac{1}{\ell}\rfloor$ or $\lfloor\frac{1}{\ell}\rfloor - 1$, depending on the relation between the sizes δ and $1-\lfloor\frac{1}{\ell}\rfloor\cdot\ell$.

Partition Graphs.

For a partition $\mathcal{P}$ we define a complete and weighted bipartite graph $G_\mathcal{P}$ as follows. The left set of vertices corresponds to the set $[n]$ of bidders, and the right set of vertices corresponds to the set of intervals in $\mathcal{P}$. For any $i \in [n]$ and $I_j \in \mathcal{P}$, we create an edge e_{ij} with weight $w(e_{ij}) = b_i(I_j)$, i.e. the bid of bidder i for interval I_j. Clearly, any matching in $G_\mathcal{P}$ induces an allocation in which each bidder i is allocated the interval to which it is matched (if one exists). We therefore refer to allocations and matchings interchangeably.

The Mechanism.

With these notions in hand, we now describe our mechanism for auctioning uniform-size plots. Given the plot size ℓ, the mechanism, to which we refer as Mechanism 1, operates as follows:

1. *Bidding:*
   - Set $\Delta = 1-\lfloor\frac{1}{\ell}\rfloor\cdot\ell$.
   - Create the partitions $\mathcal{P}_\ell^0$ and $\mathcal{P}_\ell^\Delta$.
   - For every bidder i, get a bid $b_i(I)$ for each interval $I \in \mathcal{P}_\ell^0 \cup \mathcal{P}_\ell^\Delta$ of the two partitions.
2. *Computing the Allocation:*
   - Create the partition graphs $G_{\mathcal{P}_\ell^0}$ and $G_{\mathcal{P}_\ell^\Delta}$.
   - Compute maximum weight matchings M_0^* and M_Δ^* in $G_{\mathcal{P}_\ell^0}$ and $G_{\mathcal{P}_\ell^\Delta}$ (respectively).
   - Return the heavier among M_0^* and M_Δ^*, denoted M^*.
3. *Computing the Payments:*
   - For each bidder i, compute the maximum weight matching in $G_{\mathcal{P}_\ell^0}$ with i removed, and in $G_{\mathcal{P}_\ell^\Delta}$ with i removed; denote the heavier of these matchings by M_{-i}^*.
   - Charge each bidder i a payment
   $$P_i = w(M_{-i}^*) - w(M^*) + b_i(M^*(i))\ .$$

LEMMA 1. *Mechanism 1 requires only $2/\ell$ bids from each bidder, runs in time polynomial in $n+1/\ell$, and is truthful.*

PROOF. The number of bids is apparent from the description of the mechanism. For the running time, it is easy to see that the main computational task of Mechanism 1 is the computing a maximum-weight matching in the bipartite graphs $G_{\mathcal{P}_\ell^0}$ and $G_{\mathcal{P}_\ell^\Delta}$, each having $n+\lfloor 1/\ell\rfloor$ vertices. Computing such matchings can be done in time $O((n+1/\ell)^3)$ [11], and so the total running time of the mechanism is bounded by $O(n*(n+1/\ell)^3)$.

For the truthfulness, note that Mechanism 1 in fact uses VCG payments. In addition, the allocation part of the mechanism computes the maximum among all allocations in which the intervals are either all from $\mathcal{P}_\ell^0$ or all from $\mathcal{P}_\ell^\Delta$. Therefore, Mechanism 1 is a MIR mechanism with VCG payments, and is thus truthful. □

We next show that Mechanism 1 approximates the optimal utilitarian welfare well.

LEMMA 2. *The allocation returned by Mechanism 1 approximates the optimal welfare by a factor of 2.*

PROOF. Consider the optimal allocation A^*, giving each bidder i an interval $A^*(i)$. We create two matchings—M_0 in $G_{\mathcal{P}_\ell^0}$ and M_Δ in $G_{\mathcal{P}_\ell^\Delta}$—and show that one of them must have weight that is at least half the total value of A^*. Thus, the allocation M^* (which the heaviest among all the matchings in $G_{\mathcal{P}_\ell^0}$ and $G_{\mathcal{P}_\ell^\Delta}$) clearly has at least this value.

Recall that we denote by $\Delta = 1-\lfloor\frac{1}{\ell}\rfloor\cdot\ell$ the amount of the resource that is left uncovered in any partition to intervals of size ℓ; clearly $\Delta < \ell$. Let i be a bidder that receives an interval $A^*(i)$ in the optimal allocation; we can write

$$A^*(i) = \big(j_i\cdot\ell+\delta_i\ ,\ (j_i+1)\cdot\ell+\delta_i\big)$$

with $\delta_i < \ell$. The matching M_0 will match i to $I_{j_i}^0$, i.e. to the j_i-th interval (from the left) in $\mathcal{P}_\ell^0$, if $\delta_i \le \Delta$. If $\delta_i > \Delta$, M_0 will match i to $I_{j_i+1}^0$ (i.e. to the (j_i+1)-th interval from the left in $\mathcal{P}_\ell^0$). The matching M_Δ will always match i to $I_{j_i}^\Delta$, i.e. the j_i-th interval from the left in $\mathcal{P}_\ell^\Delta$. Bidders that receive nothing in A^* will not be matched by either M_0 nor M_Δ. We illustrate this with an example in Figure 1.

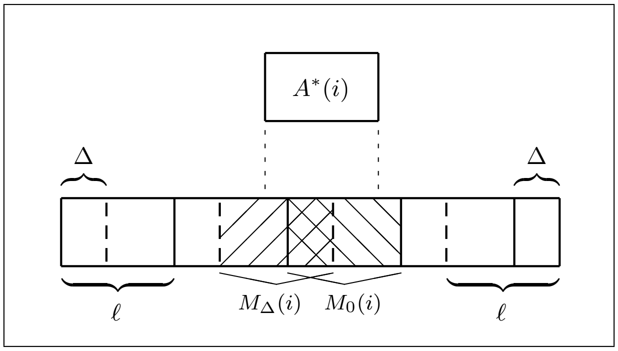

Figure 1: An example of the intervals matched to some i in M_0 and M_Δ. The full lines separate the intervals of $\mathcal{P}_\ell^0$ and the broken ones separate the intervals of $\mathcal{P}_\ell^\Delta$. In this example, bidder i receives in A^* an interval of the form $(j\cdot\ell+\delta_i\,,\,(j+1)\cdot\ell+\delta_i)$ with $\delta_i > \Delta$.

It is easy to see that M_Δ is indeed a valid matching: in order for two bidders i' and i'' to be matched to the same interval I_j^L, it must be that both $A^*(i')$ and $A^*(i'')$ begin somewhere in the interval $[j\cdot\ell\,,\,(j+1)\cdot\ell)$. However, this is clearly impossible, as A^* must allocate them both disjoint intervals of length ℓ.

To see that M_0 too is a valid matching, assume that i' and i'' are both matched in M_0 to the same interval I_j^0. Again, it is impossible that $A^*(i')$ and $A^*(i'')$ begin both within $[j\cdot\ell\,,\,(j+1)\cdot\ell)$ or both within $[(j-1)\cdot\ell\,,\,j\cdot\ell)$. Therefore, w.l.o.g. it is the case that

$$A^*(i') = \big((j-1)\cdot\ell+\delta_{i'}\ ,\ j\cdot\ell+\delta_{i'}\big)$$

and

$$A^*(i'') = \big(j\cdot\ell+\delta_{i''}\ ,\ (j+1)\cdot\ell+\delta_{i''}\big)$$

with $\delta_{i'} > \Delta$ and $\delta_{i''} \le \Delta$. However, this implies that the (non-empty) interval $(j\cdot\ell+\delta_{i''}\,,\,j\cdot\ell+\delta_{i'})$ is common to both $A^*(i')$ and $A^*(i'')$; this is again impossible, and we conclude that M_0 and M_Δ are both valid matchings.

We have thus matched every bidder i getting some interval in A^* to some interval $M_0(i) \in \mathcal{P}_\ell^0$ and some interval $M_\Delta(i) \in \mathcal{P}_\ell^\Delta$. It is also easy to observe that $A^*(i) \subseteq M_0(i) \cup M_\Delta(i)$ for all i; thus, we have that

$$\begin{aligned} v_i(A^*(i)) &\le v_i(M_0(i)) + v_i(M_\Delta(i)) \\ &= b_i(M_0(i)) + b_i(M_\Delta(i)) \end{aligned}$$

where the inequality follows from the monotonicity and sub-additivity assumptions, and the equality from truthfulness. Summing over all bidders, we get that

$$w(M_0) + w(M_\Delta) \ge v(A^*)\ ,$$

which implies that at least in one of the graphs $G_{\mathcal{P}_\ell^0}$ and $G_{\mathcal{P}_\ell^\Delta}$ there exists a matching of weight $\ge v(A^*)/2$. $\square$

Combining Lemma 1 and Lemma 2, we get:

Theorem 2. *Mechanism 1 is an efficient and truthful* 2-*approximation mechanism for the problem of auctioning a continuous resource with fixed plot size.*

Note that the algorithm also works for sub-additive valuation functions, as the proof of Lemma 2 uses only sub-additivity.

4. THE DISCRETE SETTING

We now consider the discrete setting, in which a sequence of m indivisible items $E = (e_1, \ldots, e_m)$ is auctioned, and each bidder can get a at most a single sub-sequence of E.

For this case we show a mechanism that approximates the optimal welfare without the need to assume anything except monotonicity. To make the presentation slightly simpler, assume that the number of items is $m = 2^r$ for some integer r; as we show later, this is without loss of generality.

1. *Bidding:*
   - For every $t = 0, \ldots, r$ create the 2^t-partition (without any shifting), in which each "interval" in the partition is a consecutive sequence of 2^t items. We denote the t-th partition by $\mathcal{P}_t$.
   - For every bidder i, get a bid $b_i(I)$ for each interval $I \in \bigcup_t \mathcal{P}_t$ in these partitions.
2. *Computing the Allocation:*
   - For every $t = 0, \ldots, r$ create the partition graph $G_{\mathcal{P}_t}$.
   - Compute a maximum weight matching M_t^* in the graph $G_{\mathcal{P}_t}$.
   - Among all matchings M_t^*, return the heaviest one, denoted M^*.
3. *Computing the Payments:*
   - For each bidder i consider all the graphs $G_{\mathcal{P}_t}$ with i removed; among all the matchings in these graphs, find the heaviest one, and denote it by M_{-i}^*.
   - Charge each bidder i a payment
   $$P_i = w(M_{-i}^*) - w(M^*) + b_i(M^*(i))\ .$$

Lemma 3. *Mechanism 2 requires at most $2m$ bids from each bidder, runs in time polynomial in $n+r$, and is truthful.*

Proof. The number of bids per bidder is

$$\sum_{t=0}^{r} \frac{m}{2^t} \ =\ m \cdot \sum_{t=0}^{r} \frac{1}{2^t} \ <\ 2m\ .$$

The arguments proving the rest of the lemma are indentical to the ones proving Lemma 1. $\square$

Lemma 4. *The allocation returned by Mechanism 2 approximates the optimal welfare by a factor of $r+1$.*

Proof. Let A^* be an optimal allocation (giving each agent i an interval $A^*(i)$) and let $\mathcal{I} = \bigcup_t \mathcal{P}_t$ be the set of all the intervals in all of the partitions $\mathcal{P}_t$. We show a one-to-one correspondence $f(\cdot)$ between the set of agents receiving non-empty plots in A^* and the intervals of $\mathcal{I}$ such that $v_i(A^*(i)) \le v_i(f(i))$. Since $f(\cdot)$ is one-to-one, restricting it to the intervals of any single $\mathcal{P}_t$ yields a matching in

$G_{\mathcal{P}_t}$, which we denote by M_t^f. The lemma then follows, as

$$w(M^*) \geq \frac{1}{r+1} \cdot \sum_{t=0}^{r} w(M_t^*) \geq$$
$$\geq \frac{1}{r+1} \cdot \sum_{t=0}^{r} w(M_t^f) = \frac{1}{r+1} \cdot \sum_i v_i(f(i)) \geq$$
$$\geq \frac{1}{r+1} \cdot \sum_i v_i(A^*(i)) = \frac{1}{r+1} \cdot v(A^*).$$

It thus remains to show a correspondence $f(\cdot)$ with the above properties. We construct f to map each agent i (getting a non-empty plot in A^*) to the smallest interval in $\mathcal{I}$ that contains $A^*(i)$. It is easy to observe that for each non-empty $A^*(i)$, the choice of $f(i)$ is unique; furthermore, since $f(i) \supseteq A^*(i)$, our monotonicity assumption guarantees that $v_i(f(i)) \geq v_i(A^*(i))$.

To see that $f(\cdot)$ is one-to-one, assume that there exists some interval $I \in \mathcal{I}$ and two agents $k \neq k'$ such that $f(k) = f(k') = I$. This interval clearly cannot belong to $\mathcal{P}_0$, in which each interval contains exactly one item. Thus, $I \in \mathcal{P}_{t+1}$ for some $t \geq 0$; let us explicitly refer to the "endpoints" (i.e. indices of the first and last item) of I, and write $I = [a, c]$. We know that $\mathcal{I}$ also contains the two intervals $[a, b]$ and $[b+1, c]$ (where the length of $[a, b]$ and $[b+1, c]$ is half of the length of $[a, c]$), which are both in $\mathcal{P}_t$. By the definition of f, neither of $A^*(k)$ or $A^*(k')$ is contained in $[a, b]$ or in $[b+1, c]$, but both are contained in $[a, c]$. However, this implies that both $A^*(k)$ and $A^*(k')$ must contain an item from $[a, b]$ as well as an item from $[b+1, c]$, and thus both must contain the items b and $b+1$, which is impossible as they cannot intersect. $\square$

Consider now the case in which m is not an integer power of two. In this case, we set $r = \lceil \log m \rceil$, and set the last partition $\mathcal{P}_r$ to simply contain a single interval of all the items $1, \ldots, m$. To make sure we cover all the items in the other partitions as well, we add to each partition $\mathcal{P}_t$ that does not cover all the items the last interval of the partition P_{t-1}.[3] It is easy to see that this maintains that f is one-to-one (as any interval in $\mathcal{P}_{t+1}$ is the union of at most two intervals in $\mathcal{P}_t$), and thus Lemma 4 holds for this case as well.

Combining Lemma 3 and Lemma 4, we obtain:

THEOREM 3. *Mechanism 2 is an efficient and truthful* ($\lceil \log m \rceil + 1$)*-approximation mechanism for the problem of auctioning a discrete resource.*

5. ALLOCATING A CONTINUOUS CAKE

In this section we treat the classic case, where bidders' valuation functions are non-atomic measures on the interval $[0, 1]$, and there is no limitation on the size of each subinterval a bidder can get. The mechanism works in two stages. In the first stage it chooses $n/2$ bidders at random, and uses them to split the interval $[0, 1]$ to at most $O(n^2)$ segments. In the second stage, it invokes Mechanism 2 to divide the cake between the bidders who did not participate in the first stage:

Mechanism 3

1. *Creating sub intervals:*
   - Choose $n/2$ bidders at random. Denote this set by S.
   - For every bidder $i \in S$, ask i to divide $[0, 1]$ to $2n$ intervals of equal worth to that bidder.
   - Generate a partition J by taking the union of all boundary points reported by the bidders of S.
2. *Computing the Allocation and payments:*
   - Treat every interval in J as a single indivisible item, and invoke Mechanism 2 on the bidders in $[n] \setminus S$ and on the items in J.

Note that bidders in S only help define a partition, but do not receive any cake (nor pay any money).

LEMMA 5. *Mechanism 3 requires at most $2n^2$ bids from each bidder, runs in polynomial time, and is universally truthful.*

PROOF. The bidders in $[n] \setminus S$ see a mechanism having at most n^2 fixed intervals (or discrete items), and they get sequences of them. Hence, for these bidders, the lemma follows directly from Lemma 3. For the bidders in S, the mechanism is truthful since they do not get any utility no matter what they do, and thus they have no incentive to misreport.[4] The bounds on the communication and computational complexity follow straightforwardly from Lemma 3. $\square$

Note that bidders in S are asked questions of the form "given a point a, what is the smallest point b such that $v_i(a, b) = \frac{v_i([0,1])}{2n}$?"; such queries are known as "cut queries" in the cake-cutting literature. If we restrict ourselves to mechanisms with with evaluation queries only (as in Mechanisms 1 and 2), it is easy to show that the best possible approximation ratio possible for this setting is n, which can be obtained by giving the entire cake to a single bidder.

LEMMA 6. *The allocation returned by Mechanism 3 approximates the optimal welfare by a factor of $O(\log n)$ in expectation.*

PROOF. Consider the (random) discrete instance created by the mechanism, in which the items of J are allocated to bidders from $[n] \setminus S$. Denote by A_J^* the optimal allocation for this (discrete) instance, and by $v(A_J^*)$ its welfare. We show that $v(A_J^*)$ is expected to be a constant fraction of $v(A^*)$, i.e. that

$$E\big[v(A_J^*)\big] \geq v(A^*)/\alpha$$

for some constant α. Plugging this into Lemma 4 yields the desired result. The remainder of the proof is therefore dedicated to establishing the above inequality.

[3] To eliminate ambiguity in the definition of f, we simply define it to map each agent i to the smallest interval *in the partition with the minimal t* containing $A^*(i)$.

[4] The mechanism as described provides only that truthtelling is a weakly dominant strategy for bidders in S. This is standard in auction theory (see e.g. [1, 5] for a discussion on strict and weak dominance). By slightly modifying the mechanism we can get, in addition, that for any valuation function of any single bidder, there exist valuations for the other bidders for which truthtelling is strictly dominant. This can be achieved by initially not telling the bidders if they are in S or not, and requiring all bidders to provide their divisions in stage (1) of the mechanism. Then, in stage (2), the valuations provided by the bidders not in S must agree with their previously reported valuations.

Let us call a bidder i *happy* if $v_i(A^*(i)) \geq \frac{v_i([0,1])}{2n}$, i.e. if in the optimal allocation she gets at least $\frac{1}{2n}$ from her value for the entire cake. Denote by

$$H = \left\{ i \ : \ v_i(A^*(u)) \geq \frac{v_i([0,1])}{2n} \right\} .$$

the set of happy bidders.

An easy fact it that the happy bidders must contribute at least half of the welfare in the optimal allocation, i.e.

$$\sum_{i \in H} v_i(A^*(i)) \geq \frac{v(A^*)}{2} . \tag{1}$$

To see this, assume that this is not the case; then

$$\begin{aligned} v(A^*) &= \sum_{i \in H} v_i(A^*(i)) + \sum_{i \in [n] \setminus H} v_i(A^*(i)) \\ &< \frac{v(A^*)}{2} + \sum_{i \in [n] \setminus H} v_i(A^*(i)) \\ &\leq \frac{v(A^*)}{2} + \sum_{i \in [n] \setminus H} \frac{v_i([0,1])}{2n} \end{aligned}$$

which is equivalent to $\frac{1}{n} \sum_{i \in [n] \setminus H} v_i([0,1]) > v(A^*)$. Since $[n] \setminus H$ has at most n bidders, at least one of them must have $v_i([0,1]) > v(A^*)$, and thus giving her the entire cake yields welfare strictly greater than the optimal solution, which is impossible.

Let us now order the happy bidders by the order of their plots in A^* (from left to right), and write $H = \{i_1, \ldots, i_{|H|}\}$; i.e. i_1 is the happy bidder that gets the leftmost piece, i_2 is the happy bidder getting the second-leftmost piece, etc. (We ignore all non-happy bidders that may get a piece placed between the pieces of happy bidders.)

We further call a happy bidder i_k *good* if $i_k \in [n] \setminus S$ and $i_{k-1}, i_{k+1} \in S$, i.e. if she is not in the set S but both her "neighbors" are. (If $k = 1$ or $k = |H|$ we treat it as if the "missing' neighbor is indeed in S.) We denote the set of good bidders by G.

We can now show that in the discretized instance created by the mechanism, there is an allocation A'_J giving each good bidder $i_k \in G$ a piece worth at least the value of her piece in the optimal division $v_{i_k}(A^*(i_k))$. To see that, consider the piece $A^*(i_{k-1})$ of her left happy neighbor. Since i_k is good, it must be that $i_{k-1} \in S$. Since i_{k-1} is a happy bidder, by definition we have that

$$v_{i_{k-1}}(A^*(i_{k-1})) \geq \frac{v_{i_{k-1}}([0,1])}{2n} .$$

Thus, when i_{k-1} is asked to divide the cake into pieces of value $\frac{v_{i_{k-1}}([0,1])}{2n}$, at least one of the boundaries must fall within $A^*(i_{k-1})$; this boundary can be the leftmost boundary of $A'_J(i_k)$. We can symmetrically show how to set the rightmost border of $A'_J(i_k)$ to somewhere within $A^*(i_{k+1})$. Figure 2 illustrates this with an example. Thus, we have obtained that $v_{i_k}(A'_J(i_k)) \geq v_{i_k}(A^*(i_k))$.

Combining the above, the welfare of the optimal division in the discretized instance can be lower-bounded by

$$v(A^*_J) \geq v(A'_J) \geq \sum_{i_k \in G} v_{i_k}(A^*(i_k)) \tag{2}$$

Note, however, that this value depends on the set G which is generated randomly. Clearly, for any single $i_k \in H$, we have

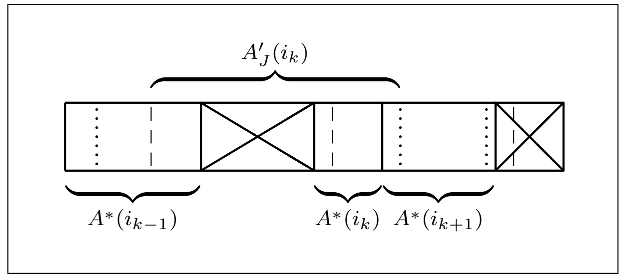

Figure 2: An example of a piece $A'_J(i_k)$. The crossed-out portions are pieces which A^* gives to non-happy bidders, which we ignore. The dotted lines mark the division of the cake into pieces of value $\frac{1}{2n}$ for bidder i_{k-1}, and the dotted lines mark the same division for bidder i_{k+1}. If i_k is "good", i.e. if $i_k \in [n] \setminus S$ and $i_{k-1}, i_{k+1} \in S$, then in the discrete instance it is possible to give i_k the interval $A'_J(i_k)$, which contains $A^*(i_k)$.

that $\Pr\left[i_k \in G\right] \geq \frac{1}{8}$. Thus, by the definition of expectation we can conclude that

$$\begin{aligned} E\left[v(A^*_J)\right] &\geq E\left[\sum_{i_k \in G} v_{i_k}(A^*(i_k))\right] \\ &= \sum_{i_k \in H} v_{i_k}(A^*(i_k)) \cdot \Pr\left[i_k \in G\right] \\ &\geq \frac{1}{8} \cdot \sum_{i_k \in H} v_{i_k}(A^*(i_k)) \\ &\geq \frac{1}{8} \cdot \frac{1}{2} \cdot v(A^*) \end{aligned}$$

where the first inequality follows from (2), and the last one from (1). □

Combining Lemma 5 and Lemma 6, we obtain:

THEOREM 4. *Mechanism 3 is an efficient and universally truthful $O(\log n)$-approximation mechanism for the problem of auctioning a continuous resource.*

6. CONCLUSION AND OPEN PROBLEMS

We studied the problem of auctioning contiguous sub-intervals of a one-dimensional good. The challenges in this setting are two fold: to determine where to cut the cake, and which agent gets what piece. The setting is characterized by an inherent geometric structure, absent in combinatorial auctions. As the problem of finding an optimal allocation of such goods is computationally hard (even without incentives), we provide approximation mechanisms.

There are many natural extensions to this work, most of which are open problem for future research, such as:

- **Variations on the Setting.** In the continuous case we assume that the bidders' valuations are additive. This is used in the first stage of the mechanism—which determines the division into sub-intervals—but it is not necessary for the second. Can this assumption be weakened?

- **Lower Bounds.** The only negative result we are aware of is that no FPTAS exists for the problem unless P=NP [2], and this result holds even without requiring truthfulness. Can a stronger lower bounds on the approximability of welfare by *truthful* mechanisms be proven?
- **Multiple Pieces per Bidder.** We have considered only allocations in which each bidder can receive at most one contiguous piece of the resource. A natural relaxation of this requirement is allowing bidders to get multiple such pieces. If the number of pieces per bidder is a constant, or given as a parameter, it is not hard to see that the allocation problem remains NP-complete. How well can truthful mechanisms approximate the welfare in that case? What can be achieved when dropping the contiguity requirement altogether (while still having a minimum allowed piece size)?
- **2-Dimensional Resources.** While in many cases it is reasonable to model the resource as having only one-dimension, other resources are inherently multi-dimensional, e.g. land. Dealing with such resources seems to require different tools and techniques, and poses interesting challenges for a mechanism designer.

7. REFERENCES

[1] D. Abreu and H. Matsushima. Virtual implementation in iteratively undominated strategies: complete information. *Econometrica: Journal of the Econometric Society*, pages 993–1008, 1992.

[2] Y. Aumann, Y. Dombb, and A. Hassidim. Computing socially-efficient cake divisions. In *AAMAS*, pages 343–350, 2013.

[3] X. Bei, N. Chen, X. Hua, B. Tao, and E. Yang. Optimal proportional cake cutting with connected pieces. In *AAAI*, 2012.

[4] S. J. Brams and A. D. Taylor. *Fair Division: From cake cutting to dispute resolution.* Cambridge University Press, New York, NY, USA, 1996.

[5] J. Chen, A. Hassidim, and S. Micali. Robust perfect revenue from perfectly informed players. In *ICS*, pages 94–105, 2010.

[6] Y. Chen, J. Lai, D. C. Parkes, and A. D. Procaccia. Truth, justice, and cake cutting. In *AAAI*, 2010.

[7] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J. Padget, S. Phelps, J. A. Rodriguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica (Slovenia)*, 30(1):3–31, 2006.

[8] P. Cramton. The fcc spectrum auctions: An early assessment. *Journal of Economics & Management Strategy*, 6(3):431–495, 1997.

[9] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial auctions.* MIT Press, 2006.

[10] S. Dobzinski and N. Nisan. Limitations of vcg-based mechanisms. In *STOC*, pages 338–344, 2007.

[11] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19(2):248–264, Apr. 1972.

[12] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35(12):1969–1985, 1999.

[13] M. Guo and V. Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *AAMAS*, pages 881–888, 2010.

[14] L. Han, C. Su, L. Tang, and H. Zhang. On strategy-proof allocation without payments or priors. In *WINE*, volume 7090 of *Lecture Notes in Computer Science*, pages 182–193, 2011.

[15] R. Johari and J. N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. *Operations Research*, 57:823–839, July-August 2009.

[16] I. Kremer and K. G. Nyborg. Divisible-good auctions: The role of allocation rules. *RAND Journal of Economics*, pages 147–159, 2004.

[17] B. Lehmann, D. J. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. In *ACM Conference on Electronic Commerce*, pages 18–28. ACM, 2001.

[18] R. T. Maheswaran and T. Basar. Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decision and Negotiation*, 13:2003, 2003.

[19] A. Maya and N. Nisan. Incentive compatible two player cake cutting. In *WINE*, volume 7695 of *Lecture Notes in Computer Science*, pages 170–183. Springer, 2012.

[20] E. Mossel and O. Tamuz. Truthful fair division. In *Algorithmic Game Theory*, pages 288–299. Springer, 2010.

[21] N. Nisan. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 1–12, 2000.

[22] N. Nisan, J. Bayer, D. Chandra, T. Franji, R. Gardner, Y. Matias, N. Rhodes, M. Seltzer, D. Tom, H. Varian, and D. Zigmond. Google's auction for tv ads. In *Automata, Languages and Programming*, volume 5556 of *Lecture Notes in Computer Science*, pages 309–327. Springer Berlin Heidelberg, 2009.

[23] N. Nisan and A. Ronen. Algorithmic mechanism design (extended abstract). In *STOC*, pages 129–140, 1999.

[24] N. Nisan and A. Ronen. Computationally feasible vcg mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.

[25] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory.* Cambridge University Press, New York, NY, USA, 2007.

[26] A. D. Procaccia. Cake cutting: not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.

[27] J. Robertson and W. Webb. *Cake-cutting algorithms: Be fair if you can.* A K Peters, Ltd., Natick, MA, USA, 1998.

[28] M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinational auctions. *Management Science*, 44(8):1131–1147, 1998.

[29] M. Tennenholtz. Some tractable combinatorial auctions. In *AAAI/IAAI*, pages 98–103, 2000.

[30] S. Yang and B. Hajek. Vcg-kelly mechanisms for allocation of divisible goods: Adapting vcg mechanisms to one-dimensional signals. *Selected Areas in Communications, IEEE Journal on*, 25(6):1237–1243, 2007.