COMP/MATH 533: Algorithmic Game Theory September 29, 2014

Lecture 8

Lecturer: Yang Cai *Scribe: Andrew Holliday and Richard Torres*

NOTE: The content of these notes has not been formally reviewed by the lecturer. It is recommended that they are read critically.

1 Introduction

In combinatorial auctions, a large number of items are auctioned concurrently and bidders are allowed to express preferences on bundles of items. This is preferable to selling each item separately when there are dependencies between different items. Some of the direct applications of combinatorial auctions include spectrum auctions and resource allocation. Combinatorial auctions are notoriously hard in both theory and practice. For this reason, the combinatorial auction problem has a role in game theory research similar to the role of the traveling-salesman problem in combinatorial optimization. New ideas about analyzing and designing mechanisms will be tested first on combinatorial auctions to gauge their performance.

2 The set-up

The setting of a combinatorial auction is the following:

There are n bidders (for instance: Bell, Rogers, Telus, etc), and a set M of m non-identical items (for example: a license for broadcasting at a certain frequency in a given region).

An ***outcome*** is a n-dimensional vector $(S_1, S_2, ..., S_n)$ with the S_i's pairwise disjoint, where each S_i denotes the set of items allocated to bidder i (her bundle). Notice that there are $(n+1)^m$ possible outcomes.

In principle, each bidder could value each outcome differently, but here we simplify our framework a bit by assuming that each bidder only has a value for their own allocation. This is, bidder i has a private value $v_i(S)$ for each subset S of M. Each valuation v_i must have "free disposal", i.e., be monotone: for $S \subseteq T \subseteq M$ we have that $v_i(S) \leq v_i(T)$, and it should be "normalized": $v_i(\emptyset) = 0$. We could make other assumptions on the valuation function in order to simplify the auction design problem. We will discuss that later on these notes.

The ***welfare*** of an outcome $(S_1, S_2, ..., S_n)$ is given by $\sum_{i=1}^{n} v_i(S_i)$.

3 Challenges of Combinatorial Auctions

Challenge 1: Preference elicitation can be impractical.

Clearly, we cannot use direct-revelation sealed-bid auction in this setting, since each bidder has 2^m numbers to specify (notice that by just taking $m = 20$ each bidder would have around 1 million numbers to specify!). To solve this issue we can use indirect mechanisms.

Indirect mechanisms are mechanisms in which the bidders reveal partial information to the mechanism. This avoids the problem of requiring the bidders to specify all of their value information at once. Examples of indirect mechanisms are the Ascending English Auction (the one you see in the movies!) or any of its variants. In the Japanese variant (one which is easy to argue about), the auction begins at some opening price, which is publicly displayed and it increases at a steady rate. Each bidder either chooses to stay in or to drop out, and once a bidder drops out she cannot return. The winner is the last bidder to stay in, and the sale price is the price at which the second-to-last bidder dropped out. In this setting, each bidder has a dominant strategy, which consists of staying till the price is higher than her value. Notice that if we apply the revelation principle on this auction we get the Vickrey auction.

Challenge 2: Welfare maximization can be intractable.

Welfare maximization is not always tractable. For instance, consider the “single-minded” case. Here, the bidders are interested only in a single specified bundle of items, having a value for this bundle (or any superset) and value zero for any other bundle. Single-minded valuations are thus very simply represented. But it turns out, that even with this relatively simple valuations it is NP-Hard to maximize welfare for single-minded bidders.

Challenge 3: VCG can have bad revenue and incentive properties (despite being DSIC).

To see this, let's consider the following example: We have 2 bidders and 2 items (item A and item B). Bidder 1 only wants both items, and his valuation is given by $v_1(AB) = 1$ and zero otherwise. On the other hand, bidder 2 is only interested in item A, having valuation $v_2(AB) = v_2(A) = 1$ and zero otherwise. In this setting, VCG gives both items to the first bidder, charging him a price of 1. Now, suppose that there is a third bidder who only wants item B, with valuation $v_3(AB) = v_3(B) = 1$ and zero otherwise. What does VCG do in this case? It gives A to the second bidder, B to the third bidder, and charges them 0!

This example shows that VCG can be vulnerable to collusion and false-name bidding. For instance, it creates an incentive for bidder 2 to create a ”false bidder” 3 - by having a friend enter the auction, for example - to get its desired item for a lower price.

Challenge 4: Indirect mechanisms are usually iterative, which offers new opportunities for strategic behavior.

When discussing Challenge 1, we introduced indirect mechanisms as a possible solution to the problem of preference elicitation. However, indirect mechanisms can also be problematic. Since they are usually iterative, proceeding in steps like the Ascending English auction, they can be manipulated by bidders in various ways. One way is for bidders to exchange information with each other using the values of their bids.

A famous example of this took place in a US spectrum auction [2]. License #378, a spectrum license for Rochester, MN, was part of an auction with US West and Macleod battling for it. After being frustrated by Macleod in this auction, US West bidded on many other licenses on which Macleod had the leading bid, setting all bids to be multiples of 1000 plus 378. The message was clear: US West really wanted license #378, and if Macleod did not stop bidding for that license, US West would try to win many of the other licenses Macleod was battling for.

4 Applications of Combinatorial Auctions: Spectrum Auctions

Spectrum auctions are arguably the killer application of combinatorial auctions. They deserve particular attention because how common they are and because of the large amounts of money involved. Various approaches have been used to attempt to design good spectrum auctions.

One obvious approach is to simply carry out a set of single-item auctions, one for each section of spectrum. This works well if a bidder's value for (say) two items is less or equal than the sum of its values for each individually (e.g., they are at least partially redundant). In this case, we call the items **substitutes**, and this can be represented in terms of the value function as: $v(AB) \leq v(A) + v(B)$ for items A and B. Welfare maximization is computationally tractable when the items are substitutes and true valuations are known. On the other hand, we call two items **complements** if a bidder's value of getting (say) two items is greater than the sum of its values for each individually (e.g., they are partially co-dependent). We can write this last case as: $v(AB) > v(A) + v(B)$. In this setting (complement items), welfare maximization remains intractable.

In real situations, items tend to be a mixture of substitutes and complements. When the problem is "mostly substitutes", selling items separately could have good performance.

How do we run the auction?

The next question is how to separate the auction into single-item auctions. We discuss two possible problems that can arise in this context:

Rookie Mistake 1:

One's first instinct may be to run the auctions in sequence - hold an auction for one item, then for another one, and so on. However, this approach proves to be problematic.

To see why, imagine a set of k identical items to be auctioned, and suppose you decide to run a sequence of Vickrey auctions on these items individually. If all n bidders bid truthfully, in the first auction the item will be sold to the bidder with the highest value, with price being the second-highest value. In the second auction, an identical item will be sold to the bidder with the second-highest value with price of the third-highest value, and so on.

So for each auction after the first, an identical item is sold at a lower price than the previous auction. A clever bidder, realizing this, has an incentive to bid lower and try to win in a later auction - so the auctions are no longer DSIC! In fact, discovering an optimal strategy for this case is nontrivial, which is what we as designers wish to avoid.

This mistake has been made in real spectrum auctions in the past. In March of 2000, Switzerland held a sequence of three auctions for individual spectrum blocks. The first two auctions were for identical blocks of 28 Mhz, while the third was for a block of 56 Mhz (twice the size of the previous two). The first item sold for 121 million Swiss francs, the second sold for 134 million, and the third for 55 million.

Rational agents would value the third block twice as much as the first two, so in an optimal auction or set of auctions it should have sold for twice as much. But in fact it sold for half as much. Clearly, this outcome was far from achieving optimal welfare and revenue.

This shows that the single-item auctions should be held *simultaneously* rather than *sequentially.*

Rookie Mistake 2:

Another mistake one could make would be to use sealed-bid single-item auctions. Again, imagine k identical items, and suppose that each bidder wants only one item. One strategy a bidder might pursue would be to bid its true value in just one of the auctions, hoping that it will win. Another strategy, would be to underbid in multiple auctions, and hope that she wins just one (or at least not too many of them).

With the second strategy, the bidder's odds of winning any one auction are lower, but her odds of winning at least one auction may be comparable to the odds of winning in the earlier strategy. There is also a chance that the bidder will get the item at a lower price than by making just one honest bid; and there is a chance she will have to pay for more than one item, which clearly she doesn't want. It is not obvious what to do in this case - this again leads to poor performance in terms of welfare and revenue.

This mistake has been made in real auctions in the past [4]. In 1990, the New Zealand's government auctioned off television broadcasting licenses using simultaneous sealed-bid Vickrey auctions. Based on the estimated worth of the licenses, the government projected that it would make around $250 million in revenue. Not only did the actual revenue - $36 million - fall far short of this, but the behavior observed in some of the auctions was bizarre. The winning bid for one license was $100,000, but the second-highest bid was just $6. Another license was won by a bidder who bid $7 million while the second-highest bid was $5,000.

Simultaneous Ascending Auctions (SAAs)

The Simultaneous Ascending Auction (SAA) consists of a set of single-item Ascending English Auctions running simultaneously with the same set of bidders. In each round of bidding, each bidder bids on any subset of the items it wishes, subject to an "activity rule" which states that the total number of items a bidder bids on must not increase from one round to the next. In practice, the Simultaneous Ascending Auction (SAA), or some modification of it, is used to carry out most spectrum auctions.

The major advantage of the SAA is that it allows for the interactions of the bidders with the mechanism to naturally "discover" the price of the item, a process called price discovery. The corresponding process for value, value discovery, also takes place, revealing information about the values of the items only as it is needed, rather than all at once. All this allows bidders to bid more wisely, altering their strategies as they auction progresses and they learn more about each other's valuations.

Despite having these advantages mentioned above, the SAAs are not perfect either. As an illustration of this, let's take a look at the following auction:

In 1999 in Germany [3], the government auctioned ten blocks of cell-phone spectrum in ten simultaneous ascending auctions, with the rule that each bid must be at least 10% greater than the previous bid. There were two companies bidding in the auction, T-Mobile and Mannesman. Mannesman's first bid was 20 million Deutsche marks on blocks 1-5 and 18.18 million on blocks 6-10. Notice that 10% more than 18.18 million is 19.99 million, slightly less than 20 million. T-Mobile saw these bids and interpreted them as an offer to split the blocks evenly for 20 million each. T-Mobile accepted the offer by bidding 20 million on blocks 6 to 10, and the auction ended there. This was arguably a close-to-optimal outcome in terms of welfare, but it was far from optimal in revenue - the government made considerably less money from the auction than expected.

References

[1] N. Nisan. T. Roughgarden, E. Tardos, V. Vazirani. *Algorithmic Game Theory,* Chapter 11. Cambridge University Press 2007.

[2] P. Cramton, J. A. Schwartz. *Collusive Bidding in the FCC Spectrum Auctions,* page 4. In *Contributions to Economic Analysis & Policy,* Volume 1, Issue 2, 2002.

[3] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach,* 3rd Edition, Chapter 17. Prentice Hall. Copyright 2010, 2003, 1995.

[4] F. Islam. *Sold to the slyest bidder.* Article, The Guardian, Sunday 23 April 2000.