CSC304: Algorithmic Game Theory and Mechanism Design
Fall 2016

Allan Borodin (instructor)
Tyrone Strangway and Young Wu (TAs)

November 9, 2016

Lecture 16

- Announcements
  - I will soon be posting a 2-3 questions for the final assignment.
  - The tests have been graded; average around 70%.
  - Usual policy as to requests for regrading.
- Todays agenda
  - Begin chapter 17 discussion of (one-sided) matching markets.

Begin matching markets

So far, in our mechanism design for auctions, we have been considering one seller (which we can view as the mechanism) and many buyers. We did give one example of a procurement auction which essentially just flips the roles of seller and buyer. And Tyrone spoke about an equilibrium analysis for many sellers and one additive or submodular buyer.

We now want to consider an auction setting with many buyers and many sellers. Namely, we want to consider *matching markets* and more specifically, we will be focusing attention on *one-sided matching markets*. And, more specifically, we wlll be restricting attention to *unit demand* buyers. This is the material discussed in Chapetr 17. Some relevant material has already occured in section 3.2 which we discussed early in the term.

Begin matching markets

So far, in our mechanism design for auctions, we have been considering one seller (which we can view as the mechanism) and many buyers. We did give one example of a procurement auction which essentially just flips the roles of seller and buyer. And Tyrone spoke about an equilibrium analysis for many sellers and one additive or submodular buyer.

We now want to consider an auction setting with many buyers and many sellers. Namely, we want to consider *matching markets* and more specifically, we will be focusing attention on *one-sided matching markets*. And, more specifically, we wlll be restricting attention to *unit demand* buyers. This is the material discussed in Chapetr 17. Some relevant material has already occured in section 3.2 which we discussed early in the term.

In a one-sided market, only one side (for us now the buyers) are strategic and have private valuations. In the more general setting of two-sided markets, sellers may also be strategic. For now, we restrict attention to one-sided markets.

One-sided matching markets

A one-sided matching market consists of the following ingrediants:

- A set of n strategic buyers and m sellers.
- Each seller has an item to sell (and hence we can alternatively think of this side as a set of items). Each buyer has a unit demand valuation function. That is, each buyer is interested in obtaining at most one item and has a value (possibly 0) for each possible item. (In a more general market, buyers have valuations for sets of items as we had in CAs.)
- The main objective is to maximize the social welfare (i.e. the sum of the valuations of the buyers for the items that have been allocated/assigned).

We let $v_{i,j}$ denote the value of buyer i fior item j. We visualize such a setting as an edge weighted bipartite graph, the graph model considered in section 2 of Chapter 3.

Weighted one-sided matching markets

As stated so far, we are considering arbitrary non-negative real values for the $\{v_{i,j}\}$. This then is the general weighted one-sided matching market setting and the underlying combinatorial allocation problem is to find a maximum weighted matching in the bipartite graph. This is refered to as *the assignment problem* in the combinatorial optimization literature.

Weighted one-sided matching markets

As stated so far, we are considering arbitrary non-negative real values for the $\{v_{i,j}\}$. This then is the general weighted one-sided matching market setting and the underlying combinatorial allocation problem is to find a maximum weighted matching in the bipartite graph. This is refered to as *the assignment problem* in the combinatorial optimization literature.

There are polynomial time optimal algorithms for this problem. Hence we could use VCG to find an optimal allocation and set prices to achieve a truthful mechanism.

However, we will take what is in some sense a more market based approach and see how prices might evolve to achieve a maximum matching without taking bids. (Recall how the Vickrey auction for a single item can be viewed as an ascending English auction. Instead the auctioneer in some sense keeps asking who is still interested in the item.)

Weighted and unweighted matching markets

For our purpose, there is not much loss in generality in thinking of the $\{v_{i,j}\}$ as rational values, and we can then scale so that all values are non-negative integers.

As a special case, we can also consider the case when all $\{v_{i,j}\}$ are in $\{0, 1\}$. This unweighted maximum matching problem corresponds to buyers either wanting or not wanting a particular item.

One can also consider buyers having preferences, and not values, for the differenct items. This becomes more of a topic in social choice which will occupy much of the last third of the course.

The ascending auction for a weighted matching market

Suppose we have n buyers and m sellers. Since we are not considering strategic sellers we can just say that we have m items for sale (by m different sellers).

Without loss of generality we will assume $n = m$. For example if there are $n < m$ buyers then add $m - n$ "dummy buyers" each having value 0 for all items. Similarly for $m < n$ items, add $n - m$ "dummy items", each having no value for any buyer.

The ascending auction for a weighted matching market

Suppose we have n buyers and m sellers. Since we are not considering strategic sellers we can just say that we have m items for sale (by m different sellers).

Without loss of generality we will assume $n = m$. For example if there are $n < m$ buyers then add $m - n$ "dummy buyers" each having value 0 for all items. Similarly for $m < n$ items, add $n - m$ "dummy items", each having no value for any buyer.

Our goal is to construct a vector of prices $\mathbf{p} = p_1, \ldots, p_n$ for the items so that every buyer will receive an item in her *demand set*. The demand set of items for agent i is $\{j : v_{i,j} - p_j \geq v_{i,k} - p_k \text{ for all } k\}$. That is, given the price vector $\mathbf{p}$, the items in the demand set are those items providing the maximum utility to the buyer.

The ascending auction for a weighted matching market continued

If we can set prices so that everyone is allocated an item in their demand set, then we will see that this matching will achieve the optimum social welfare. It will also result in an envy free allocation. It will also insure that the “market clears” in that if $n \leq m$, all items having non zero value for someone will be purchased.

Note: Hopefully we will later return to the issue of market clearance when buyers are not necessarily unit demand.

The basic idea of the ascending auction is (just like for one item in the English auction) to gradually raise prices on some items so that everyone obtains an item in their demand set.

The demand graph

Since this is a matching problem, lets express the problem in terms of graph theory.

The demand graph

Given a vector of prices $\mathbf{p}$, the demand graph $D(\mathbf{p})$ is the unweighted bipartite graph where there is an edge between buyer i and item j if and only if j is in the demand set of i.

Our problem then reduces to constructing a vector of prices $\mathbf{p}$ so that there is a *perfect matching* in the n by n demand graph $D(\mathbf{p})$. The ascending auction then raises prices until the demand graph has a perfect matching. It remains to show that this can always be done and how to compute a perfect matching when one exists.

Recalling basic graph definitions

We recall that in a graph $G = (V, E)$, a matching is a set of edges $E' \subseteq E$ such that no vertex appears in more than one edge in E'. For a bipartite graph with $V = X \cup Y$ and $E \subseteq X \times Y$, a matching E' satisifes that every vertex in X and Y has degree at most 1 in $G' = (V, E')$.

A perfect matching in a bipartite such that every vertex is in the matching and this of course requires $|X| = |Y|$. A perfect matching is of course a maximum matching but not necessarily conversely.

We will show how to efficiently find a maximum matching in a unweighted bipartite graph and we will then be able to know if that matching is perfect or not.

Moreover, if the maximum matching is not perfect, the algorithm will exhibit a *constricted set* that is preventing a perfect matching. (In our application of matching markets, we will have found a *constricted set* of buyers which will determine a corresponding set of items whose prices need to be raised.

Constricted sets

Let $G = (V, E)$ be a $n \times n$ bipartite graph. Let $V = X \cup Y$ and for any subset $S \subseteq X$, let $N(S) = \{y \in Y : (x, y) \in E \text{ for some } x \in Y\}$. That is, $N(S) \subseteq Y$ is the *neighbourhood* of S.

Constricted sets

$S \subseteq X$ is a constricted set if $|N(S)| < |S|$

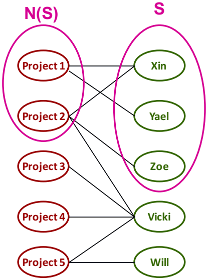

Figure: Example of a constricted set

An ascending auction algorithm for a matching market

We will do things a little different than the KP text. One simplification is that we are assuming integral valuations, so when we need to raise prices we just increases prices by 1. Our auction can raise more than one price in an iteration. But mainly the auctions follow the same idea and are based on the same basic result:

Hall's Marriage Theorem

An $n \times m$ (with say $m \leq n$) bipartite graph $G = (V, E)$ with $V = X \times Y$ has a matching of size m if and only if for all $X' \subseteq X, |N(X')| \geq |X'|$

Hall's Theorem is stated and proved in Theorem 3.2.2 of the KP text. An immediate consequence is that an $n \times n$ bipartite graph has a perfect matching if and only if there are no constricted sets.

An ascending auction template for a matching market continued

We let $\{v_{i,j}\}$ be the value of buyer i for item j. We let X be the set of buyers and Y the set of items.

An ascending auction for a matching market

Set the price vector $\mathbf{p} = (0, 0, \ldots, 0)$.
Let $D(\mathbf{p})$ be the demand graph.
Repeat until D has a perfect matching
Find a constricted set $S \subseteq X$ and raise the prices of all items in $N(S)$ by one unit. (Note: there can be many constricted sets)
Create a new demand graph for the updated prices
End Repeat

Theorem: The ascending auction terminates. A perfect matching in $D(\mathbf{p})$ is an envy-free allocation since every buyer is getting an item in their demand set.

Social welfare and termination in the ascedning auction

We will show termination by what is called a *potential argument*. And in doing so, will show that the algorithm results in an allocation that is socially optimal.

Given the current prices $(p_1, \ldots, p_n)$ and the valuation profile of the buyers, define

- The potential of a buyer i is the utility $v_{i,j} - p_j$ for some item j in buyer is demand set.
- The potential of an item/seller j is its price p_j.
- The potential P_t of the auction (at any iteration t in the algorithm) is the total sum of all buyer and seller potentials.

Social welfare and termination in the ascedning auction

We will show termination by what is called a *potential argument*. And in doing so, will show that the algorithm results in an allocation that is socially optimal.

Given the current prices $(p_1, \ldots, p_n)$ and the valuation profile of the buyers, define

- The potential of a buyer i is the utility $v_{i,j} - p_j$ for some item j in buyer is demand set.
- The potential of an item/seller j is its price p_j.
- The potential P_t of the auction (at any iteration t in the algorithm) is the total sum of all buyer and seller potentials.

Note that initially, the potential P_0 equals $\sum_i \max_j(v_{i,j}$.

Social welfare and termination in the ascedning auction

We will show termination by what is called a *potential argument*. And in doing so, will show that the algorithm results in an allocation that is socially optimal.

Given the current prices $(p_1, \ldots, p_n)$ and the valuation profile of the buyers, define

- The potential of a buyer i is the utility $v_{i,j} - p_j$ for some item j in buyer is demand set.
- The potential of an item/seller j is its price p_j.
- The potential P_t of the auction (at any iteration t in the algorithm) is the total sum of all buyer and seller potentials.

Note that initially, the potential P_0 equals $\sum_i \max_j (v_{i,j}$.

Claim: After each iteration the potential decreases by at least 1. Why?

Conclusion: The auction must terminate and a perfect matching in the resulting demand graph provides a socially optimal allocation. Why?

An example of the ascending auction

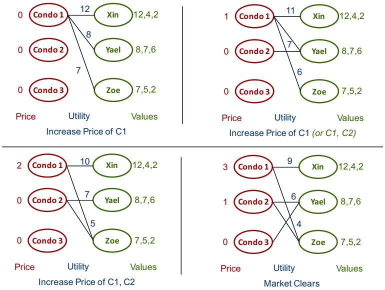

Figure: Example of the ascending auction