CSC304: Algorithmic Game Theory and Mechanism Design
Fall 2016

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

November 21, 2016

Lecture 19

- Announcements
  - I have posted 5 questions for the last assignment.
    New due date: noon, Wednesday, November 30.
  - Clarification for question 2. For termination analysis, add a minimum price resetting step (see next slide). I have clarified question 3 as to the price reduction step that I didn't mention in lectures 16 and 17.
  - I should note that students are responsible for material discussed in the lectures and tutorials whether or not that material is available in the lecture slides or text.
  - Quick clarification of two questions asked on piazza re the last assignment. For question 1, we do not have to specify $v_1(\{a, b\})$ as we assume free disposal. For question 4, I will give an example today of how lying can help. The man should not lie by saying he prefers no one as then he will get no one.
- Todays agenda
  - Mechanism design without money; continue stable matching
  - If time permits, begin voting theory. **Note:** Rushing some topics now now so that the assignment can be done.

An ascending auction template for a matching market revisited

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)
% For termination analysis reduce all prices uniformly by one unit so that lowest price is 0.
Create a new demand graph for the updated prices
End Repeat

Theorem: The ascending auction terminates. Note that the demand graphs in each round (and the potential function used for termination analysis) are the same whether or not we do the price reduction step.

Review: The Gale-Shapley Deferred Acceptance algorithm

It turns out that there is a relatively efficient was to construct a perfect (or maximum size) matching that is stable.

The Deferred Acceptance (DA) algorithm comes in two versions; namely, the Female Proposes (FPDA) and Male Proposes (MPDA) algorithms.

For simplicity, lets assume $n = n'$ so that our goal is to establish a perfect matching that is stable. (If $n \neq n'$, then the goal is to find a stable matching of size $\min(n, n')$.)

For definiteness, lets consider the FPDA. The MPDA is the same with the roles of men and women reversed.

The Female Proposes Delayed Acceptance FPDA algorithm

The FPDA algorithm

Initialize the *currently engaged CE list* to be empty;
that is, all males and females are *not engaged*.
Initialize every woman w's *already proposed* AP_w *list* to be empty
In rounds DO until every female is engaged

1. Every not engaged female w proposes to the man $m \notin AP_w$ that is most preferred in her preference ranking $\succ_w$.
2. For each man m, let $PT_m(t)$ be the set of women who have proposed to m in this round t plus the woman w he is currently engaged to (if engaged)
    - 2A. Let w^* be the woman most preferred in $P_m(t)$
    - 2B. Add (m, w^*) to CE removing (m, w) from CE if $w \neq w^*$ and (m, w) were engaged at the start of the round

End DO

FPDA example 1: round 1

Wmn	1st	2nd	3rd	4th
a	x	y	z	w
b	y	x	w	x
c	x	z	y	w
d	y	w	x	z

A * means: already proposed to that man

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	b	a	c	d
z	d	b	c	a

Round 1

Proposals:	New Engagements:
a: x	w: −
b: y	x: a
c: x	y: b
d: y	z: −

FPDA example 1: round 2

Wmn	1st	2nd	3rd	4th
a	x*	y	z	w
b	y*	x	w	x
c	x*	z	y	w
d	y*	w	x	z

A * means: already proposed to that man

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	b	a	c	d
z	d	b	c	a

Round 2

Current:	Proposals:	New Engagements:
w: –	a: –	w: d
x: a	b: –	x: a
y: b	c: z	y: b
z: –	d: w	z: c

Done - Stable:

a:x
b:y
c:z
d:w

FPDA example 2: round 1

Wmn	1st	2nd	3rd	4th
a	x	y	z	w
b	y	x	w	x
c	x	**y**	**z**	w
d	y	w	x	z

A * means: already proposed to that man

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	**c**	**b**	**a**	**d**
z	d	b	c	a

Round 1

Proposals:	New Engagements:
a: x	w: –
b: y	x: a
c: x	y: b
d: y	z: –

FPDA example 2: round 2

Wmn	1st	2nd	3rd	4th
a	x*	y	z	w
b	y*	x	w	x
c	x*	y	z	w
d	y*	w	x	z

A * means: already proposed to that man

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

Round 2

Current:	Proposals:	New Engagements:
w: –	a: –	w: d
x: a	b: –	x: a
y: b	c: y	y: ~~b~~ c
z: –	d: w	z: –

b is "jilted"

FPDA example 2: round 3

Wmn	1st	2nd	3rd	4th
a	x*	y	z	w
b	y*	x	w	x
c	x*	y*	z	w
d	y*	w*	x	z

A * means: already proposed to that man

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

Round 3

Current:	Proposals:	New Engagements:
w: d	a: –	w: d
x: a	b: x	x: ~~a~~ b
y: ~~b~~ c	c: –	y: ~~b~~ c
z: –	d: –	z: –

a is "jilted"

FPDA example 2: round 4

Wmn	1st	2nd	3rd	4th
a	x*	y	z	w
b	y*	x*	w	x
c	x*	y*	z	w
d	y*	w*	x	z

A * means: already proposed to that man

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

Round 4

Current:	Proposals:	New Engagements:
w: d	a: y	w: d
x: ~~a~~ b	b: –	x: ~~a~~ b
y: ~~b~~ c	c: –	y: ~~b~~ c
z: –	d: –	z: –

a's proposal not accepted by y
(no change)

FPDA example 2: round 5

Wmn	1st	2nd	3rd	4th
a	x*	y*	z	w
b	y*	x*	w	x
c	x*	y*	z	w
d	y*	w*	x	z

A * means: already proposed to that man

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

Round 5

Current:	Proposals:	New Engagements:
w: d	a: z	w: d
x: ~~a~~ b	b: –	x: ~~a~~ b
y: ~~b~~ c	c: –	y: ~~b~~ c
z: –	d: –	z: a

Stable:
a:z
b:x
c:y
d:w

Why does FPDA work?

We need to show that the FPDA always terminates and ends in a stable matching. Both termination and stability follow from a few observations:

- The sequence of women to which any man is engaged is strictly improving This is obvious since men are only accepting better and better offers
- The sequence of men proposed to by any woman gets worse in each new proposal This is again obvious since a woman who has to propose again must go to the next best person in her list.
- These observations not only guarantee termination but show that at most n^2 proposals are needed. Why? There are instances where the number of proposals come pretty close to this bound.
- The claim is that there cannot be a blocking pair (m, w) for the match μ produced by the FPDA algorithm. Why?

Why we cannot have a blocking pair

Suppose (m, w) is a blocking pair. Then w prefers m to $\mu(w)$. This means she would have proposed to m before proposing to $\mu(w)$. But then what would have happened if m prefers w to $\mu(m)$?

Why we cannot have a blocking pair

Suppose (m, w) is a blocking pair. Then w prefers m to $\mu(w)$. This means she would have proposed to m before proposing to $\mu(w)$.
But then what would have happened if m prefers w to $\mu(m)$?

Case 1: If w proposed before $\mu(m)$, then m would not have accepted $\mu(m)$'s proposal since he would be engaged to either w or someone even more preferred than w.

Case 2: If w proposed after $\mu(m)$, then he would have rejected/jilted $\mu(m)$.

This then shows that μ could not have been the matching so we conclude there is no blocking pair.

Properties of Deferred Acceptance

Theorem: FPDA (or MPDA) requires at most n^2 proposal

We are assuming that there are n men and n women. At each round, there is at least one new proposals So some women moves down in her preference list, proposing to a man one rank down. Since there are n women and n men the process must terminate within n^2 proposals.

With a little care and for a reasonable computation model, the overall time complexity is $O(n^2)$.

There are instance where the number of proposal comes “close” to this bound.

An example with many proposals

Man	1st	2nd	3rd	4th	5th
v	a	b	c	d	e
w	b	c	d	a	e
x	c	d	a	b	e
y	d	a	b	c	e
z	a	b	c	d	e

Wmn	1st	2nd	3rd	4th	5th
a	w	x	y	z	v
b	x	y	z	v	w
c	y	z	v	w	x
d	z	v	w	x	y
e	v	w	x	y	z

On this instance FPDA uses 21 proposals and 17 rounds.

FPDA is "female-optimal"

As stated the stable matching problem does not have a concept of social welfare. But for whom is a particular stable matching good or best?

Note that there can be many stable matchings. In general, the FPDA and MPDA will produce different matchings. When they are the same, there is a unique stable matching.

We can ask the following question. When restricted to stable matchings, who is the best match for a given women? It turns out that for the matching produced by the FPDA algorithm, every women is matched to her most desirable man

Optimality Theorem

Within the class of stable matchings, every women is matched to her most desirable man in the matching produced by the FPDA algorithm.

We this say that FPDA is *female-optimal* (even though it is the men who are in some sense making final decisions). This is a theorem and hence can be proved. Proof? Similarly, MPDA is *male-optimal*.

How good is the match for the men in the FPDA matching?

Just as we can define the best match for a women when restricted to stable matchings, we can define the worst match for a man when restricted to stable matchings.

Pessimality Theorem

Within the class of stable matchings, every man is matched to his least desirable women in the matching produced by the FPDA algorithm.

We say that FPDA is *male-pessimal*. Similarly, MPDA is female-pessimal.

Moral: Setting the agenda (i.e. in what order to propose) can be more influential than who makes the final decision. This phenomena exists in many settings (e.g. the rules of a tournament).

What is a consequence of this female optimality and male pessimality for FPDA?

Since every women is matched by the best man she can obtain in a stable matching, there is no reason for any women to every not be truthful when proposing.

But given that men are receiving their worst choice in the FPDA, it is perhaps not surprising that some men can benefit by not being truthful (i.e by rejecting a women even when she is better than his present choice).

In fact, unless there is a unique stable matching (i.e. when FPDA and MPDA result in the same matching), there is always at least one man who can benefit by not being truthful in the FPDA. And (anwering a question in class), it is computationally easy to find such a man and know what he should do (if he knows the results of the MPDA).

Note: There is no mechanism that is truthful for both men and women. Of course, when there is a unique mechanism (using FPDA or MPDA) then neither the men or women have reason to not be truthful.

An example showing how a man can lie to improve match

Here is an example where if everyone is acting truthfully, man $m1$ is matched with woman $w2$ in the FPDA matching.

$w1$: $m2 \succ_{w1} m1 \succ_{w1} m3 \succ_{w1} m4$	$m1$: $w1 \succ_{m1} w2 \succ_{m1} w3 \succ_{m1} w4$
$w2$: $m4 \succ_{w2} m1 \succ_{w2} m2 \succ_{w2} m3$	$m2$: $w2 \succ_{m2} w1 \succ_{m2} w3 \succ_{m2} w4$
$w3$: $m1 \succ_{w3} m3 \succ_{w3} m2 \succ_{w3} m4$	$m3$: $w3 \succ_{m3} w1 \succ_{m3} w2 \succ_{m3} w4$
$w4$: $m4 \succ_{w3} m3 \succ_{w3} m2 \succ_{w3} m1$	$m4$: $w4 \succ_{m4} w3 \succ_{m4} w2 \succ_{m4} w1$

Round 1		Round 2		Round 3	
Props:	Eng'd:	Props:	Eng'd:	Props:	Eng'd:
w1: m2	w1: m2	w1: –	w1: m2	w1: –	w1: m2
w2: m4	w2: –	w2: m1	w2: m1	w2: –	**w2: m1**
w3: m1	w3: m1	w3: –	w3: –	w3: m3	w3: m3
w4: m4	w4: m4	w4: –	w4: m4	w4: –	w4: m4

In the next slide, we show how some deceit by man $m1$ can improve his match.

How a particular man in this example can lie to improve his match

If instead of being truthful, man $m1$ lies in round 2 by rejecting $w2$s proposal (even though she is preferred to his current match $w3$), he eventually gets rewarded by ending up with his most preferred partner $w1$.

$w1: m2 >_{w1} m1 >_{w1} m3 >_{w1} m4$
$w2: m4 >_{w2} m1 >_{w2} m2 >_{w2} m3$
$w3: m1 >_{w3} m3 >_{w3} m2 >_{w3} m4$
$w4: m4 >_{w3} m3 >_{w3} m2 >_{w3} m1$

$m1: w1 >_{m1} w2 >_{m1} w3 >_{m1} w4$
$m2: w2 >_{m2} w1 >_{m2} w3 >_{m2} w4$
$m3: w3 >_{m3} w1 >_{m3} w2 >_{m3} w4$
$m4: w4 >_{m4} w3 >_{m4} w2 >_{m4} w1$

Round 1		Round 2		Round 3		Round 4	
Props:	Eng'd:	Props:	Eng'd:	Props:	Eng'd:	Props:	Eng'd:
w1: m2	w1: m2	w1: –	w1: m2	w1: –	w1: –	w1: m1	**w1: m1**
w2: m4	w2: –	w2: m1	w2: –	w2: m2	w2: m3	w2: –	w2: m3
w3: m1	w3: m1	w3: –	w3: m1	w3: –	w3: m3	w3: –	w3: m3
w4: m4	w4: m4	w4: –	w4: m4	w4: –	w4: m4	w4: –	w4: m4

Some practical considerations: Many to one matchings

As mentioned, many if not most applications of stable matching concern many to one matchings. For example, consider the application of residents applying to hospitals where now each hospital has a quota of some n_i positions. (This would likely be more targeted as to the nature of the residency; e.g. 3 positions in dermatology.)

As is the usual case, the residents propose (i.e. apply) to the hospitals (by giving their ranking of hospitals). A hospital program with a quota of n_i positions would then maintain in each round a set of up to the n_i best applicants (wrt. their ranking of residents) who have thus far been temporarily matched.

Notice that as before the hospitals can potentially benefit by not adhering to their ranking. Moreover, a hospital can now potentially also gain by misstating its quota.

Many other practical considerations

This is still an active field of research. Basically, the assumption that everyone knows the preferences of all potential matches is very unrealistic.

Here are some of the issues:

- As mentioned before, one important consideration is when one (or both) sides of the matching can have "complementarities ". This happens in the resident-hospital application when residents are couples. This is more generally the problem with the limitation of individuals having separate preference lists. Instead now, residents can now declare themselves as a couple and then list pairs of hospitals for their preferences.

  Unfortunately, there are now instances where stability cannot be obtained and it is NP hard to even determine if a stable match exists or to obtain one when one does exist or to minimize the "amount of instability". There are new and effective approaches to solving this problem in practice.

  There is always a Pareto optimal matching.

Many other practical considerations continued

- Participants may not know their full preference list and in practice do not need to know a full preference list. This leads to different approaches when having such incomplete knowledge.
  1. One can try to minimize, by elicitation, the amount of additional information needed to find a stable match.
  2. One can consider various ways to complete a partial ranking so as to achieve some worst case objective.
- Participants may only have probabilistic beliefs about their preferences and interviews are needed to be able to resolve rankings. But participants have limited budgets and must restrict themselves to a few interviews. How should participants choose interviews so as to maximize their expected utility (with respect to some utility measure).

I plan to post some additional slides by Joanna Drummond, a Toronto PHD student who is an expert on such stable matching issues.