8803 Connections between Learning, Game Theory, and Optimization

Maria-Florina Balcan Lecture 18: October 26, 2010

Reading: Algorithmic Game Theory book, Chapter 11, in particular, Sections 11.1, 11.2.

In this lecture, we study combinatorial auctions, where a large number of items are auctioned concurrently and bidders are allowed to express preferences on bundles of items. We focus on combinatorial auctions with single minded bidders and provide a mechanism that is computationally efficient, incentive compatible, and provides a $\sqrt{m}$ approximation to the total social welfare, as good as theoretically possible in polynomial time.

Combinatorial Auctions

The combinatorial auction setting is formalized as follows. There is a set X of m indivisible goods that are concurrently auctioned among n bidders/players. Each player $i \in \{1, \ldots, n\}$ has a valuation function mapping a subset of products to their value to the player, so $v_i : 2^X \to \mathbb{R}$.

Typical assumptions, that we will make today, are monotonicity (also known as "free disposal"), i.e. $v_i(S) \leq v_i(T)$, $\forall\, S \subseteq T$, and $v(\emptyset) = 0$.

The whole point of defining a player valuation is that the value of a bundle of items need not be equal to the sum of his values for the items in the bundle. For disjoint sets $S \cap T = \emptyset$, we say that S, T are complements if $v(S \cup T) \geq v(S) + v(T)$ (for example S, T are a right shoe and a left shoe) and we say that S, T are substitutes if $v(S \cup T) \leq v(S) + v(T)$ (for example S, T are margarine and butter).

We will assume *quasilinear* utilities, i.e. if the bidder wins a bundle S and pays p then its utility is $v_i(S) - p$. We also assume *no externalities*, i.e. that a bidder only cares about the items he receives and not about how the other items are allocated among the other bidders.

In general we have representation and communication concerns: the valuation functions are exponential sized objects since they specify a value for *each* bundle. So we in general we have to deal with the question of how can enough information be transferred to the seller so that a reasonably good decision can be made? For the rest of the lecture today we will consider a class of valuations for which representation is easy, exact optimization is challenging, but the best possible approximation is achieved by a truthful mechanism.

The Single-Minded Case

Single minded bidders are only interested in a single specified bundle of items, and get a specified scalar value if they get this whole bundle (or any superset) and get zero value for any other bundle. Formally:

Definition 1. *A valuation v is called* single-minded *if there exists a bundle of items S^* and a value $v^* \in \mathbb{R}^+$ such that $v(S) = v^*$ for all $S \supseteq S^*$ and $v(S) = 0$ for all other S. A single-minded bid is a pair (S, v).*

Welfare Maximization: The Algorithmic Problem

Let us first consider the algorithmic allocation problem (ignoring incentives) for single-minded bidders.

Input (S_i^*, v_i^*) for $i = 1, \ldots, n$.

Output A subset of winning bids $W \subseteq \{1, \ldots, n\}$ with maximum social welfare $\sum_{i \in W} v_i^*$ such that every two winners are compatible, i.e. $S_i^* \cap S_j^* = \emptyset$ for $i \neq j \in W$.

The problem of maximizing social welfare is the weighted packing problem in disguise and thus it is NP-hard, as we show now via a simple reduction from Independent Set.

Theorem 1. *The allocation problem among single-minded bidders is NP-hard. More precisely, its decision version, that of deciding whether the optimal allocation has social welfare of at least k (where k is an additional part of the input) is NP-complete.*

Proof: We describe a reduction form the decision version of the NP-complete problem Independent Set: given an undirected graph $G = (V, E)$ and $k \in \mathbb{Z}^+$, does G have an independent set of size k? An independent set is a subset of vertices that have no edge between any two of them.

Given an Independent Set instance, we construct an allocation problem instance as follows. The set of items will be E, the set of edges in G. We will have a player for each vertex in the graph. For vertex $i \in V$, we will have the desired bundle of i be the set of adjacent vertices $S_i^* = \{e \in E : i \in e\}$ and the value $v_i^* = 1$.

Notice that a set of winners W in the allocation problem satisfies $S_i^* \cap S_j^* = \emptyset, \forall i, j \in W$ if and only if the set of vertices corresponding to W is an independent set in the original graph. The social welfare obtained by W is exactly the size of this independent set. Thus an independent set of size at least k exists if and only if the social welfare of the optimal allocation is at least k. This implies the NP-hardness of the decision version of maximizing social welfare for single-minded bidders. As the problem is clearly in NP (one can trivially check that a set of winners has social welfare at least k and that the winners are compatible), it is NP-complete. ■

Not only finding the maximum independent set size is NP-complete, but it is known that *approximating* it to within a factor of $n^{1-\varepsilon}$ (for any fixed ε) is NP-complete. Since in our reduction, the social welfare was equal to the size of the independent set, we get the same hardness result for maximizing social welfare among single-minded bidders.

Since the number of edges is at most the number of players squared, i.e. $m \leq n^2$, we get:

Theorem 2. *Approximating the optimal allocation among single-minded bidders to within a factor better than $m^{1/2-\varepsilon}$ is NP-hard.*

If we want to obtain a reasonable (say, $\sqrt{m}$) approximation, this ruins our hopes of using the Vickrey-Clarke-Groves (VCG) mechanism since VCG can only be used when we can exactly compute the optimal allocation.

We will now see a clever truthful mechanism for single-minded bidders that achieves the best approximation factor, $\sqrt{m}$.

Welfare Maximization: An Incentive Compatible Mechanism

We will now bring incentives back in to the picture and we will consider bids (S_i, v_i) for $i = 1, \ldots, n$ from the players, that may not coincide with their true (S_i^*, v_i^*). In the mechanism below, we will show that it is a dominant strategy for each bidder to bid truthfully, i.e. we will have $(S_i, v_i) = (S_i^*, v_i^*)$.

Initialization: let $W \leftarrow \emptyset$ and reorder the bids such that $\frac{v_1}{\sqrt{|S_1|}} \geq \cdots \geq \frac{v_n}{\sqrt{|S_n|}}$

for $i = 1..n$ **do**

  if i is compatible with the winners so far, i.e. $S_i \cap \left(\bigcup_{j \in W} S_j\right) = \emptyset$ **then**

    add i to the set of winners: $W \leftarrow W \cup \{i\}$

The algorithm outputs the set W of winners plus payments for each player defined as follows.

Payments. Charge each player $i \in W$

$$p_i = \frac{v_j}{\sqrt{\frac{|S_j|}{|S_i|}}}$$

where j is the smallest index such that $S_i \cap S_j \neq \emptyset$ and $S_k \cap S_j = \emptyset$ for all $k < j, k \neq i$. Bidder i's payment p_i is thus the minimum bid he could have made (given its set S_i) for which he still wins. If no such j exists for i, then let $p_i = 0$. Bidder j is the one who lost exactly because of i: had bidder i not participated in the auction, j would have won.

We will show that this greedy mechanism is essentially the best we can hope for.

Theorem 3. *The greedy mechanism is efficiently computable, incentive compatible and produces a $\sqrt{m}$ approximation to the social welfare.*

Computational efficiency is obvious. Incentive compatibility follows directly from Lemma 1.

Lemma 1. *A deterministic mechanism for single-minded bidders in which losing bidders pay* 0 *is incentive compatible if and only if it satisfies the following two conditions:*

monotonicity *If i wins for a bid (S_i, v_i) then i still wins for any bid (S_i', v_i') in which he offers more for fewer items, i.e. $v_i' > v_i$ and $S_i' \subsetneq S_i$*

critical payment *i's payment when bidding set S_i is the minimum value v_i^c such that a bid (S_i, v_i^c) still wins. $p_i(S_i, v_i) = \inf\{v_i' : i$ wins for $(S_i, v_i')\}$.*

We first argue that the greedy mechanism satisfies Lemma 1 then prove Lemma 1. The mechanism is monotonic since any bidder i can only advance in the ordering for a higher bid v_i' or for fewer items $|S_i'|$. The critical payment condition is met since i wins as long as he appears in the greedy order before j. The payment computed is exactly the value at which the transition between i being before and after j in the greedy order occurs: indeed $v_i \geq \frac{v_j}{\sqrt{\frac{|S_j|}{|S_i|}}} \iff \frac{v_i}{\sqrt{|S_i|}} \geq \frac{v_j}{\sqrt{|S_j|}}$.

Proof:[of Lemma 1] We will only prove sufficiency, i.e. that monotonicity and critical payment imply incentive compatibility. We will fix the bids of players other than i and drop the index i for notational simplicity. We assume that monotonicity and critical payment hold and aim to show

$$u(S^*, v^*) \geq u(S', v') \,\forall\, (S', v'), \tag{1}$$

where $u(S, v)$ is the utility for bidding (S, v). Let $\tilde{p}(S) = \inf\{v_i' : i \text{ wins for } (v_i', S)\}$.

First observe that under the given conditions a truthful bidder will never receive negative utility, i.e. $u(S^*, v^*) \geq 0$. His utility is 0 for losing (losers pay 0). For winning, his value must be at least the critical value $v \geq \tilde{p}(S^*)$ by definition.

We can assume that (S', v') wins and $S' \supseteq S^*$: otherwise Eq. (1) is obvious.

We first show that $u(S^*, v') \geq u(S', v')$. For this we will prove that $\tilde{p}(S') \geq \tilde{p}(S^*)$. For all $x < \tilde{p}(S^*)$, a bid (S^*, x) loses (by definition of critical value). By monotonicity, we get that (S', x) loses as well. Thus $\tilde{p}(S') \geq x$, $\forall\, x < \tilde{p}(S^*)$, implying $\tilde{p}(S') \geq \tilde{p}(S^*)$. Thus by bidding (S^*, v') instead of (S', v') the bidder still wins and his payment will not increase, implying our claim that $u(S^*, v') \geq u(S', v')$.

It is left to show that bidding (S^*, v^*) is never worse than the winning bid (S^*, v'). Let $\tilde{p} = \tilde{p}(S^*)$.

Suppose first that (S^*, v^*) wins. If $v' \geq \tilde{p}$ then $u(S^*, v') = u(S^*, v') = v^* - \tilde{p}$. If $v' < \tilde{p}$ then (S^*, v') loses by definition of $\tilde{p}$ implying that $u(S^*, v') = 0 \leq u(S^*, v^*)$.

Suppose now (S^*, v^*) loses. Then $v^* \leq \tilde{p}$. If (S^*, v') wins then $u(S^*, v') = v^* - \tilde{p} \leq 0$. ■

Since the greedy mechanism is truthful we can assume that $v_i = v_i^*$ and $S_i = S_i^*$. The only thing left to establish is the approximation factor achieved by the greedy algorithm, which turns out to be the optimal one, $\sqrt{m}$.

Lemma 2. *Let OPT be an allocation (i.e. set of winners) with maximum value of $\sum_{i \in OPT} v_i^*$ and let W be the output of the greedy algorithm. Then*

$$\sum_{i \in OPT} v_i^* \leq \sqrt{m} \sum_{i \in W} v_i^*$$

i.e. the greedy algorithm has an approximation factor of $\sqrt{m}$.

Proof: For each $i \in W$, let $OPT_i = \{j \in OPT : j \geq i, S_j^* \cap S_i^* \neq \emptyset$ be the collection of players j in OPT that did not enter W because of bidder i (in addition to i itself).

Clearly $OPT \subseteq \bigcup_{i \in W} OPT_i$ (any $j \in OPT$ is either in W or is knocked out because of some $i \in W$). This implies that $\sum_{i \in OPT} v_i^* \leq \sum_{i \in W} \sum_{j \in OPT_i} v_j^*$.

Thus it is enough to show that $\forall\, i \in W, \sum_{j \in OPT_i} v_j^* \leq \sqrt{m} v_i^*$.

Note that every $j \in OPT_i$ appeared after i in the greedy ordering and thus $v_j^* \leq \frac{v_i^*}{\sqrt{|S_i^*|}} \sqrt{|S_j^*|}$. We get

$$\begin{aligned}\sum_{j \in OPT_i} v_j^* &\leq \frac{v_i^*}{\sqrt{|S_i^*|}} \sum_{j \in OPT_i} \sqrt{|S_j^*|} \\ &\leq \frac{v_i^*}{\sqrt{|S_i^*|}} \sqrt{OPT_i} \sqrt{\sum_{j \in OPT_i} |S_j^*|}\end{aligned}$$

where the last inequality follows from the Cauchy-Schwarz inequality.

Every S_j^* for $j \in OPT_i$ intersects S_i^*. Since OPT is a valid allocation, the sets S_j^* for $j \in OPT_i$ must be disjoint implying that $|OPT_i| \leq |S_i^*|$.

As OPT is an allocation, $\sum_{j \in OPT_i} |S_j^*| \leq m$ implying $\sum_{j \in OPT_i} v_j^* \leq \sqrt{m} v_i^*$, as desired. ■

Note: Note that ordering bidders based on v_i or $v_i/|S_i|$ would preserve incentive compatibility, but provide a much worse approximation ratio.