

CS364A: Algorithmic Game Theory

Lecture #8: Combinatorial and Wireless Spectrum Auctions*

Tim Roughgarden[†]

October 16, 2013

1 Selling Items Separately

Recall that a combinatorial auction has n bidders and m non-identical items, with bidder i having a private valuation $v_i(S)$ for every bundle $S \subseteq M$ of items. Asking each bidder to report 2^m bids is absurd unless m is very small. Thus, for the first time in the course, we have no choice but to design and analyze indirect mechanisms, and especially iterative mechanisms that query bidders for relevant valuation information on a “need-to-know” basis. This entails relaxing both the DSIC guarantee and full welfare maximization — we will miss these properties, but have no alternative.

What other mechanisms can we try? Given that we need to sell multiple items, and don’t want to elicit valuations for every bundle, the simplest mechanisms to try are those that sell the items separately, using some type of single-item auction for each. We could certainly implement such an auction if desired — all we need is one bid per bidder per item, which is arguably the minimum imaginable.

We’ll pin down the precise auction format shortly, but first we should ask a more basic question: could selling items separately conceivably work, even in principle? There is lots of beautiful and clarifying theory on this question, some of which we’ll cover later. For now we summarize the main take-aways from this theory.

There is a fundamental dichotomy between combinatorial auctions in which items are *substitutes*, and those in which items are *complements* — with the former being far easier, in theory and in practice, than the latter. Roughly speaking, items are substitutes if you get diminishing returns from them — having one item only makes others less valuable. For two items A and B , for example, the substitutes condition means that $v(AB) \leq v(A) + v(B)$. In a spectrum auction context, two licenses for the same area with equal-sized frequency

*©2013, Tim Roughgarden.

[†]Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

ranges are usually substitute items. Theory indicates that selling items separately has a chance to work well when items are (mostly) substitutes. For starters, welfare maximization is a computationally tractable problem when items are substitutes and the true valuations are known. In addition, the undesirable properties of the VCG mechanism pointed out last lecture and in the exercises evaporate when items are substitutes, generalizing the single-item case. But even though substitute items are the “easy” case, we’ll see that it’s easy to screw up when trying to sell them separately.

Items are complements if there are synergies between them, so that possessing one makes others more valuable. With two items A and B , this translates to the property $v(AB) > v(A) + v(B)$. Complements arise naturally in wireless spectrum auctions, as some bidders want a collection of licenses that are adjacent, either in their geographic areas or in their frequency ranges. With complements, welfare maximization (without incentive constraints) is already a very difficult problem; see also Problem 5. We cannot expect a simple auction format like separate single-item auctions to perform well in such environments.

The items in spectrum auctions, and most real-world combinatorial auctions, are a mixture of substitutes and complements. If the problem is “mostly substitutes,” then separate single-item auctions might already perform well, if properly implemented. If not, then additional ideas are needed; see Section 3.

2 Simultaneous Ascending Auctions

There are numerous ways to organize separate single-item auction. Next we discuss two of the design decisions that seem to matter a lot in practice; see Cramton [2] and Milgrom [4, Chapter 1] for more details.

Rookie mistake #1: Hold the single-item auctions sequentially, one at a time.

To see why holding auctions sequentially is probably a bad idea, consider the especially easy case of identical items, where each bidder wants at most one. This problem can be solved easily via a single auction that allocates all of the items (e.g., by extending the Vickrey auction to this setting). Suppose instead we hold a sequence of single-item auctions. Concretely, consider two identical items, sold via back-to-back Vickrey auctions, and suppose you are a bidder with a very high valuation — you expect to win any auction that you participate in. What should you do? First, suppose everyone else bids straightforwardly, meaning that, if they haven’t won an item yet, then they participate in the next auction and bid their true valuation. If you participate in the first auction, you would win and pay the second-highest valuation. If you skip it, the bidder with the second-highest valuation would win the first auction and disappear, leaving you to win the second auction at a price equal to the third-highest original valuation. Of course, now that we realize that it is not a dominant strategy for bidders to bid straightforwardly in a sequence of Vickrey auctions, we have to reason about how they might be strategizing. Summarizing, it’s hard to bid intelligently in a sequence of Vickrey auctions because you have to guess the expected selling price of future auctions, and this in turn makes the auctions’ outcomes unpredictable, with the possibility

of low welfare and revenue.

In March 2000, Switzerland auctioned off 3 blocks of spectrum via a sequence of Vickrey auctions. The first two auctions were for identical items, 28 MHz blocks, and sold for 121 million and 134 million Swiss francs, respectively. This is already more price variation than one would like for identical items. But the kicker was that in the third auction, where a larger 56 MHz block was being sold, the selling price was only 55 million! The bids were surely far from equilibrium, and both the welfare and revenue achieved by this auction are suspect.¹

The discussion and history lessons above suggest holding single-item auctions for multiple items *simultaneously*, rather than sequentially. But there is still the question of the auction format for each single-item auction.

Rookie mistake #2: Use sealed-bid single-item auctions.

In 1990, the New Zealand government auctioned off essentially identical licenses for television broadcasting using simultaneous (sealed-bid) Vickrey auctions. It is again difficult for bidders to figure out how to bid in such an auction. Imagine, for example, that there are 10 licenses and you want one of them (but not more). How should you bid? One legitimate strategy is to pick one of the licenses (at random, say) and go for it. Another strategy is to bid less aggressively on multiple licenses, hoping that you get one at a bargain price, and that you don't win too many extra licenses that you don't want. The difficulty is trading off the risk of winning too many licenses with the risk of winning too few.

The difficulty of bidding and coordinating in a simultaneous sealed-bid auction makes the auction format vulnerable to outcomes with low welfare and revenue. For example, suppose there are three bidders and two identical items, and each bidder wants only one. The obvious extension of the Vickrey auction sells the two licenses to the bidders with the highest valuations, each at a price equal to the smallest valuation. In a simultaneous sealed-bid auction, if each bidder targets only one license, then one of the licenses is likely to have only one bidder and will thus be given away for free (or more generally, sold at the reserve price).

The revenue in the 1990 New Zealand auction was only \$36 million, a paltry fraction of the projected \$250 million. In contrast, most spectrum auctions over the past 20 years have met or exceeded projected revenues. On one license, the high bid was \$100,000 while the second-highest bid (and selling price) was \$6! On another, the high bid was \$7 million and the second-highest was \$5,000. To add insult to injury, the high bid were made available to the public, who could then see just how much money was left on the table! A later New Zealand auction kept the simultaneous sealed-bid format but switched to first-price auctions — this switch probably failed to prevent the miscoordination and consequent welfare and revenue losses that plagued the previous auction, but it did make these losses less evident to the public.

Simultaneous *ascending* auctions (SAAs) form the basis of most spectrum auctions run over the last 20 years. We discuss the basic format first, and then some of the bells and

¹In addition to the questionable auction format, it didn't help matters that there were some strategic mergers of potential bidders before the auction, leading to less competition than expected.

whistles that have been added on over the years. Conceptually, SAAs are like a bunch of single-item English auctions being run in parallel in the same room, with one auctioneer per item. More precisely, each round, each bidder can place a new bid on any subset of items that it wants, subject to an *activity rule*. The activity rule forces all bidders to participate in the auction from the beginning and contribute to the price discovery discussed below. The details of an activity rule can be complex, but the gist is to require that the number of items that a bidder bids on only decreases over time as prices rise. Generally, the high bids and bidders are visible to all — even though this can encourage signaling and retaliatory bids (recall USWest vs. McLeod last lecture). The first round with no new bids ends the auction.

The main reason that SAAs work better than sequential or sealed-bid auctions is *price discovery*. As a bidder acquires better information about the likely selling prices of licenses, it can implement mid-course corrections — abandoning licenses for which competition is more fierce than anticipated, snapping up unexpected bargains, and rethinking which packages of licenses to assemble. The format typically resolves the miscoordination problems that plague simultaneous sealed-bid auctions. For instance, suppose there are two identical items and three bidders. Every round, some bidder will be losing both auctions. When it jumps back in, it makes sense to bid for the currently cheaper item, and this will keep the prices of the two items roughly the same.

Another bonus of the SAA format is that bidders only need to determine valuations on a need-to-know basis. We’ve been assuming that valuations are known to bidders at the beginning of the auction, but in practice determining the valuation for a bundle of items can be costly, involving research, expert advice, and so on. In sharp contrast to direct-revelation auctions, a bidder can often navigate an SAA with only coarse estimates for most valuations and precise estimates for the bundles that matter.

Generally, SAAs are believed to perform well, meaning they achieve good welfare and revenue. This assertion is not easy to test after an auction, since valuations remain unknown and bids are incomplete and potentially non-truthful. However, there are a number of “sanity checks” that suggest good auction performance. First, there should be little or no resale of items after the auction, and any reselling should take place at a price comparable to the auction’s selling price. This indicates that speculators did not play a significant role in the auction. Second, similar items should sell for similar prices (cf., the Swiss and New Zealand auctions). Third, revenue should meet or exceed projections. Fourth, there should be evidence of price discovery — for example, prices and provisional winners at the mid-point of the auction should be highly correlated with final selling prices and winners. Finally, the packages assembled by bidders should be sensible, such as groups of licenses that are adjacent geographically or in frequency range.

SAAs have two big vulnerabilities. The first problem is *demand reduction*, and this is relevant even when items are substitutes. Demand reduction occurs when a bidder asks for fewer items than it really wants, to lower competition and therefore the prices paid for the items that it gets.

To illustrate, suppose there are two identical items and two bidders. The first bidder has valuation 10 for one of the items and valuation 20 for both. The second bidder has valuation

8 for one of the items and does not want both (i.e., its valuation remains 8 for both). Giving both items to the first bidder maximizes the welfare, at 20. The VCG mechanism would earn revenue 8 on this example. Now consider how things play out in an SAA. Bidder 2 would be happy to have either item at any price less than 8. Thus, bidder 2 drops out only when both items have price at least 8. If bidder 1 stubbornly insists on winning both items, its utility will be $20 - 16 = 4$. If, on the other hand, bidder 1 targets just one item, then each of the bidders will get one of the items at a near-zero price. Bidder 1's utility is then close to 10. In this example, demand reduction leads to a loss of welfare and revenue, relative to the VCG mechanism's outcome. There is ample evidence of demand reduction in many spectrum auctions.

The second big problem with SAAs, which is relevant when items are complements (including in many spectrum auctions), is the *exposure problem*. As an example, consider two bidders and two non-identical items. Bidder 1 only wants both items — they are complementary items for the bidder — and its valuation is 100 for them (and 0 otherwise). Bidder 2 is willing to pay 75 for either item. The VCG mechanism would give both items to bidder 1, for a welfare of 100, and would generate revenue 75. In a SAA, though, bidder 2 will not drop out until the price of each item reaches 75. Bidder 1 is in a no-win situation: to get both items it would have to pay 150, more than its value. The scenario of winning only one item for a non-trivial price could be even worse. On the other hand, if bidder 2's value for each item was only 40, then bidder 1 should just go for it. But how can bidder 1 know which scenario is closer to the truth? The exposure problem makes bidding in an SAA difficult for a bidder for whom items are complements, and it often leads to risk-averse, tentative bidding by such bidders.

3 Bells and Whistles

A difficult and controversial question is whether or not to augment the basic SAA format by package bidding — bidding on sets of items in addition to individual items — and, if so, how. The primary reason to allow package bidding is to alleviate the exposure problem when items are complements, to free up bidders who desire bundles of items to bid aggressively for them. There are also scenarios where package bids can remove the incentive for demand reduction.

The conservative viewpoint, which dominated practice until relatively recently, is that package bids add complexity to a quite functional auction format and might do more harm than good. Limited forms of package bidding have been incorporated into spectrum auction designs only over the past 5–10 years.

One design approach is to tack on one extra “proxy” round after the SAA where bidders can submit package bids on any subsets of items that they want, subject to an activity rule; see Ausubel and Milgrom [1] for details. These package bids compete with each other as well as the winning bids on individual items from the SAA phase of the auction. The final allocation is determined by a welfare-maximization computation, treating bids as true values. The biggest issue with this approach is that computing the final prices is tricky.




Figure 1: After some of the TV broadcasters are bought out, the remaining ones will be repacked via channel reassignment to free up a contiguous portion of the spectrum.

The VCG payment rule is not used because of its poor revenue and incentive properties (see Lecture 7 and the exercises). A more aggressive payment rule, which is not DSIC but does have other good incentive properties, is used instead. Typical behavior of bidders with this relatively complex pricing rule does not seem to be well understood.

A second approach is to predefine a limited set of allowable package bids, rather than allowing bidders to propose their own. Ideally, the predefined package bids should be well-aligned with what bidders actually want, yet structured enough to permit reasonably simple allocation and payment rules. Hierarchical packages — for example, allowing bids on individual licenses, on regional bundles of licenses, and on nationwide bundles — have emerged as a sweet spot for this design approach [3]. The biggest issue with predefined package bids is that they can do more harm than good when they are poorly matched with bidders' goals. For example, imagine that you're a bidder who wants the items ABCD, but the available packages are ABEF and CDHI — what's your bidding strategy?

4 The Cutting Edge

We've reached the state-of-the-art of wireless spectrum auctions, so let's conclude with a peek into the future: an upcoming FCC double auction, to take place possibly in 2014.²

Wireless spectrum doesn't grow on trees. At this point, giving someone a new allocation of spectrum generally requires taking it away from someone else. Soon, the FCC plans to do precisely this, using a reverse auction (cf., Exercise 7) to free up spectrum by buying out TV broadcasters and then a forward auction to resell the spectrum to companies that can put it to more valuable use. The forward auction will likely be implemented as an SAA with bells and whistles, as usual; the reverse auction is completely new.

In addition, the FCC will repack the remaining broadcasters so that the freed up frequency is contiguous. For example, they might buy out a number of TV broadcasters across the nation who were using a UHF channel somewhere between 38 and 51, and reassign all of the remaining broadcasters to have a channel between 38 and 45, leaving the part of the spectrum corresponding to channels 46–51 free for new users (see Figure 1).

In a very cool development, the current frontrunner for the reverse auction format is a greedy approximate welfare-maximizing allocation rule, not unlike those we discussed for

²The auction format is still under discussion, and this section is only the author's best guess as to what will be adopted.

Knapsack auctions in Lecture 4. In the proposed model, each bidder i (a TV broadcaster) has a private valuation v_i for its broadcasting license. That is, v_i is the “minimum acceptable offer” for buying out i .³ Letting N denote the set of bidders, a set $S \subseteq N$ of winning bidders — where “winning” means being bought out — is feasible if the remaining bidders $N \setminus S$ can be repacked in the target range (e.g., channels 38–45).⁴ For instance, if $S = N$ then all bidders are bought out and the entire spectrum is freed up, so S is certainly feasible. When $S = \emptyset$, no spectrum is freed up, an infeasible outcome. Checking whether or not a given set S is feasible is a medium-size NP-hard problem — essentially the graph coloring problem, since two TV stations with overlapping geographic areas cannot be assigned the same or adjacent channels — so solving it requires state-of-the-art algorithmic technology. As of this writing, SAT solvers and integer programming solvers are battling it out, striving to solve these “feasibility-checking” problems as fast as possible (ideally, in seconds).

We give a direct-revelation description of the proposed class of allocation rules, although they can (and likely will be) implemented via an iterative auction with descending, bidder-specific prices. Descending implementations are preferred to sealed-bid implementations because, empirically, bidders find them easier to play. The allocation rule starts with the trivial feasible set (all bidders), and then iteratively removes bidders from the current feasible set until a minimal feasible set is reached. A greedy scoring rule is used to choose which bidder to remove in each iteration. One might call this a “reverse greedy algorithm,” since it deletes bidders starting from the entire set, rather than forward greedy algorithms which iteratively add bidders starting from the empty set (cf., Knapsack auctions in Lecture 4). Milgrom and Segal [5] call these *deferred allocation rules*.

- Set $S = N$. [Initially feasible.]
- While there is an $i \in S$ such that $S \setminus \{i\}$ remains feasible:
 - (*) Delete some such i from S . [I.e., i will not be bought out.]
- Return S .

Step (*) is obviously underdetermined, and it’s easy to think of various heuristics to try, like deleting the bidder with the highest bid (i.e., least willing to be bought out), the bidder with the highest per-capita bid, etc. The exact choice of the greedy rule will likely be guided by the welfare achieved by different rules on synthetic data.

³This single-parameter model assumes that each TV station is owned by a different strategic agent. This assumption is not entirely true in practice, but it makes the model much easier to reason about.

⁴One interesting question is how to set this target. The bigger the target, the bigger the expenses per unit of spectrum in the reverse auction and the smaller the revenues per unit of spectrum in the forward auction, since increasing supply should decrease the price. An ideal target would equalize the price per spectrum in the forward and reverse auctions — or perhaps with a somewhat higher price in the forward auction, so that auction expenses are recovered by the auction’s net revenue. One approach that is being discussed seriously is to use the proposed reverse auction to estimate the entire supply curve — the cost of acquiring spectrum for each possible target — and then match supply and demand accordingly during the forward auction.

If we implement (*) using a scoring function, deleting the bidder i with the largest score (subject to $S \setminus \{i\}$ being feasible), and if this scoring function is increasing in a bidder's bid and independent of the bids of the other active players, then the deferred allocation rule is monotone (bidding lower can only cause you to win). See Exercise Set #4. By Myerson's Lemma, paying critical bids — the largest bid that a winning bidder could have made and still gotten bought out — yields a DSIC auction.

Remarkably, deferred allocation rules have a number of good incentive properties above and beyond DSIC, which are not shared by their forward-greedy cousins [5]; see also Problem Set #3.

References

- [1] L. Ausubel and P. Milgrom. Ascending proxy auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 3. MIT Press, 2006.
- [2] P. Cramton. Simultaneous ascending auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 4. MIT Press, 2006.
- [3] J. K. Goeree and C. A. Holt. Hierarchical package bidding: A paper & pencil combinatorial auction. *Games and Economic Behavior*, 70(1):146–169, 2010.
- [4] P.. Milgrom. *Putting Auction Theory to Work*. Cambridge, 2004.
- [5] P. Milgrom and I. Segal. Deferred-acceptance heuristic auctions. Working paper, August 2013.