How Does Computer Science Inform Modern Auction Design? Scribe: Kiran Vodrahalli
Tim Roughgarden, Stanford Distinguished Lecture: 11/01/17

1 Introduction

For the past fifteen or twenty years, computer scientists have been learning a lot about economics and game theory. A lot of emerging computer science applications people cared about were interactions of different parties: relay networks, online markets for advertising, etc. Ideas seem to be flowing now in the opposite direction: From computer science to game theory/economics. I am going to tell you about a major auction which just finished, and highlight the role of computer science, both in its design and in its analysis.

2 The F.C.C. Auction

Congress in 2012 authorized F.C.C. to sell the spectrum to wireless people with auction. Using auction wasn't new, but spectrum was already owned by television owners. The goal was to repurpose the spectrum, take away from television and re-sell to more useful people.

The auction had two phases: A reverse auction, where government buys back licenses, and then a forward auction, which is the usual method. The reverse auction was totally new! This ended several months ago, and it turned out pretty well! They bought back licenses for 10 billion dollars, and sold for 20 billion dollars. A majority of it was applied towards the deficit. The bill that authorized the auction was called the "Middle Class Tax Relief and Job Creation Act".

The act was passed in Fall 2012 and was run in March 2016.

2.1 The Reverse Auction

It was proposed by Milgrom and Segal (was selected by FCC after a bunch of people made bids on different proposals): A descending clock auction. The goal was to make it as easy as possible for people to participate in. Mom-and-pop television broadcasters were parts of the target. For the auction to be a success, it was necessary to get healthy participation. So they wanted the auction to be really simple. It's an iterative auction, it works in rounds. In a round, you might be asked a question of the following form: Would you be willing to sell your license for, say, a million dollars. If you say no, you are kicked out of the auction forever more. You get to stay on the air. You can also say "yes". You don't necessarily get that price, but you get to live another day. You'll see a set of ever-decreasing prices until you leave. When the auction terminates, the goverment will buy your license. I'll specify how it ends.

The auction itself has a limited form of collusion-proof-ness, but collusion is hard to deal with, so it's dealt with legal channels. Broadcasters were not allowed to talk to anyone.

So the auction starts with very very high bidding prices – high enough such that anyone would be ecstatic to get that price. Of course the value drops over time. When does the auction stop? When prices are low as possible, subject to a clear enough spectrum. The spectrum is broken into chunks, allocated into TV channels. We want to clear a target number of channels. The decision the FCC made early on in the design process was that clearing a channel means clearing it nation-wide: There should not be a single network broadcasting nation-wide. Maybe an example goal is from channels $38 - 51$, clear 10 of them.

An issue is that buyouts are scattered across channels. So if you drop out of the auction, you are guaranteed to remain on the air, but you are not guaranteed to maintain the channel you were on. So what is the real problem? If you want to clear 10 channels nationwide, you want to buy up enough licenses such that you can re-arrange the spectrum so that you clear up a bunch. This is a repacking problem.

Each license is equipped with geographic region and frequency. Think of each station as a ball. Different colors correspond to different frequencies. After removing some balls, it's possible to re-assign colors such that you have a smaller number of colors (frequencies). So this is a graph coloring problem, which is NP-complete. So how is this actually going to work? Now, NP-completeness is not a death sentence, but you have to work a bit harder. So how big are these graph coloring problems? Not huge, but enough to be non-trivial: 2000 vertices, 10000 instances. How hard is that with current infrastructure? Not too bad to solve one instance. But this auction needs to solve thousands of these graph coloring problems every day! If you are going to give someone a lower offer, you need to be prepared for them to say no. They drop out of the auction, and now you are responsible for packing them into the number of channels. You have to speculatively think about them dropping out before you hear their offer. So you have to check what will happen: For every round, for every participant, you have to solve a graph coloring problem. The F.C.C. gave the auction designers of 1 minute to solve these graph coloring problems. To tackle this problem they brought on a computer scientist, who was tasked with solving these graph coloring problems. They encoded as an instant of SAT. So as a baseline, they used the best SAT solvers (winners of the most recent SAT competition) and see how well they do. They were able to solve 85% of them in less than a minute. What does it mean if you fail to solve in a minute? Then it's telling you don't know if it's possible to repack. In order to avoid the epic fail, you have to assume it's not possible to do the repacking: that's the conservative approach. Can't repack, better not lower the person's offer. If you had found the repacking though, you could have paid the person a lower price. Here is a true connection between the runtime of the algorithm and the cost.

Now why not 2 min budget? They argued about it, but F.C.C. wanted it to be as quick as possible. Even with 1 min, it took a year. It ran Monday through Friday.

Now at the very first moment, if the graph coloring problem says they cannot be repacked, then they will stay at that price frozen and not be lowered, since in between, only other people's prices will be lowered and the problem will only get harder (except in the case that the outer loop lowers the tolerance for how much spectrum you want to clear).

The F.C.C. decided to be opaque, so the players had no knowledge. Iterative auctions are much easier for people to play; answering yes-no questions is much easier than getting people to commit. This is why it's better than asking directly the minimum price you would accept.

So anyways, I was talking about how there is motivation to get from 85% to 100%. So the team tuned the solvers for the types of examples they knew they would face. They also did some caching tricks: Several examples are similar. Then, they were able to solve 99%. The point here is clear: Cutting edge techniques for NP-complete problems were both necessary and sufficient to execute a good auction.

The question was maybe you should preferentially prefer some stations over others, since them dropping out might impact the difficulty of the packing problem. This came into how they calculated the initial offer prices. Higher price if you make the problem harder to solve.

2.2 The Forward Auction

The computer science toolbox is good for explaining why auctions work. Let me begin with a cautionary tale. In New Zealand in the 1990s, the government decided to release 10 new television channels. They decided to use an auction called simultaneous second-price auctions – this is good for selling one item. In fact, eBay uses such an auction. Highest bidder wins and pays second highest bid overall – there is no reason to strategize.

But in parallel, it loses all of these nice properties. Imagine you had to write down ten numbers, one for each bidder. Imagine you only want one of these licenses. What should you do? Point is, it's not that obvious. Maybe pick a true value for only one auction, selected at random. If you think it's not that expensive of an auction, maybe you bid low on several of the offers and hope you get lucky. It's unclear how to behave, so it's likely you'll get poor outcomes. Their projected revenue was 250 million, actual was 36 million. On one channel, the high bid was $100,000$. The second highest, and the selling price, was 6 dollars. So people are very reticent to making changes to best practices in forward auction design.

So the English auction, someone names higher and higher prices. You keep your hand up – it's the one you see in the movies. You could imagine to sell 10 licenses that you have a auctioneer for each price. This is called a simultaneous ascending auction. It usually works decently, but there are some problems:

(a) Demand reduction. Let us say there are only two bidders. One of them has value 6 per license, and the second bidder wants one license with value 5. What's going to happen in simultaneous auction? The second bidder will always stick around until the price hits 5. This means first bidder gets value of 12 - 5*2 = 2. The first bidder can stick around and wait it out. But you can do something smarter: Give up on one auction, and get the other one at a much cheaper price. Then the first bidder gets value 6 - 0 = 6.

(b) Exposure problem with item synergies. You may want to guess whether you will be able to win an item which is a substitute for the current item, and use that to bid.

There are some folklore beliefs: If there are no item synergies, then the efficiency caused by demand reduction is not a deal breaker. If you do have strong item synergies, then the belief that simple auctions like simulteanous ascending auctions are deal-breakers.

When we talk about the outcome of the auction, we are talking about an equilibrium. We woul dlike ana pproximation guarantee for the equilibria of some game. Since the dawn of algorithmic game theory, a lot of computer scientists have thought about this – the price of anarchy, which Christos defined early. I was obsessed with this as a grad student. Obsession channeled properly can be really good for research.

3 Standard Model for Auctions with Mulitple Bidders

You have n bidders, m items. Each bidder i has a nonnegative valuation $v_i(S)$ for each subset S of items (2^m parameters). Bidder i wants to maximize $v_i(S_i) - $payment. Once you have a fixed auction, you have a game in the sense of game theory. How should we assess the quality of an equilibrium for this game? We like to look at social welfare. The social welfare of the allocation $S_1, \cdots, S_n$: $\sum_i v_i(S_i)$. We do not know what equilibria look like, so can there be hope for proving approximation guarantees for this kind of equilibria? The answer is yes. There are a lot of ways you can define "without strong complements", and definition of the auction, and approximately optimal.

The theory of price of anarchy that has emerged over past 15 years is surprisingly coherent. There's a user-friendly toolbox for proving these approximation arguments: Extension theorem for smooth games: Smoothness arguments. It is more of a philosophy, which involves two steps. You first think about a very simple case. Maybe you assume you know the valuations, or that you are only going to look at pure strategy eqilibria. The second part involves proving an extension theorem which relates the simple setting to the complex setting. You can impose theorems on the way you prove the approximation guarantee in the simple setting.

One way to assert things are not strong complements, you have subadditive valuations: $v_i(S + T) \leq v_i(S) + v_i(T)$ for all S, T.

Theorem 3.1 (Feldman/Fu/Gravin/Lucier [13). *]. Every equilibrium of a simultaneous first-price auction (SIA) has welfare at least* 50% *of the maximum possible.*

It turns out that you get 63% have submodular valuations. The take-away is that without strong complements, simple auctions work pretty well – we have a quantitative statement that supports this rule of thumb from before.

The theory of revenue-maximizing auctions is not used in this setting. The theory has trouble figuring out exactly what to do. The whole point of this is not to make all this money, but the spectrum is being allocated in the best way to serve society. In the F.C.C. auction, there was guarantee: go through a phase where you decide to clear less spectrum in the outer loop.

The second practical rule of thumb was that simple auctions are not good enough when you have complementary goods. This has the flavor of an impossibility theorem, a complexity theory statement.

Theorem 3.2 (Hassidim/Kaplan/Mansour/Nisan 11). . *With general valuations, SIAs can have equilibria with welfare arbitrarily smaller than the maximum possible.*

Also,

Theorem 3.3 (Roughgarden 14)**.** . *With general valuations, every simple auction can have equilibria arbitrarily smaller than the maximum possible. Here, "simple" means subexponential in m number of bids per player.*

Note that the number of bids is a way to measure complexity. For instance, in the original SIA, it's linear in m. If you bid on fairs, m^2. etc.

What is the proof of this theorem? In a moral sense, the result is true because equilibria have two properties that constrain them to be not that much more powerful than efficient communication protocols: They are guaranteed to exist and easy to verify. These allow you to translate results from complexity theory into results on auctions.

4 Takeaways

Computer science and auction theory have a lot of connections. We say in the reverse auction that designs were motivated by computational constraints: SAT solvers for fast solutions for repacking problems, and we used descending iterative auctions as reverse greedy algorithms. On the other side, the price of anarchy toolbox proves that simple auctions work well without complements. Communication complexity explains why simple auctions perform poorly without strong complements.

Reverse auctions are single parameter setting (resistant to manipulation) – people only have one license. In the forward auction, multiple parameter setting: people have a bunch of values in their heads. Not as resistant to strategy.

The execution of reverse auction will proceed in different ways/ different trajectories. You visit players in order – could you do things better if you picked the order? Or solve them all in parallel, and do whichever finishes first. They didn't do anything smart though, they had lots of ideas after it was over to implement. To solve a problem you need to know who's already dropped out.

The initial prices for networks was a heuristic combination (geometric mean) of the degree of the station (how much it overlaps with other stations) and the market price for the valuation.

5 My Questions

(a)