

Computational Social Choice: Autumn 2010

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

Plan for Today

Today's lecture will be a (brief) introduction to the problem 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 the material on these slides is based on the book by 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 woman should have an incentive 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 \leadsto 1 round (n proposals)
 - all men have the same preferences $\leadsto \frac{n(n+1)}{2}$ 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.

Observation: 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) are not fair. But what is *fair*?

- One option is to implement the stable matching that *minimises* the *regret* of the person worst off. (The regret of a person is the 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.

Arguably, neither of these definitions of fairness is fully convincing (but it is also not clear what would be a better solution).

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.

In the literature, this is known as the stable marriage problem with *incomplete preferences*. Note that incompleteness here does *not* mean that preferences are taken to be partial orders.

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 mates 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:

$$\begin{array}{l} m_1 : w_1 \succ w_2 \mid \\ w_1 : m_2 \succ m_1 \mid \end{array} \quad \begin{array}{l} m_2 : w_2 \succ w_1 \mid \\ w_2 : m_1 \succ m_2 \mid \end{array}$$

\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. \checkmark

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:

$$\begin{array}{ll} m_1 : w_1 \mid w_2 & m_2 : w_1 \succ w_2 \\ w_1 : m_1 \sim m_2 & w_2 : m_2 \mid m_1 \end{array}$$

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.

Applications

Matching theory has a number of important applications, e.g.:

- Matching students to universities (but note that each university will usually accept more than one student).
- Matching junior doctors to hospitals. Additional complications arise when we want to allow married couples (two junior doctors) to express joint preferences.
- Kidney transplants: matching patient-donor pairs. Acceptability relates to blood type etc. Preferences can be used to optimise expected compatibility. Longer cycles, going beyond two-sided matching are also possible (but too long is risky).

Depending on the application, we might be interested in fairness issues, stability, strategy-proofness, algorithmic efficiency (or any combinations of the above).

COMSOC Concerns

Research on two-sided matching has always combined algorithmic issues with social choice-theoretic (or at least game-theoretic) concerns. There has been work on the computational complexity of particular matching problems since the 1980s. Some examples of recent work:

- Parametrised complexity of computing maximal stable matchings with ties (Marx and Schlotter, 2010)
- Hardness of manipulation (Pini et al., 2010)
- Compact preference representation: using CP-nets in two-sided matching (Pilotto et al., 2009)

D. Marx and I. Schlotter. Parameterized Complexity and Local Search Approaches for the Stable Marriage Problem with Ties. *Algorithmica*, 58:170–187, 2010.

M.S. Pini, F. Rossi, K.B. Venable, and T. Walsh. Manipulation Complexity and Gender Neutrality in Stable Marriage Procedures. *JAAMAS*. In press (2010).

E. Pilotto, F. Rossi, K.B. Venable, and T. Walsh. Compact Preference Representation in Stable Marriage Problems. Proc. ADT-2009.

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 (more than two sides, additional constraints, ...)

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.