

Designing the Perfect Auction

Distributed algorithmic mechanism design is a field at the intersection of computer science and economics.

IN JANUARY, THE U.S. Federal Communications Commission (FCC) launched one of the biggest auctions in history, selling off what some in the wireless industry have called the “beachfront property” of the electromagnetic spectrum. Its auction of the 700MHz frequencies, the largest and most valuable slice of spectrum to come available in years, brought in more than \$19 billion.

With the auction, the FCC inaugurated the first use in a major spectrum auction of “package bidding,” in which bidders are allowed to bid either on individual state licenses or regional packages of licenses. Although package bidding makes a lot of sense—for example, some bidders might be interested in buying state A only if they can be guaranteed to also get state B—until now, the FCC had not offered an auction design that gave bidders enough flexibility while keeping combinatorial and computational complexity in check.

The auction is a high-profile example of distributed algorithmic mechanism design, a field that combines economics and algorithm design. Economic mechanism design is

concerned with how to design a market or market-like institution so that it will achieve a desired goal, such as allocating goods efficiently, maximizing profit, or achieving an equitable distribution. Mechanism design is, in a sense, the inverse of game theory: in game theory, one is given the rules of a game and the goal is to predict the outcome; in mechanism design, one

is given a set of desired outcomes and the goal is to design a game that will achieve them.

In 2007, Leo Hurwicz, Eric Maskin and Roger Myerson were awarded the Nobel Prize in Economics for their work in economic mechanism design. Their theoretical work underpins a host of practical applications, including eBay’s auctions, the auctions used by Google and Yahoo! to sell ad slots, the matching system used to pair medical residents and hospitals, the California electric power exchange, the rules governing trades on NASDAQ and other financial markets, and, of course, the FCC spectrum auctions.

As auctions have grown more complex and more economic transactions



occur online, computational considerations in mechanism design are taking on a growing significance. Distributed algorithmic mechanism design, or just algorithmic mechanism design, is emerging as a field in its own right.

“Mechanism design is one of the major intellectual interfaces between computer science and economics, as well as one of the most vibrant areas of economics,” says Christos Papadimitriou, a computer scientist at the University of California, Berkeley.

For Your Consideration

Broadly speaking, there are two main strands in the literature. The first involves bringing computational considerations to the economics mechanism design literature. The second involves bringing incentive considerations to the computer science literature.

As an example of the first issue, consider the recent spectrum auction by the Federal Communications Commission mentioned earlier. In this auction, the right to use spectrum in various locations was sold to mobile phone companies and other potential users. The valuation that a buyer places on spectrum in a particular location may depend strongly on whether or not it wins spectrum in other locations.

In theory, each buyer could assign a different value to each possible subset of the geographic locations being sold. How does one design an auction that will yield reasonable outcomes in such a “combinatorial auction”? In such auctions, it turns out that the so-called “winner determination problem” is, in general, NP-complete. However, researchers working in algorithmic mechanism design have discovered various approximation algorithms and special cases that allow for reasonably good solutions in practical examples.

As an example of the second issue, consider the famous “stable marriage problem” in which one wants to design an algorithm to match up men and women. Each man has a ranking over the women, and each woman has a ranking over the men. A stable assignment is one such that no couple would prefer to leave their current mates to form a new couple. Although this particular description may sound somewhat frivolous, there are much more serious examples, such as matching

A good starting point for studying distributed algorithmic mechanism design is a simple auction.

up hospitals and residents or organ donors and recipients.

It turns out that stable assignments always exist, and there are a number of algorithms that compute them. However, these algorithms assume that the participants are truthfully revealing their rankings. Do they actually have the appropriate incentives to do so? It turns out that some algorithms provide such incentives to men, and some provide such incentives to women, but there is no algorithm that provides incentives for both sides of the market to be truthful.

Ideas from economics can shed light on many computer science problems that arise from user interactions, such as computer viruses and spam, says Preston McAfee, a researcher at Yahoo! Research in Burbank, CA. “I think there’s a growing recognition that problems of bad behavior are incentive problems in the realm of game theory, rather than technological problems in the realm of traditional computer science,” he says.

Understanding the effect of incentives on how algorithms perform is “the latest and most momentous twist” on the question of computation’s limits, Papadimitriou says.

“With classical algorithms, you get your inputs and then compute away, and the answer comes out,” he says. “In this new context, you have to get your inputs by peering into the souls of selfish agents trying to promote themselves.”

A Simple Auction

A good starting point for studying distributed algorithmic mechanism design is a simple auction. A seller has one item to sell and n buyers have values

v_1, \dots, v_n for this item. The seller may have a reserve price r , which is the minimum price at which he is willing to sell the item. Typically there will also be a bid increment, the minimum amount by which a bid may be changed.

The goal is to design an online auction that will achieve some desired goals. There are many types of auctions that could be used. They include:

English auction. The seller starts at r and progressively raises the price by the bid increment until all but one of the buyers drops out. This is the most common form of auction.

Dutch auction. The seller starts at a high price and progressively lowers the price by the bid increment until a buyer shouts out “buy.” This sort of auction is used to sell flowers in the Netherlands.

First-price sealed bid. The buyers write down a bid and seal it in an envelope. The envelopes are opened and the item is awarded to the highest bidder at the price he or she bid. This form is commonly used for construction contracts.

Second-price sealed bid. The buyers write down a bid and seal it in an envelope. The envelopes are opened and the item is awarded to the highest bidder at the second-highest price. This auction was used by stamp collectors in the 19th century to sell stamps by mail.

It turns out there are some relationships among these auctions. For example, with fully rational players, the outcome of the English auction is the same, up to the bid increment, as the outcome of the second-price sealed bid auction. This is, perhaps, not so surprising upon reflection, as the English auction ends up awarding the item to the bidder who is willing to go the highest, but he or she only has to pay the bid of the second highest bidder plus a possible bid increment.

To make the argument slightly more precise observe that the payoff to a bidder with value v_1 is $v_1 - b_2$ where b_2 is the bid of the second-highest bidder. There are three cases to consider:

► $v_1 > b_2$. In this case bidder 1 wants to win. But the bidder can do so by reporting $b_1 = v_1$.

► $v_1 < b_2$. In this case bidder 1 wants to lose. But the bidder can do so by reporting $b_1 = v_1$.

► $v_1 = b_2$. In this case, bidder 1 is indifferent about winning or losing, so the bidder may as well report $b_1 = v_1$.

In each case it is optimal for bidder 1 to report his or her true value, regardless of what the bidder thinks other bidders will do. This is known as a *dominant* strategy in game theory. If everyone reports their true value, the item ends up being awarded to the bidder with the highest value, which is the efficient outcome in the sense of maximizing the value of the assignment.

The auction used by eBay is basically a form of second-price auction; the bidder who programs his or her bidding agent with the highest value wins, but only has to pay the second highest bid.

Note that in both of the examples mentioned—the 19th century stamp collectors and the eBay auction—the underlying motivation for adopting this auction form was communication costs. The stamp collectors did not want to mail bids back and forth and eBay buyers did not want to log on every time they wanted to change their bid.

Combinatorial Auctions

To continue with auctions, let us imagine a much more complex problem in which many items are to be sold. Let x represent an assignment of goods to bidders and let $v_a(x)$ represent agent a 's valuation—the agent's willingness to pay—for a given assignment. In principle, each agent may care not only about what he or she gets in the assignment, but also what everyone else gets. The seller does not know the bidders' valuation functions.

The auction design goal is to assign the items to the agents in a way that maximizes the sum of the individual valuations of the assignment.

Perhaps surprisingly, this mechanism design problem can be solved in much the same way as the single item auction. We simply ask each person to report their valuation functions. Next, we find the assignment that maximizes the sum of the reported valuations. The payment that agent a makes is the difference between the maximal value to the other agents if agent a is present and the maximal value if agent a is removed from the calculation. Roughly speaking, each agent has to pay the

cost that his or her presence imposes on the other agents.

To see how this generalizes the previous simple auction, note that in the simple auction the price that the highest bidder has to pay is the cost he or she imposed on the other agents; if the highest bidder weren't present, the second highest bidding agent would receive the item. This mechanism is known as the Vickrey-Clarke-Groves or VCG mechanism. It provides incentives to report true values for virtually any sort of problem. Of course, it also has flaws. For example, it does not typically generate the maximum amount of revenue for the seller.

Search Engine Ad Auctions

Google, MSN, and Yahoo! all use an auction to sell ad space on their search engines. Advertisers bid for positions on a search results page with the highest bidder receiving the most prominent position. The second highest bidder gets the second most prominent position and so on. Each advertiser pays a price per click based on the bid of the advertiser below him or her.

It turns out that there is no dominant strategy in this game when more than two positions being auctioned off. However, it is possible to find outcomes that are "stable" in the sense that no agent wants to change his or her bid, given the bids of the other advertisers.

Designing online auctions has been an evolutionary process, says Alvin Roth, an economist at Harvard University. "Google's design came out of some earlier designs, and getting the right design has been an important part of its success," he says. "It has helped create a market that didn't exist before."

Distributed algorithmic mechanism design offers an interesting theoretical framework for incorporating incentives into algorithmic design. It also offers exciting opportunities for interdisciplinary collaboration as well as being highly relevant to important practical problems, such as auctioning off the popular 700MHz frequencies. ■

Hal R. Varian is the chief economist of Google.

Berkeley, CA-based science and technology writer Erica Klarreich provided additional reporting.

Information Technology

Video Search, Intel Style

Researchers at Intel labs in China and the U.S. are developing a video search technology that will enable users to search images in videos by person and object. Intel's video search technology divides videos on a frame-by-frame basis and uses image and face-recognition software to identify and categorize faces, voices, objects, locations, and movements. Next, the videos are reassembled to facilitate video search.

With Intel's video search technology, users will no longer have to fast forward through or watch an entire video, but can instantly cut to a particular scene or scenes. In addition to enabling users to instantly analyze videos, Intel's objective is to create a visual computing platform in which people can interact with a personal computer in a life-like, 3D environment.