Computational Social Choice: Spring 2015

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Matching

Voting and fair allocation of goods were both about agents having (different kinds of) preferences over possible alternatives.

Matching is about agents that have preferences over each other.

Many important applications:

- Matching junior doctors to hospitals
- Matching school children to schools
- Kidney exchanges (different model from what we'll discuss)

Matching research is a very early example for taking an algorithmic approach to a problem coming from economics.

Nobel Prize in Economics for Lloyd Shapley and Alvin Roth in 2012.

Gemeente Amsterdam

- Wonen
- Werken
- Vrije tijd

- Actueel
- Veelgevraagd
- Bestuur en organisatie
- Alle onderwerpen

- Contact
- English
- Adressengids
- Nieuwsbrief

Home > Onderwijs & jeugd > Voortgezet onderwijs > Centrale matching   Zoeken...

Centrale matching

9 december 2014

In het kort

Informatie over matching, de nieuwe methode van toewijzen van leerlingen aan VO-scholen.

Op deze pagina

- Centrale Matching
- Hoe werkt matching?
- Matching: tijdspad
- Matching: Deferred Acceptence

Centrale Matching

Dit jaar gaat het toewijzen van leerlingen aan VO-scholen anders dan in de jaren daarvoor. Amsterdam gaat gebruik maken van het systeem 'centrale matching'.

Hoe werkt matching?

1. Je krijgt voor elke school in het voorgezet onderwijs in Amsterdam een lotingnummer.
2. Je wordt tijdelijk geplaatst op de school die het hoogst op je voorkeurslijst staat.
3. Elke school kijkt of er genoeg plek is voor alle leerlingen die tijdelijk bij de school zijn geplaatst. Als dit het geval is, dan houdt de school al deze leerlingen vast. Als dat niet het geval is, dan houdt de school de leerlingen vast die voorrang hebben. De lotingnummers voor de school worden gebruikt om te bepalen welke leerlingen niet op de school geplaatst kunnen worden.

Plan for Today

This will be a (brief) introduction to the theory of *two-sided matching*:

- two groups of agents have preferences over possible matchings between them; and
- we need to find a "good" matching

Most of today's material is based on Roth and Sotomayor (1990).

A.E. Roth and M.A.O. Sotomayor. *Two-sided Matching: A Study in Game-theoretic modeling and analysis*. Cambridge University Press, 1990.

The Stable Marriage Problem

We are given:

- n *men* and n *women*
- each has a linear *preference* ordering over the opposite sex

We seek:

- a *stable* matching of men to women: no man and women should want to divorce their assigned partners and run off with each other

The Gale-Shapley Algorithm

Theorem 1 (Gale and Shapley, 1962) *There exists a stable matching for any combination of preferences of men and women.*

The *Gale-Shapley 'deferred acceptance' algorithm* for computing a stable matching works as follows:

- In each round, each man who is not yet engaged proposes to his favourite amongst the women he has not yet proposed to.
- In each round, each woman picks her favourite from the proposals she's receiving and the man she's currently engaged to (if any).
- Stop when everyone is engaged.

D. Gale and L.S. Shapley. College Admissions and the Stability of Marriage. *American Mathematical Monthly*, 69:9–15, 1962.

Analysis

The Gale-Shapley algorithm is correct and efficient:

- The algorithm always *terminates*.
- The algorithm always returns a *stable* matching. For if not, the unhappy man would have proposed to the unhappy woman ...
- The algorithm has *quadratic complexity:* even in the worst case, no man will propose twice to the same woman. For instance:
  - each man has a different favourite $\rightsquigarrow$ 1 round (n proposals)
  - all men have the same preferences $\rightsquigarrow$ $\frac{n(n+1)}{2}$ proposals

M-Optimal and W-Optimal Matchings

A stable matching is called *M-optimal* if every man likes it at least as much as every other stable matching.

A stable matching is called *W-optimal* if every woman likes it at least as much as every other stable matching.

It is possible to prove that the matching returned by the Gale-Shapley algorithm (with men proposing) is M-optimal (and W-pessimal).

Fairness

M-optimal matchings (returned by the Gale-Shapley algorithm) arguably are not fair. But what is *fair*?

- One option is to implement the stable matching that *minimises* the *regret* of the person worst off (regret = number of members of the opposite sex they prefer to their assigned partner).

  Gusfield (1987) gives an algorithm for min-regret stable matchings.

- Similarly, we can implement the stable matching that maximises *average satisfaction* (i.e., that minimises average regret).

  Irving et al. (1987) give an algorithm for this problem.

D. Gusfield. Three Fast Algorithms for Four Problems in Stable Marriage. *SIAM Journal of Computing*, 16(1):111–128, 1987.

R.W. Irving, P. Leather, and D. Gusfield. An Efficient Algorithm for the “Optimal” Stable Marriage. *Journal of the ACM*, 34(3):532–543, 1987.

Stable Marriages under Incomplete Preferences

In an important generalisation of the simple stable marriage problem, people are allowed to specify which members of the opposite sex they consider *acceptable*, and they only report a strict ranking of those.

- Now the assumption is that a man/woman would rather remain *single* than marry a partner they consider unacceptable.
- Now a matching is *stable* if no couple has an incentive to run off together and if no individual has an incentive to leave their assigned partner and be single.
- The *Gale-Shapley algorithm* can easily be extended to this setting: simply stipulate that men don't propose to unacceptable women and women don't accept unacceptable men.

This is called the stable marriage problem with *incomplete preferences*.

Impossibility of Strategy-Proof Stable Matching

Call a matching mechanism *strategy-proof* if it never gives either a man or a woman an incentive to misrepresent their preferences.

Theorem 2 (Roth, 1982) *There exists no matching mechanism that is both stable and strategy-proof.*

The proof on the next slide uses only two men and two women, but it relies on a manipulation involving agents misrepresenting which partners they find *acceptable*. Alternative proofs, using three men and three women, involve only changes in preference (not acceptability).

A.E. Roth. The Economics of Matching: Stability and Incentives. *Mathematics of Operations Research*, 7:617–628, 1982.

Proof

Suppose there are two men and two women with these preferences:

$$m_1 : w_1 \succ w_2 \mid \qquad m_2 : w_2 \succ w_1 \mid$$
$$w_1 : m_2 \succ m_1 \mid \qquad w_2 : m_1 \succ m_2 \mid$$

$\leadsto$ 2 stable matchings: $\{(m_1, w_1), (m_2, w_2)\}$ and $\{(m_1, w_2), (m_2, w_1)\}$
So any stable mechanism will have to pick one of them.

- Suppose the mechanism would pick $\{(m_1, w_1), (m_2, w_2)\}$. Then w_2 has an incentive to pretend that she finds m_2 unacceptable, as then $\{(m_1, w_2), (m_2, w_1)\}$ becomes the only stable matching.
- Suppose the mechanism would pick $\{(m_1, w_2), (m_2, w_1)\}$. Then m_1 has an incentive to pretend that he finds w_2 unacceptable, as then $\{(m_1, w_1), (m_2, w_2)\}$ becomes the only stable matching.

Hence, for any possible stable matching mechanism there is a situation where someone has an incentive to manipulate. ✓

Preferences with Ties

We can further generalise the stable marriage problem by also allowing for *ties*, i.e., by allowing each agent to have a weak preference order over (acceptable) members of the opposite sex.

Two observations:

- We can still *compute a stable matching* in polynomial time:
  (1) arbitrarily break the ties
  (2) apply the standard Gale-Shapley algorithm
- Now (for the first time in this lecture) different stable matchings of the same problem may have *different size*. Example:

$$m_1 : w_1 \mid w_2 \qquad m_2 : w_1 \succ w_2$$
$$w_1 : m_1 \sim m_2 \qquad w_2 : m_2 \mid m_1$$

Both $\{(m_2, w_1)\}$ and $\{(m_1, w_1), (m_2, w_2)\}$ are stable.

Complexity of Computing Maximal Stable Matchings

Recall that computing *some* stable matching is still polynomial. But as there may be exponentially many of them, this doesn't mean that we can compute a most preferred stable matching efficiently. Indeed:

Theorem 3 (Manlove et al., 2002) *Deciding whether a stable matching with a cardinality exceeding K exists is NP-complete for marriage problems with incomplete preferences and ties.*

Proof: Omitted.

Note that the above is the decision variant of the problem of computing a matching of *maximal cardinality*.

D.F. Manlove, R.W. Irving, K. Iwama, S. Miyazaki, and Y. Morita. Hard Variants of Stable Marriage. *Theoretical Computer Science*, 276(1–2):261–279, 2002.

Variants

Variants and generalisations are applicable to many scenarios:

- *Residents-Hospitals Problem:*

  Matching of junior doctors (residents) to hospitals.

  Many-to-one variant of stable mariage problem with incomplete preferences, with each hospital having a certain capacity.

- *School Choice:*

  Matching of school children to schools.

  Like residents-hospitals problem, but schools have 'priorities' rather than preferences (e.g., based on distance to home of child or whether sibling already at school).

  Main difference is interpretational: schools not economic agents.

Summary

We have seen several variants of two-sided matching problems:

- basic marriage problem; extension to incomplete preferences; extension to preferences with ties
- we have hinted at possible extensions (many-to-one, ...)

We have discussed various desirable properties:

- stability: no agent(s) have an incentive break the matching
- fairness: possibly expressed in terms of "regret"
- strategy-proofness: incompatible with stability
- possibly conditions on cardinality: can lead to intractability
- algorithmic efficiency

We have seen how the 'deferred acceptance' algorithm of Gale and Shapley can be used to compute stable matchings efficiently.