An Efficient Algorithm for the "Optimal" Stable Marriage Problem

Dimitris Chatzidimitriou

Corelab
NTUA

June 28, 2013

Outline

Outline

What is a stable matching (or marriage)?

- An instance of size n of the stable matching problem consists of n men's and n women's lists of all the members of the opposite sex ranked in order of preference.
- A complete matching is a set of n pairs (or couples) (m, w), in which each man m and each woman w appears in only one pair.
- A stable matching is a complete matching between men and women such that no man and woman who are not a couple, prefer each other more than their partners in the matching. More formally: $\nexists$ couples $(m_1, w_1), (m_2, w_2)$ such that m_1 prefers w_2 to w_1 and w_2 prefers m_1 to m_2.

Notice that the preference lists must be **complete**!

Basic Concepts

What is a stable matching (or marriage)?

- An instance of size n of the stable matching problem consists of n men's and n women's lists of all the members of the opposite sex ranked in order of preference.
- A complete matching is a set of n pairs (or couples) (m, w), in which each man m and each woman w appears in only one pair.
- A stable matching is a complete matching between men and women such that no man and woman who are not a couple, prefer each other more than their partners in the matching. More formally: $\nexists$ couples $(m_1, w_1), (m_2, w_2)$ such that m_1 prefers w_2 to w_1 and w_2 prefers m_1 to m_2.

Notice that the preference lists must be **complete**!

Basic Concepts

What is a stable matching (or marriage)?

- An instance of size n of the stable matching problem consists of n men's and n women's lists of all the members of the opposite sex ranked in order of preference.
- A complete matching is a set of n pairs (or couples) (m, w), in which each man m and each woman w appears in only one pair.
- A stable matching is a complete matching between men and women such that no man and woman who are not a couple, prefer each other more than their partners in the matching. More formally: $\nexists$ couples $(m_1, w_1), (m_2, w_2)$ such that m_1 prefers w_2 to w_1 and w_2 prefers m_1 to m_2.

Notice that the preference lists must be **complete**!

The Gale-Shapley Algorithm

In 1962 Gale and Shapley proved that for each instance of the problem there exists at least one stable matching. In fact they gave a greedy algorithm to compute one!

The algorithm is very simple: Each man "proposes", in order, to all the women, pausing when a woman agrees to consider his proposal, but continuing if a proposal gets rejected. When a woman receives a proposal, she rejects it if she already has a better one, otherwise she agrees to consider it rejectin the one she had (if any).

This results in a male optimal solution! A female optimal solution can also be obtained by reversing the roles of men and women in the above algorithm. Obviously the two solutions match if and only if that is the only stable matching.

The Gale-Shapley Algorithm

In 1962 Gale and Shapley proved that for each instance of the problem there exists at least one stable matching. In fact they gave a greedy algorithm to compute one!

The algorithm is very simple: Each man "proposes", in order, to all the women, pausing when a woman agrees to consider his proposal, but continuing if a proposal gets rejected. When a woman receives a proposal, she rejects it if she already has a better one, otherwise she agrees to consider it rejectin the one she had (if any).

This results in a male optimal solution! A female optimal solution can also be obtained by reversing the roles of men and women in the above algorithm. Obviously the two solutions match if and only if that is the only stable matching.

The Gale-Shapley Algorithm

In 1962 Gale and Shapley proved that for each instance of the problem there exists at least one stable matching. In fact they gave a greedy algorithm to compute one!

The algorithm is very simple: Each man “proposes”, in order, to all the women, pausing when a woman agrees to consider his proposal, but continuing if a proposal gets rejected. When a woman receives a proposal, she rejects it if she already has a better one, otherwise she agrees to consider it rejectin the one she had (if any).

This results in a male optimal solution! A female optimal solution can also be obtained by reversing the roles of men and women in the above algorithm. Obviously the two solutions match if and only if that is the only stable matching.

But what about the "socially optimal" stable matching?

We could of course compute all the stable matchings and then compare them to find the best.

...but the maximum number of stable matchings in an instance of size n grows exponentially with n!

We will present an $O(n^4)$-time algorithm by Irving, Leather and Gusfield to find the socially optimal stable matching.

But what about the "socially optimal" stable matching?

We could of course compute all the stable matchings and then compare them to find the best.

...but the maximum number of stable matchings in an instance of size n grows exponentially with n!

We will present an $O(n^4)$-time algorithm by Irving, Leather and Gusfield to find the socially optimal stable matching.

But what about the "socially optimal" stable matching?

We could of course compute all the stable matchings and then compare them to find the best.

...but the maximum number of stable matchings in an instance of size n grows exponentially with n!

We will present an $O(n^4)$-time algorithm by Irving, Leather and Gusfield to find the socially optimal stable matching.

But what about the “socially optimal” stable matching?

We could of course compute all the stable matchings and then compare them to find the best.

...but the maximum number of stable matchings in an instance of size n grows exponentially with n!

We will present an $O(n^4)$-time algorithm by Irving, Leather and Gusfield to find the socially optimal stable matching.

Outline

Notation

- The preference lists can be specified by two *nxn* ranking matrices:

$$mr(i,k) = j \quad \text{if woman } k \text{ is the } j\text{-th choice of man } i$$

$$wr(i,k) = j \quad \text{if man } k \text{ is the } j\text{-th choice of woman } i$$

- Suppose that, for a given stable matching instance,

$$S = \{(m_1, w_1), \ldots, (m_n, w_n)\}$$

  is a stable matching.
- We define the value $c(S)$ of matching S as:

$$c(S) = \sum_{1}^{n} mr(m_i, w_i) + \sum_{1}^{n} wr(m_i, w_i)$$

- A stable matching S is optimal if it has **minimum** possible value $c(S)$.

Shortlists

The sequence of proposals in the Gale-Shapley Algorithm has two very important implications:

- If m proposes to w, then there is no stable matching in which m can "do better" than w.
- If w receives a proposal from m, then there is no stable matching in which w can "do worse" than m.

These observations suggest that we should remove m and w from each other's list, if w receives a proposal from someone she likes better than m.

By doing so consecutively we result in the *(male oriented) shortlists.*

Shortlists

The sequence of proposals in the Gale-Shapley Algorithm has two very important implications:

- If m proposes to w, then there is no stable matching in which m can "do better" than w.
- If w receives a proposal from m, then there is no stable matching in which w can "do worse" than m.

These observations suggest that we should remove m and w from each other's list, if w receives a proposal from someone she likes better than m.

By doing so consecutively we result in the *(male oriented) shortlists.*

Shortlists

The sequence of proposals in the Gale-Shapley Algorithm has two very important implications:

- If m proposes to w, then there is no stable matching in which m can "do better" than w.
- If w receives a proposal from m, then there is no stable matching in which w can "do worse" than m.

These observations suggest that we should remove m and w from each other's list, if w receives a proposal from someone she likes better than m.

By doing so consecutively we result in the *(male oriented) shortlists.*

Example

An example of an instance of size 8:

1: 3 1 5 7 4 2 8 6	1: 4 3 8 1 2 5 7 6
2: 6 1 3 4 8 7 5 2	2: 3 7 5 8 6 4 1 2
3: 7 4 3 6 5 1 2 8	3: 7 5 8 3 6 2 1 4
4: 5 3 8 2 6 1 4 7	4: 6 4 2 7 3 1 5 8
5: 4 1 2 8 7 3 6 5	5: 8 7 1 5 6 4 3 2
6: 6 2 5 7 8 4 3 1	6: 5 4 7 6 2 8 3 1
7: 7 8 1 6 2 3 4 5	7: 1 4 5 6 2 8 3 7
8: 2 6 7 1 8 3 4 5	8: 2 5 4 3 7 8 1 6
male preference lists	female preference lists

And the resulting shortlists:

1: 3 1 5 7 4	1: 4 3 8 1 2
2: 1 3 4 8 7	2: 3 7 5 8
3: 7 4 3 1 2 8	3: 7 5 8 3 6 2 1
4: 5 8 6 1 4 7	4: 6 4 2 7 3 1 5
5: 4 2 8 7 3 6 5	5: 8 7 1 5 6 4
6: 6 5 7 4 3	6: 5 4 7 6
7: 8 6 2 3 4 5	7: 1 4 5 6 2 8 3
8: 2 7 1 3 5	8: 2 5 4 3 7
male shortlists	female shortlists

Outline

Rotations

- A **rotation** ρ of size r is a sequence

$$\rho = (m_0, w_0), \dots, (m_{r-1}, w_{r-1})$$

  of man/woman pairs such that, for each $i \in [0, r-1]$,
  1. w_i is first in m_i's shortlist
  2. w_{i+1} is second in m_i's shortlist ($i+1$ is taken modulo r)

  Such a rotation is said to be **exposed** in the shortlists.

- Observation: Unless we have reached the female-optimal solution, at least one rotation is exposed.

- Significance: Starting from a stable solution, if each m_i in an exposed rotation exchanges his partner w_i for w_{i+1}, the resulting solution is also stable!

- If, given a rotation $\rho = (m_0, w_0), \dots, (m_{r-1}, w_{r-1})$, we remove each succesor x of m_{i-1} in w_i's shortlist and also w_i from x's list for each i, then the rotation ρ is said to have been **eliminated**.

Rotations

- A **rotation** ρ of size r is a sequence

$$\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$$

  of man/woman pairs such that, for each $i \in [0, r-1]$,
  1. w_i is first in m_i's shortlist
  2. w_{i+1} is second in m_i's shortlist ($i+1$ is taken modulo r)

  Such a rotation is said to be **exposed** in the shortlists.

- Observation: Unless we have reached the female-optimal solution, at least one rotation is exposed.

- Significance: Starting from a stable solution, if each m_i in an exposed rotation exchanges his partner w_i for w_{i+1}, the resulting solution is also stable!

- If, given a rotation $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$, we remove each succesor x of m_{i-1} in w_i's shortlist and also w_i from x's list for each i, then the rotation ρ is said to have been **eliminated**.

Rotations

- A **rotation** ρ of size r is a sequence

$$\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$$

  of man/woman pairs such that, for each $i \in [0, r-1]$,
  1. w_i is first in m_i's shortlist
  2. w_{i+1} is second in m_i's shortlist ($i+1$ is taken modulo r)

  Such a rotation is said to be **exposed** in the shortlists.

- Observation: Unless we have reached the female-optimal solution, at least one rotation is exposed.

- Significance: Starting from a stable solution, if each m_i in an exposed rotation exchanges his partner w_i for w_{i+1}, the resulting solution is also stable!

- If, given a rotation $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$, we remove each succesor x of m_{i-1} in w_i's shortlist and also w_i from x's list for each i, then the rotation ρ is said to have been **eliminated**.

Rotations

- A **rotation** ρ of size r is a sequence

$$\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$$

  of man/woman pairs such that, for each $i \in [0, r-1]$,
  1. w_i is first in m_i's shortlist
  2. w_{i+1} is second in m_i's shortlist ($i+1$ is taken modulo r)

  Such a rotation is said to be **exposed** in the shortlists.

- Observation: Unless we have reached the female-optimal solution, at least one rotation is exposed.

- Significance: Starting from a stable solution, if each m_i in an exposed rotation exchanges his partner w_i for w_{i+1}, the resulting solution is also stable!

- If, given a rotation $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$, we remove each succesor x of m_{i-1} in w_i's shortlist and also w_i from x's list for each i, then the rotation ρ is said to have been **eliminated**.

Rotations cont'd

- In fact, **every stable matching** M **for a given stable marriage instance can be obtained by eliminating a fixed number of rotations**. Note that even though this set of rotations R_M is fixed, the order of their elimination is not!

- Also, since each pair (m, w) can appear in **at most** one rotation and any rotation contains at least two pairs, **the total number of rotations is** $O(n^2)$.

Rotations cont'd

- In fact, **every stable matching M for a given stable marriage instance can be obtained by eliminating a fixed number of rotations**. Note that even though this set of rotations R_M is fixed, the order of their elimination is not!

- Also, since each pair (m, w) can appear in **at most** one rotation and any rotation contains at least two pairs, **the total number of rotations is $O(n^2)$**.

Rotation Poset

- A rotation π is said to be an **explicit predecessor** of $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$ if, for some $i \in [0, r-1]$ and some woman y ($\neq w_i$), π eliminates (m_i, y) and m_i prefers y to w_i.
- Obviously a rotation cannot become exposed until all of its explicit predecessors are eliminated.
- The reflexive transitive closure $\leq$ of the explicit predecessor relation is a **partial order** on the set of rotations, called the **rotation poset**, and $\pi < \rho$ if and only if π must be eliminated before ρ becomes exposed.
- A **closed set** in a poset $(P, \leq)$ is a subset C of P such that

$$\rho \in C \;\&\; \pi < \rho \Rightarrow \pi \in C.$$

Rotation Poset

- A rotation π is said to be an **explicit predecessor** of $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$ if, for some $i \in [0, r-1]$ and some woman y ($\neq w_i$), π eliminates (m_i, y) and m_i prefers y to w_i.

- Obviously a rotation cannot become exposed until all of its explicit predecessors are eliminated.

- The reflexive transitive closure $\leq$ of the explicit predecessor relation is a **partial order** on the set of rotations, called the **rotation poset**, and $\pi < \rho$ if and only if π must be eliminated before ρ becomes exposed.

- A **closed set** in a poset $(P, \leq)$ is a subset C of P such that

$$\rho \in C \;\&\; \pi < \rho \Rightarrow \pi \in C.$$

Rotation Poset

- A rotation π is said to be an **explicit predecessor** of $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$ if, for some $i \in [0, r-1]$ and some woman y ($\neq w_i$), π eliminates (m_i, y) and m_i prefers y to w_i.
- Obviously a rotation cannot become exposed until all of its explicit predecessors are eliminated.
- The reflexive transitive closure $\leq$ of the explicit predecessor relation is a **partial order** on the set of rotations, called the **rotation poset**, and $\pi < \rho$ if and only if π must be eliminated before ρ becomes exposed.
- A **closed set** in a poset $(P, \leq)$ is a subset C of P such that

$$\rho \in C \;\&\; \pi < \rho \Rightarrow \pi \in C.$$

Rotation Poset

- A rotation π is said to be an **explicit predecessor** of $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$ if, for some $i \in [0, r-1]$ and some woman y ($\neq w_i$), π eliminates (m_i, y) and m_i prefers y to w_i.

- Obviously a rotation cannot become exposed until all of its explicit predecessors are eliminated.

- The reflexive transitive closure $\leq$ of the explicit predecessor relation is a **partial order** on the set of rotations, called the **rotation poset**, and $\pi < \rho$ if and only if π must be eliminated before ρ becomes exposed.

- A **closed set** in a poset $(P, \leq)$ is a subset C of P such that

$$\rho \in C \;\&\; \pi < \rho \Rightarrow \pi \in C.$$

Rotation Poset cont'd

From the previous example:

Rotation	*Immediate Predecessors*
$\rho_1 = (1, 3), (2, 1)$	—
$\rho_2 = (3, 7), (5, 4), (8, 2)$	—
$\rho_3 = (4, 5), (7, 8), (6, 6)$	—
$\rho_4 = (1, 1), (6, 5), (8, 7)$	ρ_1, ρ_2, ρ_3
$\rho_5 = (2, 3), (3, 4)$	ρ_1, ρ_2
$\rho_6 = (4, 8), (7, 6), (5, 2)$	ρ_2, ρ_3
$\rho_7 = (3, 3), (8, 1)$	ρ_4, ρ_5
$\rho_8 = (2, 4), (5, 8), (6, 7)$	ρ_4, ρ_5, ρ_6
$\rho_9 = (1, 5), (5, 7), (8, 3)$	ρ_7, ρ_8
$\rho_{10} = (3, 1), (7, 2), (5, 3), (4, 6)$	ρ_9

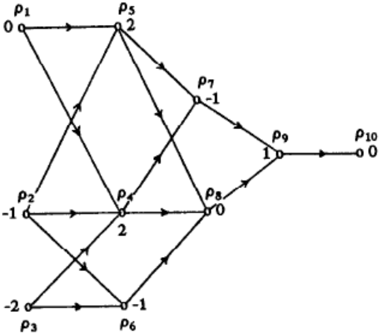

It should now be obvious that:

Theorem

The stable matchings of a given stable marriage instance are in one-to-one correspondence with the closed subsets of the rotation poset.

Rotation Poset cont'd

From the previous example:

Rotation	Immediate Predecessors
$\rho_1 = (1,3), (2,1)$	—
$\rho_2 = (3,7), (5,4), (8,2)$	—
$\rho_3 = (4,5), (7,8), (6,6)$	—
$\rho_4 = (1,1), (6,5), (8,7)$	ρ_1, ρ_2, ρ_3
$\rho_5 = (2,3), (3,4)$	ρ_1, ρ_2
$\rho_6 = (4,8), (7,6), (5,2)$	ρ_2, ρ_3
$\rho_7 = (3,3), (8,1)$	ρ_4, ρ_5
$\rho_8 = (2,4), (5,8), (6,7)$	ρ_4, ρ_5, ρ_6
$\rho_9 = (1,5), (5,7), (8,3)$	ρ_7, ρ_8
$\rho_{10} = (3,1), (7,2), (5,3), (4,6)$	ρ_9

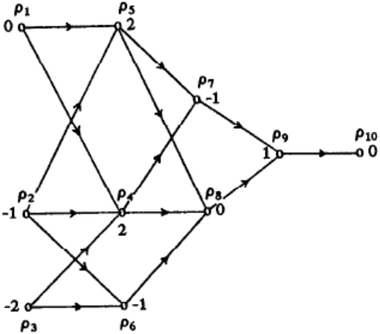

It should now be obvious that:

Theorem

The stable matchings of a given stable marriage instance are in one-to-one correspondence with the closed subsets of the rotation poset.

Weights

Given a rotation $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$, we define its **weight** $w(\rho)$ by

$$w(\rho) = \sum_{0}^{r-1} [mr(m_i, w_i) - mr(m_i, w_{i+1})] + \sum_{0}^{r-1} [wr(w_i, m_i) - wr(w_i, m_{i-1})]$$

($i-1, i+1$ are taken mod r)

Lemma

Let S be a stable matching obtained by starting from the shortlists and eliminating a particular sequence of rotations. Suppose that ρ is a rotation exposed in S, and let S' be the stable matching obtained from S by eliminating ρ. Then,

$$c(S') = c(S) - w(\rho).$$

Which leads to...

Weights

Given a rotation $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$, we define its **weight** $w(\rho)$ by

$$w(\rho) = \sum_{0}^{r-1} [mr(m_i, w_i) - mr(m_i, w_{i+1})] + \sum_{0}^{r-1} [wr(w_i, m_i) - wr(w_i, m_{i-1})]$$

$(i-1, i+1$ are taken mod $r)$

Lemma

Let S be a stable matching obtained by starting from the shortlists and eliminating a particular sequence of rotations. Suppose that ρ is a rotation exposed in S, and let S' be the stable matching obtained from S by eliminating ρ. Then,

$$c(S') = c(S) - w(\rho).$$

Which leads to...

Weights

Given a rotation $\rho = (m_0, w_0), \ldots, (m_{r-1}, w_{r-1})$, we define its **weight** $w(\rho)$ by

$$w(\rho) = \sum_{0}^{r-1} [mr(m_i, w_i) - mr(m_i, w_{i+1})] + \sum_{0}^{r-1} [wr(w_i, m_i) - wr(w_i, m_{i-1})]$$

($i-1, i+1$ are taken mod r)

Lemma

Let S be a stable matching obtained by starting from the shortlists and eliminating a particular sequence of rotations. Suppose that ρ is a rotation exposed in S, and let S' be the stable matching obtained from S by eliminating ρ. Then,

$$c(S') = c(S) - w(\rho).$$

Which leads to...

Weighted Posets

Lemma

If S is the stable matching obtained from the shortlists by eliminating the rotations $\rho_1, \ldots, \rho_t$ and S_0 the male-optimal solution, then

$$c(S) = c(S_0) - \sum_{1}^{t} w(\rho_i).$$

Note that $\rho_1, \ldots, \rho_t$ must form a closed subset of the rotation poset.

To summarize, after we assign to each rotation ρ its integer weight $w(\rho)$, all that is left is to find the maximum-weight closed subset of the weighted poset.

Weighted Posets

Lemma

If S is the stable matching obtained from the shortlists by eliminating the rotations $\rho_1, \ldots, \rho_t$ and S_0 the male-optimal solution, then

$$c(S) = c(S_0) - \sum_{1}^{t} w(\rho_i).$$

Note that $\rho_1, \ldots, \rho_t$ must form a closed subset of the rotation poset.

To summarize, after we assign to each rotation ρ its integer weight $w(\rho)$, all that is left is to find the maximum-weight closed subset of the weighted poset.

Lemma

If S is the stable matching obtained from the shortlists by eliminating the rotations $\rho_1, \ldots, \rho_t$ and S_0 the male-optimal solution, then

$$c(S) = c(S_0) - \sum_{1}^{t} w(\rho_i).$$

Note that $\rho_1, \ldots, \rho_t$ must form a closed subset of the rotation poset.

To summarize, after we assign to each rotation ρ its integer weight $w(\rho)$, all that is left is to find the maximum-weight closed subset of the weighted poset.

Outline

An Alternative Rotation Graph

The construction of the weighted rotation poset P would require $O(n^6)$ time, so instead we will construct a sparse directed graph P' representing in some way P, through which we will obtain the minimum-weight closed subgraph of P.

We denote P' the subgraph of P consisting of all the nodes of P (one for each rotation), but only of those edges defined by these two rules:

1. If (m, w) is a member of some rotation π and w' is *the first woman below w in m's list such that (m, w') is a member of some other rotation ρ*, then P' contains a $(\pi \rightarrow \rho)$ edge.
2. If (m, w) is not a member of any rotation, but is eliminated by some rotation π and w is *the first woman above w' in m's list such that (m, w) is a member of some rotation ρ*, then P' contains a $(\pi \rightarrow \rho)$ edge.

An Alternative Rotation Graph

The construction of the weighted rotation poset P would require $O(n^6)$ time, so instead we will construct a sparse directed graph P' representing in some way P, through which we will obtain the minimum-weight closed subgraph of P.

We denote P' the subgraph of P consisting of all the nodes of P (one for each rotation), but only of those edges defined by these two rules:

1. If (m, w) is a member of some rotation π and w' is *the first woman below w in m's list such that (m, w') is a member of some other rotation* ρ, then P' contains a $(\pi \to \rho)$ edge.
2. If (m, w) is not a member of any rotation, but is eliminated by some rotation π and w is *the first woman above w' in m's list such that (m, w) is a member of some rotation* ρ, then P' contains a $(\pi \to \rho)$ edge.

An Alternative Rotation Graph cont'd

Lemma (1)

P' has at most $O(n^2)$ edges and, given the rotations, can be constructed in $O(n^2)$-time.

Proof.

In the rules above, each creation of an edge is associated with a pair (m, w). But no pair is associated more than once, hence P' has at most $O(n^2)$ edges.
To construct P', assuming the rotations are known, we first scan each rotation, noting the pairs that are in it and labeling the pairs it eliminates by the name of the rotation. This needs $O(n^2)$-time. Then we scan each man's preference list, keeping track of the most recently encountered pair contained in some rotation, and applying the above two rules. Each scan takes $O(n)$-time, hence P' will be constructed in $O(n^2)$ time. □

An Alternative Rotation Graph cont'd

Lemma (2)

The transitive closure of P' is P, hence P' preserves the closed subsets of P.

Proof.

Obviously $P' \subseteq P$, so we only need to show that a closed subset in P is also closed in P'. For this it suffices to show that if a rotation π explicitly precedes ρ, then ther is a directed (π, ρ)-path in P'. Let (m_i, x) be a pair eliminated by π such that $(m_i, w_i) \in \rho$ and m_i prefers x to w_{i+1}. Now, if there is a woman w above x in m_i's list such that (m_i, w) is a pair in some rotation, then let w be the first such woman and let $(m_i, w) \in \sigma$. Then successive applications of rule 1 give a directed (σ, ρ)-path of ≥ 0 edges in P' and rule 2 gives a $(\pi \to \sigma)$ edge in P'. If there is no such woman w, then m_i must prefer x to his mate in every other stable matching, but π must be exposed in at least one such stable matching M and x must prefer m_i to her mate in M, so M cannot be stable. □

Finding All the Rotations

We still have to find efficiently the set of all rotations. The following method takes $O(n^3)$-time, but it has been later refined to just $O(n^2)$.

Let S be a stable matching and assume that the preference lists have been reduced for S. Then for each man m we denote by $s(S, m)$ the second woman in m's list (if any) and by $s'(S, m)$ her mate in S. Furthermore, we denote by $G(S)$ the directed graph consisting of n nodes (one for each man), where for every man m_i there is an edge $(m_i \to s'(S, m_i))$.

It is clear that there is a bijection between the rotations exposed in S and the directed cycles in $G(S)$.

We still have to find efficiently the set of all rotations. The following method takes $O(n^3)$-time, but it has been later refined to just $O(n^2)$.

Let S be a stable matching and assume that the preference lists have been reduced for S. Then for each man m we denote by $s(S, m)$ the second woman in m's list (if any) and by $s'(S, m)$ her mate in S. Furthermore, we denote by $G(S)$ the directed graph consisting of n nodes (one for each man), where for every man m_i there is an edge $(m_i \to s'(S, m_i))$.

It is clear that there is a bijection between the rotations exposed in S and the directed cycles in $G(S)$.

We still have to find efficiently the set of all rotations. The following method takes $O(n^3)$-time, but it has been later refined to just $O(n^2)$.

Let S be a stable matching and assume that the preference lists have been reduced for S. Then for each man m we denote by $s(S, m)$ the second woman in m's list (if any) and by $s'(S, m)$ her mate in S. Furthermore, we denote by $G(S)$ the directed graph consisting of n nodes (one for each man), where for every man m_i there is an edge $(m_i \to s'(S, m_i))$.

It is clear that there is a bijection between the rotations exposed in S and the directed cycles in $G(S)$.

Finding All the Rotations cont'd

Lemma

All cycles in $G(S)$ can be identified in $O(n)$ time.

Proof.

Each node has outdegree ≤ 1, therefore $G(S)$ has $\leq n$ edges, which means that each edge is in at most one cycle. We can find all the cycles by using DFS in $O(n)$ time. □

It follows that all the rotations exposed in S can be also found in $O(n)$ time.

Finding All the Rotations cont'd

Lemma

All cycles in $G(S)$ can be identified in $O(n)$ time.

Proof.

Each node has outdegree ≤ 1, therefore $G(S)$ has $\leq n$ edges, which means that each edge is in at most one cycle. We can find all the cycles by using DFS in $O(n)$ time. □

It follows that all the rotations exposed in S can be also found in $O(n)$ time.

Finding All the Rotations cont'd

Theorem

All the rotations can be found in $O(n^3)$-time.

Proof.

We define a node ρ in P' to be at *level 0* iff ρ has no predecessors in P', and otherwise to be at *level i* iff it has at least one predecessor at each of levels $0, \ldots, i-1$ but none at level i. Clearly, the rotations exposed in S_0 (the male-optimal matching) are exactly the ones at level 0. If we eliminate those, the rotations exposed in the resulting stable matching S_1 are exactly the ones at level 1, etc. Let S_i be the stable matching obtained by eliminating all the rotations in levels $0, \ldots, i-1$.Over the entire computation, the time needed to eliminate all the rotations, reduce the resulting lists and update $s(S_i, m)$ and $s'(S_i, m)$ is $O(n^2)$. For each level i, each graph $G(S_i)$ is built and all cycles identified in $O(n)$-time. Then, since there can be at most $O(n^2)$ levels, the theorem follows. □

Finding the max-weight closed subset of P'

We will use network flow to find a maximum-weight closed subset of P' (henceforth denoted as W^*).

From P' we construct a capacitated $s - t$ flow graph $P'(s, t)$ as follows:

- A source node s and a sink node t are added to P'.
- For every node ρ_i of negative weight, a $(s \to \rho_i)$ edge is added, with capacity $|w(\rho_i)|$.
- For every node ρ_j of positive weight, a $(\rho_j \to t)$ edge is added, with capacity $w(\rho_j)$.
- The capacity of every original edge of P' is set to ∞.

Note that if $w(\rho) = 0$ then P' contains neither edge $(s \to \rho)$ nor $(\rho \to t)$.

$P'(s,t)$ should look somewhat like this:

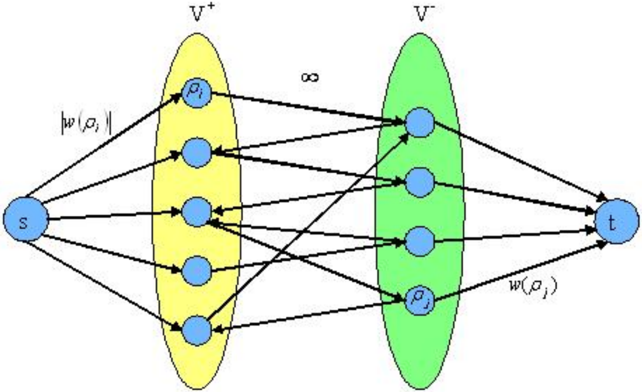

Finding the max-weight closed subset of P'

We will prove the following theorem:

Theorem

Let X be the set of edges crossing a minimum $s - t$ cut in $P'(s, t)$ and denote the capacity of X by $w(X)$. The positive nodes of W^* (denoted as W^*_+) are exactly the positive nodes whose edges into t are uncut by X. These nodes and all the nodes that reach them in P' define a maximum-weight closed subset in P'.

We denote by V^+ and V^- the sets of positive and negatives nodes in P' respectively and by $N(W)$ the set of all negative predecessors of nodes in $W \subseteq V^+$.

Finding the max-weight closed subset of P'

Proof.

Any negative node in a maximum-weight closed subset C must precede at least one positive node in C, hence the set of positive nodes of W^* is:
$W_+^* = \{W \subseteq V^+ |\ w(W) - |w(N(W))| \text{ is maximized}\} =$
$= \{W \subseteq V^+ |\ w(V^+) - (w(W) - |w(N(W))|) \text{ is minimized}\} =$
$= \{W \subseteq V^+ |\ w(V^+ \smallsetminus W) + |w(N(W))| \text{ is minimized}\}$.

• Now let $W \subseteq V^+$. In the graph $P'(s,t)$, if every edge from s to $N(W)$ is cut along with any edge from $(V^+ \smallsetminus W)$ to t, then all (s,t)-paths are cut. Hence, $w(X) \leq w(V^+ \smallsetminus W) + |w(N(W))|$ for any $W \subseteq V^+$.

• Conversely, if we let $W_+ \subseteq V^+$ consist of the positive nodes whose edges to t were uncut by X, then, by definition, X cuts all edges from $(V^+ \smallsetminus W_+)$ to t and also all edges from s to $N(W_+)$ (since X is an $s-t$ cut of finite capacity and all $V^- \to V^+$ edges have infinite capacity). Hence, $w(X) = w(V^+ \smallsetminus W_+) + |w(N(W_+))| \leq w(V^+ \smallsetminus W) + |w(N(W))|$ for any $W \subseteq V^+$, and it follows that $W_+ = W_+^*$. □

Finding the max-weight closed subset of P'

Given the positive nodes W_+^* in a maximum-weight closed subset W^* of P', we can construct an optimal stable matching in $O(n^2)$-time by the following procedure:

- We find the predecessors of W_+^* in P' by scanning backwards from the nodes in W_+^*, marking unmarked nodes that are reached ang again scanning backwards from them. This takes $O(n^2)$-time, since P' has at that many edges.
- Now all the marked nodes constitute W^*, so all we have to do is simply start with the male-optimal stable matching and traverse P' top down, by levels, eliminating all marked rotations. Again, due to the size of P', this step takes $O(n^2)$-time.

The resulting stable matching is optimal.

Finding the max-weight closed subset of P'

In order to find the minimum $s - t$ cut in $P'(s,t)$ we will use the maximum-flow algorithm of Sleator and Trajan, which needs $O(VE \log V)$-time. Since $P'(s,t)$ has $V = O(n^2)$ nodes and $E = O(n^2)$ edges, that takes $O(n^2 \log n)$. In order to reduce that to $O(n^4)$ we will use the fact that the minimum cut in $P'(s,t)$ has capacity bounded by $O(n^2)$.

When all the capacities are integral, the running time of the Ford-Fulkerson algorithm (and of the Sleator-Trajan as a subsequent algorithm) is $O(EK)$, where K is the maximum $s - t$ flow value. Since all the capacities in $P'(s,t)$ are integral and both E and K are $O(n^2)$, it follows that both the max-flow and the min-cut can be found in $O(n^4)$-time.

Finding the max-weight closed subset of P'

In order to find the minimum $s - t$ cut in $P'(s,t)$ we will use the maximum-flow algorithm of Sleator and Trajan, which needs $O(VE \log V)$-time. Since $P'(s,t)$ has $V = O(n^2)$ nodes and $E = O(n^2)$ edges, that takes $O(n^2 \log n)$. In order to reduce that to $O(n^4)$ we will use the fact that the minimum cut in $P'(s,t)$ has capacity bounded by $O(n^2)$.

When all the capacities are integral, the running time of the Ford-Fulkerson algorithm (and of the Sleator-Trajan as a subsequent algorithm) is $O(EK)$, where K is the maximum $s - t$ flow value. Since all the capacities in $P'(s,t)$ are integral and both E and K are $O(n^2)$, it follows that both the max-flow and the min-cut can be found in $O(n^4)$-time.

An $O(n^4)$-Time Solution

To summarize, we can:

- find all the rotations in $O(n^3)$
- construct $P'(s,t)$ in $O(n^2)$
- find a minimum $s-t$ cut X of $P'(s,t)$ in $O(n^4)$
- use X to find a minimum-weighted closed subset W^* of P' in $O(n^4)$
- eliminate each rotation in W^* in $O(n^2)$ to obtain a stable matching

Note that other than the $O(n^4)$ network flow computation, the rest of the algorithm needs just $O(n^2)$, so if a more specialized flow-algorithm than the Sleator-Tarjan could be found for $P'(s,t)$ or other sparse subgraphs of P that preserve the closed subset, it would immediately speed up the whole algorithm.

Generalization

One last interesting observation is that we could easily generalize the stable matching problem so that each person i not only has a rank ordering of the members of the opposite sex, but also a numerical preference $w(i,j)$ for each such person j.

Again these preferences define a rank ordering, so the notion of a stable matching is unchanged. We now ask for the stable matching that maximizes $\sum_1^n p(m_i, w_i) + \sum_1^n p(w_i, m_i)$. This problem can be solved in $O(n^4 \log n)$-time using the same method.

...and a fun fact about this problem:
It has been shown that there is no matching mechanism that is both **strategy-proof** and produces **stable** matchings.

One last interesting observation is that we could easily generalize the stable matching problem so that each person i not only has a rank ordering of the members of the opposite sex, but also a numerical preference $w(i, j)$ for each such person j.

Again these preferences define a rank ordering, so the notion of a stable matching is unchanged. We now ask for the stable matching that maximizes $\sum_{1}^{n} p(m_i, w_i) + \sum_{1}^{n} p(w_i, m_i)$. This problem can be solved in $O(n^4 \log n)$-time using the same method.

...and a fun fact about this problem:
It has been shown that there is no matching mechanism that is both **strategy-proof** and produces **stable** matchings.

Thank you!

¡Muchas gracias!

Ευχαριστώ πολύ!