Notes for 22c196

May 1, 2012

Auctions and Matching

An *auction* is a price setting mechinasm, or a way to exchange goods between individuals. Often, goods have a predetermined priceassigned to them; whether a buyer is shopping for a new shirt, groceries, or any other commonly available good, the price is fixed and static. Granted, a savvy shopper will "shop around" in an effort to find the best price between stores, but each store has a price that incorporates the production cost, distributor's cut, etc.

Naturally, the buyer wishes to minimize his spending, and the store to maximize its profit. This is mitagitated when, in the above example, we assume the items are commonly available. There is an understanding of what price a desired item merits. With 1-of-a-kind items, however, the buyer can no longer shop around. The buyer wants a price that is low enough that the buyer believes they are geting a good deal, and the seller wants a price that is high enough to make an acceptable profit.

An auction is a mechanism to establish such a price. We will focus primarily on English auctions (ascending price) and Dutch auctions (descending price). In English auctions, buyers bid the price they are willing to pay. In Dutch auctions, sellers lower the asking price until a buyer is willing to buy. We can concieve variants, such as the London Bus Systems where the value of an item increases when sold in combination with another item. (The collection is valued greater than the sum of its parts.) Likewise, auctions could consist of many or single buyers, many or single sellers, etc.

While auctions are often dynamic in nature, we consider static counterparts. Each potential buyer will submit a bid, and all bids are reviewed simultaniously. In a *First Price Auction*, the highest bidder wins the item and pays the value of her bid. In a *Second Price Auction*, the highest bidder wins the item and pays the value of *the second highest* bid. (Second Price Auctions are also refered to as Vickery Auctions, in honor of the 1996 Nobel Prize winner Vickery that developed them. Intuitively (we show this formally below), Second Price Auctions incentivize bidders to bid their true value of the item; a winning bid will always pay less than its value (exclusing identical bids). In First Price Auctions, however, bidders are encouraged to "shade" their bids in an effort to pay less.

We begin with the following assumptions:

- We assume there are many buyers for a single item.
- Each buyer i has some intrinsic valuations v_i for the item. These values are hidden from other bidders.
- Each buyer i will bid the value b_i. These values are also unknown to the other bidders.

Bidding Strategy for Second Price Auctions

A *bidding strategy* is a function that maps the hidden valuations to the offered bid. We show that an honest bidding strategy dominates in Second Price Auctions. We proceed casewise:

Case 1: Suppose bidder i bids honestly, with $b_i = v_i$.

If bidder i loses, she has payoff zero; nothing lost, nothing gained.

If bidder i wins, she pays the second highest bid, $b_k < b_i$. Note, this implies $b_k < v_i$. Then bidder i has payoff $v_i - b_k > v_i - v_i = 0$. Then the bidder has positive payoff.

Case 2: Suppose bidder i overbids, with $b_i > v_i$.

If bidder i loses, she has payoff zero.

If bidder i wins, she pays the second highest bid, $b_k < b_i$. If $b_k \geq v_i$, bidder i is still happy: $v_i - b_k \geq v_i - v_i = 0$. However, if $b_k > v_i$, bidder i experiences payoff $v_i - b_k < 0$. (The buyer has overpaid and has Buyer's remorse.)

In all cases, this strategy yields payoff bounded by the honest strategy in Case 1.

Case 3: Suppose bidder i shades her bid, with $b_i < v_i$.

If bidder i wins, she pays the second highest bid, $b_k < b_i$. Note, this implies $b_k < v_i$. Then bidder i has payoff $v_i - b_k > v_i - v_i = 0$. Then the bidder has positive payoff.

If bidder i loses, she has payoff zero. Let b_j denote the winning bid. If $b_j > v_i$, bidding honestly would not have helped. If, howveer, $b_j < b_j < v_i$, bidder i has received a payoff of zero, where bidding honestly would have won her the item and resulted in a positive payoff.

In all cases, this strategy yields payoff bounded by the honest strategy in Case 1.

Since both Case 2 and Case 3 are dominated by Case 1, we have that bidding honestly is the dominating strategy for Second Price Auctions.

Bidding Strategy for First Price Auctions

If bidder i loses, she has payoff of zero. If bidder i wins, she has payoff $v_i - b_i$.

- If bidder i bids honestly, she has payoff zero in both cases.
- If bidder i bids $b_i > v_i$, bidder i has negative payoff.
- If bidder i bids $b_i < v_i$, but still bids high enough to win, she receives payoff $v_i - b_i > 0$.

The only scenario in which bidder i receives positive payoff is to shade her bet. If she bids close to her true value and wins, her payout is small and she experiences buyer's remorse; perhaps she could have won bidding less. If she bids far less than her true value, she runs the likelihood of losing the auction. The question becomes the level at which to shade. The bidder must find some way to judge the competition: the number of bidders and estimate their values or bids.

We have assumed that all bidders are playing the same strategy, that all valuations v_i are secret, and that the seller will sell. In this case, we find an equilibrium strategy to shade the bid $b_i = \frac{n-1}{n} v_i$.

Seller Performance

Assume n bidders have valuations are drawn uniformly at random in the interfal $[0, 1]$. Index bidders in increasing bid value: $b_1 < .... < b_n$. It follows from the random valuations that the expected bid of bidder i can be given $E[v_k] = \frac{k}{n+1}$.

Now, in a Second Price Auction, bidder n will win the bid and pay the price of bidder $n-1$, which we have at the expected price of $\frac{n-1}{n+1}$.

Conversely, in a First Price Auction, bidder n will pay their own bid. Bidder n has valuation v_i expected to be $\frac{n}{n+1}$. However, bidders will shade their bids. There exists an equilibruim strategy of shading by a factor of $\frac{n-1}{n}$. This this case, the winning bid will give the seller a profit of $\frac{n}{n+1} \cdot \frac{n-1}{n} = \frac{n-1}{n+1}$.

We have seen that both First Price Auctions and Second Price Auctions will yield the same profit to the seller when valuations and bids are hidden. However, if a seller can convince the bidder to shade lass than the above equilibrium strategy, the seller will prefer a First Price Auction and hope for a better price. If the seller cannot "rub" the buyer, a seller will accept a Second Price Auction. This is called "Revenue Equivalence."

It should be noted that similar calculations can be made for other auction types. For instance, an *All Pay Auction* is one in which a bidder pays their bid value, regardless of whether or not theirs is the higest bid. Formally, they have a losing payoff of $-b_i$ and a winning payoff of $v_i - b_i$. These and other auction types are used to control bidder behavior. For instance, if we had to pay a fee to apply for a grant, we would apply for fewer. It should also be noted that Revenue Equivalence exists for these types of auctions.

We assumed that the seller is always willing to sell. If a seller withdraws, the seller can learn a buyer's valuations. Presumably, the seller can then take advantage of this information. However, the seller, presumably, has his own valuation of the item begin sold and would refuse to sell for little or no profit.

To account for this, we can extend Second Price Auctions with a *reserve price*. In this auction, the seller acts as a bidder as well; the bidder "bids" a reserve price r. Because this is a Second Price Auction, we have honest bidding as the dominating strategy. As such, we accept $r = u$, where u is the (seller's) true valuation of the price. Conversely, in a First Price Auction, the seller sets $r \geq u$, akin to a bidder's desire for shading. Just as a bidder does not shade so much that they risk losing the auction, so too will a seller have insentive to place $r \approx u$, to increase the likelihood that he is able to sell the item.

One can also construct All Pay auctions with reserve. Refer to Chapter 9 in the text.

Matching

We shift our focus now to matching algorithms in bipartite graphs. We assume there are $2n$ nodes (n on each side of the bipartite graph). This assumption is relaxed later. Edges denote preference.

For instance, one matching problem could match a set of X faculty to be assigned to the set of Y offices. Each faculty member $x \in X$ denotes which offices $y \in Y$ are acceptable. This generates the edge (x, y). We assume that all edges are of uniform weight. A faculty member either prefers an office or does not; offices are not ranked. We then want to maximize some utility, in this case "rank weighted happiness." Each faculty member has some ranking, i.e. Full Professors are ranked higher than Associate Professors, so their happiness is of higher priority.

A *perfect matching* is one in which every faculty would be assigned an office of their preference. The Augmented Path Algorithm is one algorithm that can detect perfect matchings.

(1) Select some edge.

(2) Find a path linking to unselected nodes, either directly, or using one of the already selected edges.

(3) Deselect the previously selected edges used.

(See Figure 2 for an example).

When a perfect matching exists, this algorithm finds one in $O(|V|\ |E|)$. However, there are certainly cases in which perfect matchings do not exist.

Theorem 1 (Matching Theorem)**.** *If there is a perfect matching if and only if there is no* constricted set*.*

A constricted set is one in which i nodes have (between them all) $k < i$ neighbors. Refer to Figure 1 for an example.

There exist methods for weighted edges, as well. Refer to the Hungarian algorithm ($O(n^4)$) or the Munkres algorithm ($O(n^3)$). Both, however, presuppose the presense of a central planner. The solution is found by the algorithm acting upon the graph, not by nodes (agents, or individuals) with only local information within it. We extend to local methods using auctions next time.

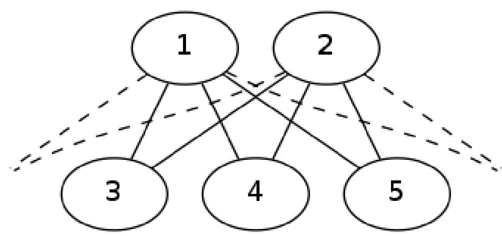

Figure 1: An example constricted set. Here, nodes 1 and 2 may have other neighbors, but nodes 3, 4, and 5 have only 2 neighbors between them, providing a constricted set.

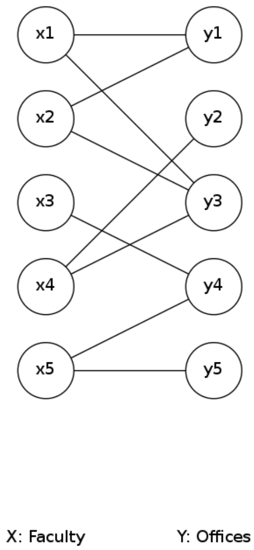

(a) Possible Input

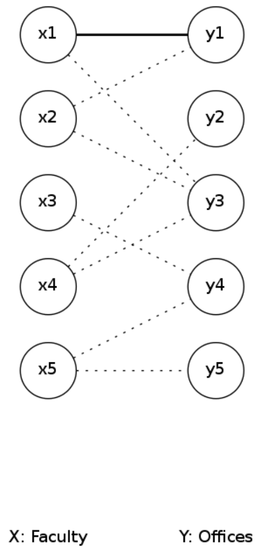

(b) Augmented Path (1)

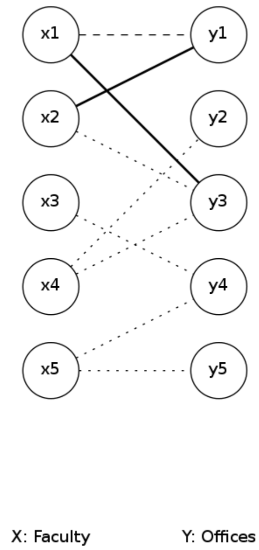

(c) Augmented Path (2)

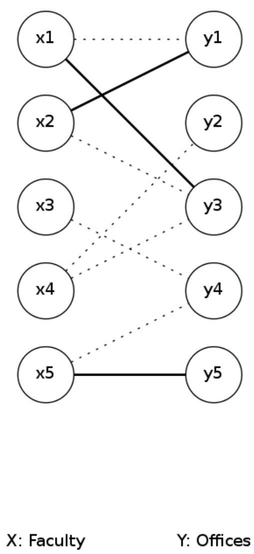

(d) Augmented Path (3)

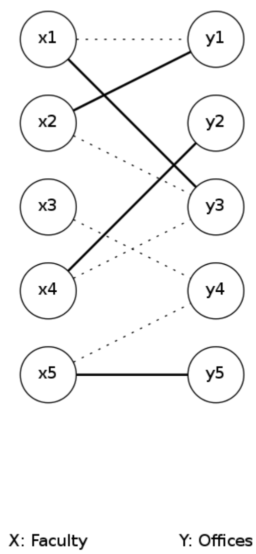

(e) Augmented Path (4)

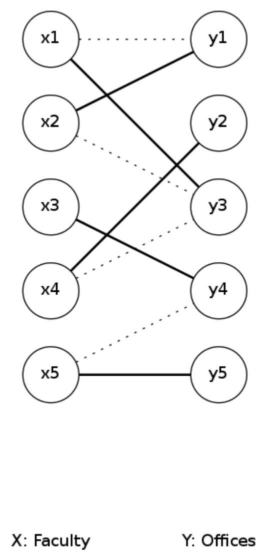

(f) Augmented Path (5)

Figure 2: Here is one example progression of the Augmented Path Algorithm. In Step 2, we remove the current edge and add two new edges, one incident on each of the previous edge's endpoints. In each subsequent step, no such removals are necessary.