A Note on Distributed Stable Matching

Alex Kipnis Boaz Patt-Shamir

presented by Ofir Reginiano
May 2010

Matching

- men & women
- each man ranks the women
- each woman ranks the men
- a matching (marriage) M
  matches each man with a woman

example

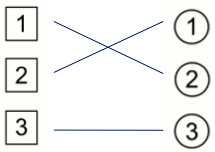

Preferences

□→○			
1	1	3	2
2	3	2	1
3	1	2	3

○→□			
①	2	1	3
②	3	1	2
③	1	3	2

Stable Matching

- (i, j) blocking:
  both i and j prefer each other
  to their matches under M
- stable matching:
  a matching without
  any blocking pair

example

Stable matching

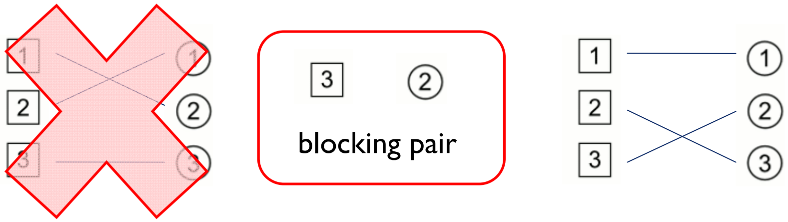

Preferences

	□→○				○→□		
1	1	3	2	1	2	1	3
2	3	2	1	2	3	1	2
3	1	2	3	3	1	3	2

outline

- Motivation
- Problem Statement
  - DSM
- Time Lower Bound
  - Mailing Problem
  - DSM >> Mailing Problem
- ε-DSM
  - Algorithm
  - Time Upper Bound

Motivation

- Gale and Shapley (1962)
  - SM exists
  - O (n^2)
- Gusfield and Irving (1989)
  - ties
  - unacceptable pairs
- 3 way SM
  - *men, women, dogs*
- Stable Roommates (SR) problem
- Scheduling network switches

DSM – Distributed Stable Matching

- G = (V,E) undirected bipartite
- each node ranks its neighbors
- matching = a set of disjoint edges

- synchronous model
- in each round, each processor can send a message to each of its neighbors
- CONGEST model –
  a message may contain up to B bits
  (e.g. B = θ(log n))

The Mailing Problem

- disjoint sets S,R ⊆ V
  (senders and receivers)

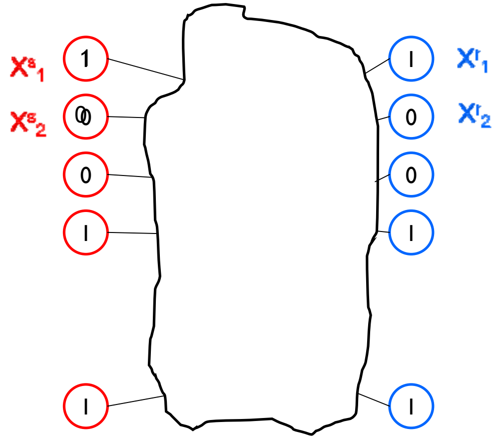

The Graph Family G

- $G_{\Gamma,m,p}$ m≥11 odd n= θ(Γ˙m)

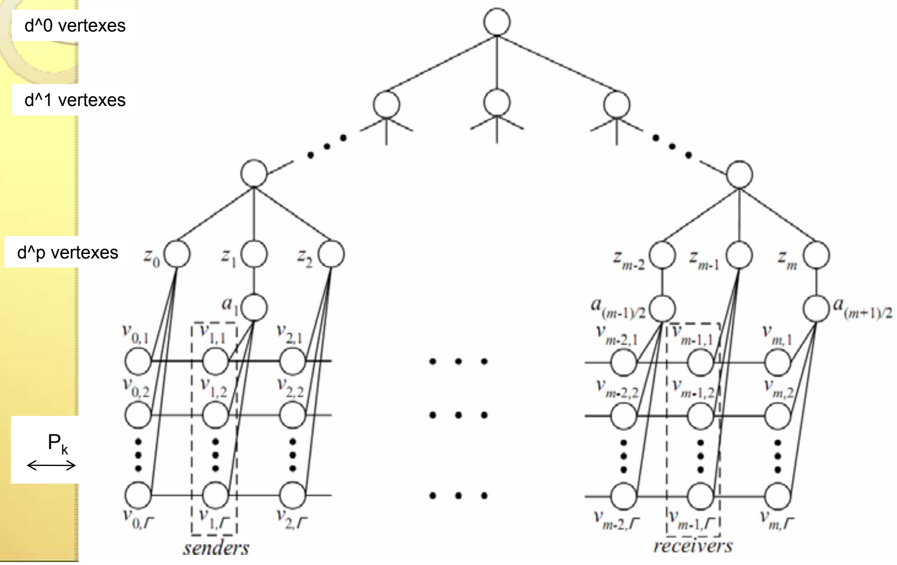

$G_{\Gamma,m,p}$ properties

- $G_{\Gamma,m,p} \in$ G is bipartite
- diameter: d = 2p+4

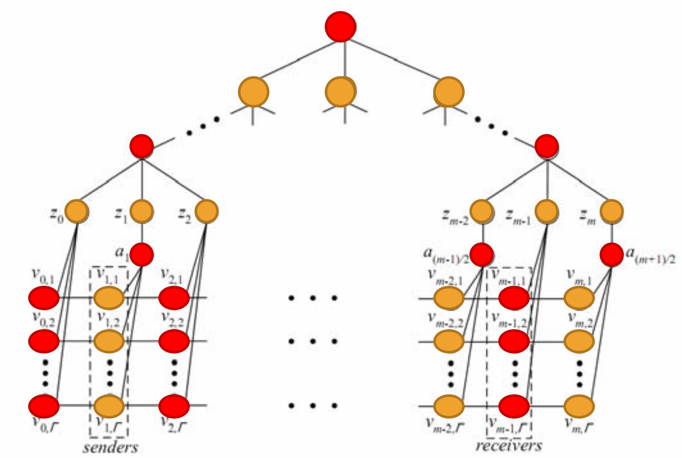

Mailing Problem in $G_{\Gamma,m,p}$

- $m = \left(\frac{n}{p \cdot B}\right)^{1/2-\frac{1}{2(2p+1)}}$

  Any deterministic protocol
  that delivers at least c˙Γ bits correctly
  (0.8<c≤1 const)
  requires Ω(m) rounds

special cases

$$m = \left(\frac{n}{p \cdot B}\right)^{1/2 - \frac{1}{2(2p+1)}}$$

- p = 1…log(n)
- p = 1 → d = 6

$\left(\frac{n}{\log n}\right)^{\frac{1}{3}}$ rounds

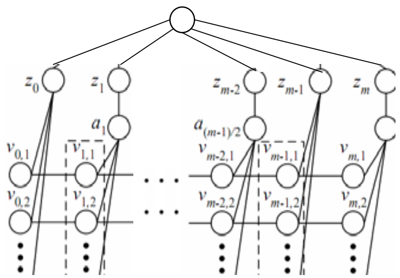

- p = log(n)

$\left(\frac{n}{\log^2 n}\right)^{\frac{1}{2}}$ rounds

alternating path

- blue edges ⊆ M and
  green edges ∩ M = ∅ ,
  or vice versa

Reduce Mailing Problem to DSM Problem

- in the d-ary tree:
  - vertices with even distance from the root: their leftmost child – most preferred
  - others: their parent – most preferred
- on the paths P_j :

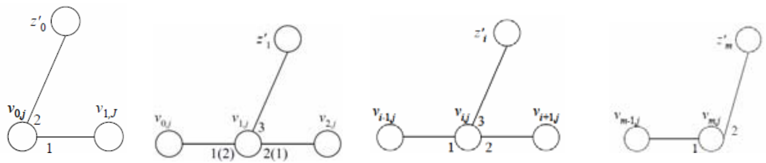

Lemma: Under any stable matching, P_j is an alternating path for all $1 \le j \le \Gamma$

case 1: i ≠ 1

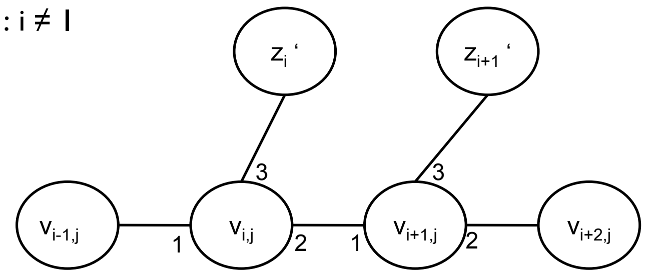

in M

not in M

Lemma: Under any stable matching,

P_j is an alternating path for all $1 \le j \le \Gamma$

case 2: i = 1

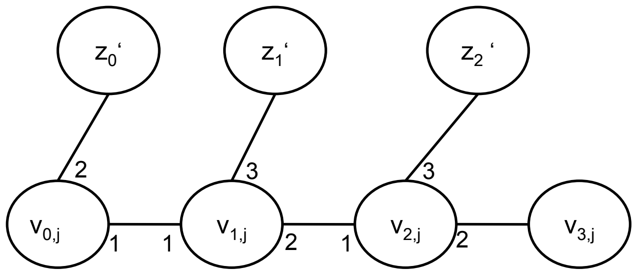

in M

not in M

Lemma: Under any stable matching,

P_j is an alternating path for all $1 \le j \le \Gamma$

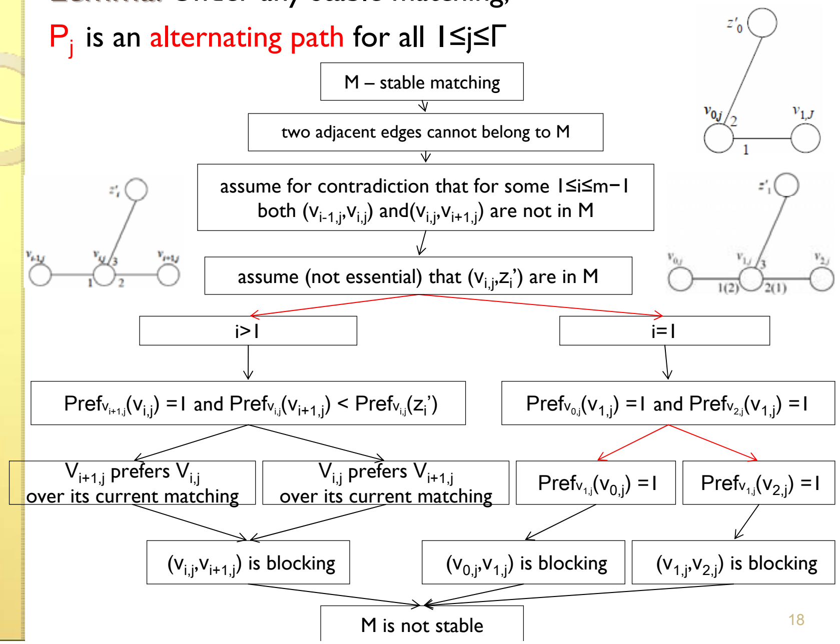

corollaries

- only 2 possible matchings along each P_j
  - determined only by the preference of $v_{1,j}$
  - do not depend on the tree

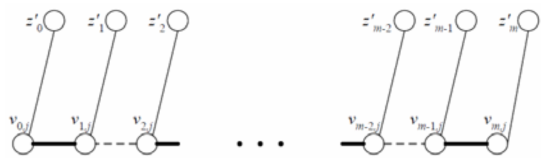

- Any deterministic algorithm for DSM requires Ω(D) rounds (D – diameter)

ρ-approximation

- ρ-approximation matching if it induces at most (1−ρ)n blocking pairs
- ρ = 1 → "regular" case

A Time Lower Bound on DSM Problem

- any deterministic ρ-approximation
- $\rho = 1 - o(1/\sqrt{n})$
- diameter $D \in \{6, 8, 10, \dots\}$
- The algorithm requires

$$\Omega\left(D + \left(\frac{n}{D \cdot B}\right)^{\frac{1}{2} - \frac{1}{2(D-3)}}\right)$$

communication rounds

Proof

reduction: Mailing problem → DSM

- input: $X^{s}_1, \dots, X^{s}_{\Gamma}$ to the mailing problem in G
- construct a DSM instance:
  - preferences of all nodes, except the senders, are fixed
  - $v_{1,j}$ depends if $X^{s}_j = 1$ (or 0)

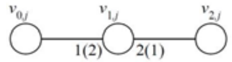

  - when DSM terminates, each node knows its match
- output transformation: receivers $v_{m-1,j}$ set $X^{r}_j = 1$ ←→ $(v_{m-1,j}, v_{m,j})$ is in the matching

Proof – cont.

- suppose M is a ρ-approximate stable matching
- # blocking pairs $\leq (1-\rho)n = o(\sqrt{n})$
- # paths Pj without any blocking pair $\geq \Gamma - o(\sqrt{n})$
- # paths Pj which are alternating $\geq \Gamma - o(\sqrt{n})$

- output variables in the receivers are correct
- $\Gamma = \theta(n/m) >> \sqrt{n}$
- # correct bits delivered using the reduction $\geq \Gamma(1 - o(1))$

- running time under these specific graphs is

$$\Omega\left(\left(\frac{n}{D \cdot B}\right)^{\frac{1}{2} - \frac{1}{2(D-3)}}\right)$$

Corollary - Time Lower Bound

- Any deterministic ρ-approximation
  for DSM
  graph's diameter $\theta(\log n)$
  $\rho = 1 - o(1/\sqrt{n})$
  requires $\Omega\left(\sqrt{\frac{n}{B \cdot \log n}}\right)$ rounds

even if we allow
$\sqrt{n}$ blocking pairs
in the output !

some notations…

- *$Pref_v$*
  preference – ordered list of all neighbors of node v.
- *$Pref_v(u) = i$* u is the *i*-th node in *$Pref_v$*
  $Pref_v(u)$ = 1 (most preferred) , …
  ∞ (not a neighbor)
- (u, v) ∈ M ←→ *M(v) = u* and M(u) = v
- v is unmatched under M → $Pref_v(M(v)) = \infty$

Extension: ε-DSM

- $0 \leq \varepsilon \leq 1$
- Stability is violated only when there is a pair that can improve their matches "a lot" by matching to each other
- (u, v) ε-blocking pair with respect to M if
  $Pref_v(u) < Pref_v(M(v)) - \varepsilon\,|Pref_v|$ and
  $Pref_u(v) < Pref_u(M(u)) - \varepsilon\,|Pref_u|$
- exact case: $\varepsilon \leq 1/(n+1)$

example

ε = 0
ε = 1/3

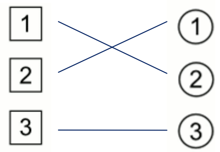

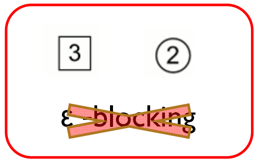

Preferences

□→○				○→□			
1	1	3	2	1	2	1	3
2	3	2	1	2	3	1	2
3	1	2	3	3	1	3	2

$G^{\varepsilon\text{-DSM}}$

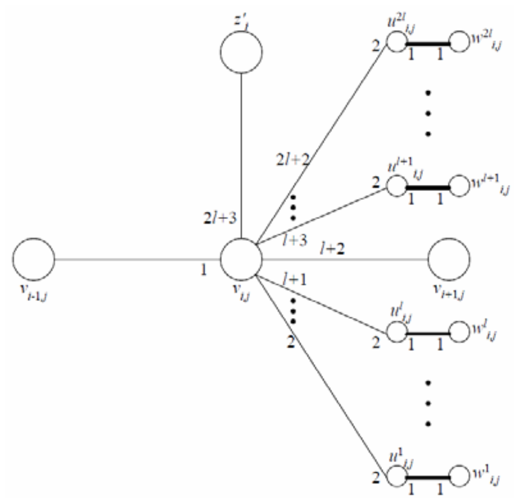

Lemma: for any $\varepsilon < \frac{l+1}{2l+3}$

P_j is an alternating path for all $1 \le j \le \Gamma$

e.g. case: $(i+1) \neq m$

$i \neq 1$

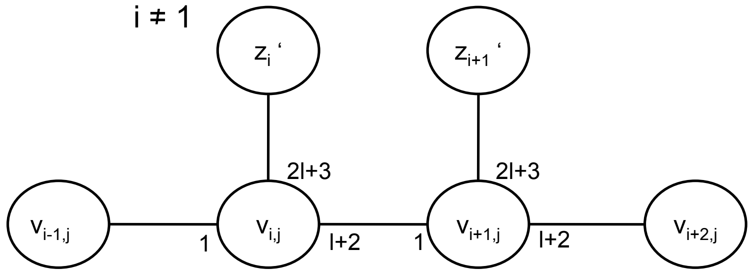

in M $\quad \mathrm{Pref}_{v_{i,j}}(z_i') - \mathrm{Pref}_{v_{i,j}}(v_{i+1,j}) = l+1 > \varepsilon \cdot |\mathrm{Pref}_{v_{i,j}}|$

not in M $\quad \mathrm{Pref}_{v_{i+1,j}}(M(v_{i+1,j})) - \mathrm{Pref}_{v_{i+1,j}}(v_{i,j}) \ge l+1 > \varepsilon \cdot |\mathrm{Pref}_{v_{i+1,j}}|$

Lemma: for any $\varepsilon < \frac{l+1}{2l+3}$

P_j is an alternating path for all $1 \le j \le \Gamma$

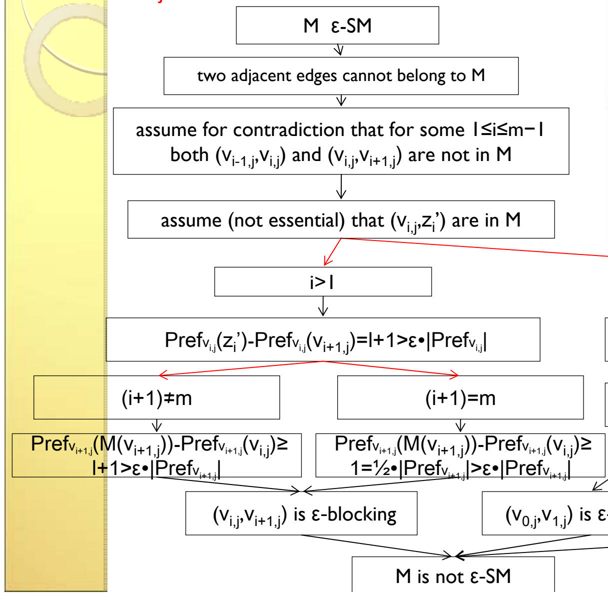

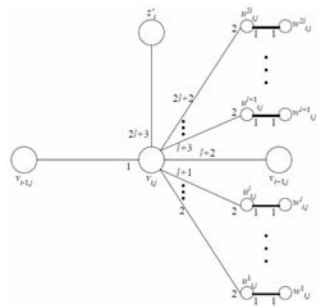

ε-DSM result

- Any deterministic ρ-approximation
  ρ=1−o(1/√n)
  diameter D∈{10,12,14, . . }
  ε < ½ constant
  requires $\Omega\left(D + \left(\frac{n}{D \cdot B}\right)^{\frac{1}{2} - \frac{1}{2(D-3)}}\right)$
  rounds

extension to ε-stability
(ε ~ ½)

ε-DSM algorithm

- an extension of Gale and Shapley alg.
- a termination detection
- bipartite graph – 2 node sets:
  - active – may only propose
  - passive – may only reject
- proceed in double rounds

u active

if $M(u) = null$ **then**
  $v \leftarrow$ *most preferred passive node on* Pref_u
  propose to v
  $M(u) \leftarrow v$
for *all* $v \in reject_u$ **do**
  if $M(u) = v$ **then**
    $M(u) \leftarrow null$
  remove v *from* Pref_u
$reject_u \leftarrow null$

v passive

$candidate \leftarrow$ *most preferred active node on* $proposers_v$
remove candidate from $proposers_v$
for *all* $u \in proposers_v$ **do**
  reject u
  remove u *from* Pref_v

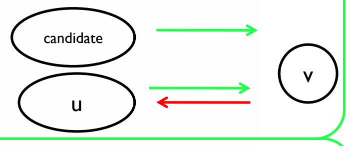

if $M(v) = null$ **then**
  $M(v) \leftarrow candidate$
else if $\mathrm{Pref}_v(candidate) < \mathrm{Pref}_v(M(v)) - \varepsilon \cdot degree_v$ **then**
  reject $M(v)$
  remove $M(v)$ *from* Pref_v
  $M(v) \leftarrow candidate$
else
  reject candidate
  remove candidate from Pref_v

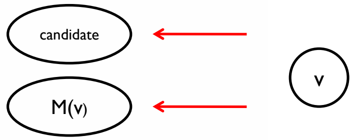

for *all* $u \in \mathrm{Pref}_v$ **do**
  if $\mathrm{Pref}_v(u) \geq \mathrm{Pref}_v(M(v)) - \varepsilon \cdot degree_v$ **then**
    reject u
    remove u *from* Pref_v

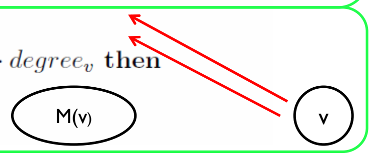

Round : 1

http://en.wikipedia.org/wiki/Stable_marriage_problem

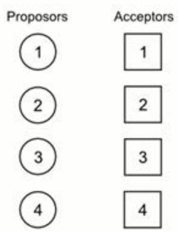

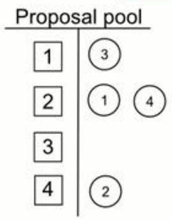

• 1-4 propose, as none are currently tentatively attached

Preferences

○→□ Proposor Table

(1)	2	1	3	4
(2)	4	1	2	3
(3)	1	3	2	4
(4)	2	3	1	4

□→○ Acceptor Table

[1]	1	3	2	4
[2]	3	4	1	2
[3]	4	2	3	1
[4]	3	2	1	4

Round : 1

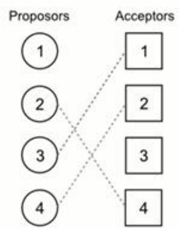

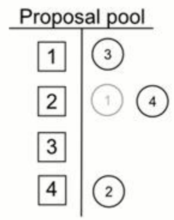

• 1 accepts 3's proposal—no better offer.
• 2 accepts 4's proposal as 4 is more prefereable to 1.
• 3 recieves no offer.
• 4 accepts 2's proposal—no better offer.

Preferences

○→□ Proposor Table

(1)	2	1	3	4
(2)	4	1	2	3
(3)	1	3	2	4
(4)	2	3	1	4

□→○ Acceptor Table

[1]	1	3	2	4
[2]	3	4	1	2
[3]	4	2	3	1
[4]	3	2	1	4

Round : 2

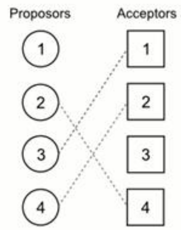

Proposal pool

1	1
2	
3	
4	

- 1, the only un-attached membe makes its offer to 1, its first preference not previously proposed to.

Preferences

○→□ Proposor Table

1	2	1	3	4
2	4	1	2	3
3	1	3	2	4
4	2	3	1	4

□→○ Acceptor Table

1	1	3	2	4
2	3	4	1	2
3	4	2	3	1
4	3	2	1	4

Round : 2

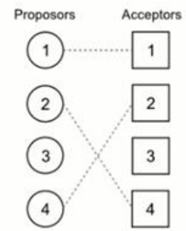

Proposal pool

1	1
2	
3	
4	

- 1 drops 3's propsal in favour of as this is higher in its preference table. 3 returns to the proposal pool.

Preferences

○→□ Proposor Table

1	2	1	3	4
2	4	1	2	3
3	1	3	2	4
4	2	3	1	4

□→○ Acceptor Table

1	1	3	2	4
2	3	4	1	2
3	4	2	3	1
4	3	2	1	4

Round : 3

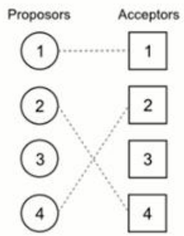

Proposal pool

1	
2	
3	3
4	

- 3 proposes to 3

Preferences

○→□ Proposor Table

1	2	1	3	4
2	4	1	2	3
3	1	3	2	4
4	2	3	1	4

□→○ Acceptor Table

1	1	3	2	4
2	3	4	1	2
3	4	2	3	1
4	3	2	1	4

Round : 3

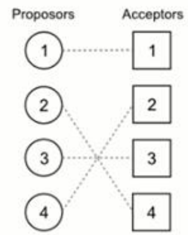

Proposal pool

1	
2	
3	3
4	

- 3 accepts 3, not having a better offer

Preferences

○→□ Proposor Table

1	2	1	3	4
2	4	1	2	3
3	1	3	2	4
4	2	3	1	4

□→○ Acceptor Table

1	1	3	2	4
2	3	4	1	2
3	4	2	3	1
4	3	2	1	4

Finish

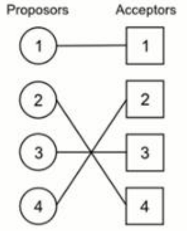

- No two members {P,A} would prefer one-another over their current pairing

Preferences

○→□ Proposor Table

(1)	2	1	3	4
(2)	4	1	2	3
(3)	1	3	2	4
(4)	2	3	1	4

□→○ Acceptor Table

[1]	**1**	3	2	4
[2]	3	**4**	1	2
[3]	4	2	**3**	1
[4]	3	**2**	1	4

Alg. Termination

- when there are no more rejected nodes
- termination detection mechanism:
  - each rejected node broadcasts (on a shortest-path tree) that the protocol has not terminated yet
  - a node halts if D (graph's diameter) consecutive steps pass without any such broadcast received

Alg. Time Complexity

ε-DSM can be solved on any graph in

$$O(D + \min(|E|, \varepsilon^{-1} n))$$

Proof:

- O(D) to find the value of D and to detect termination
- before stabilization of the matching
  - in each round at least one passive node improves the rank of its match by at least a fraction ε.
    → the total number of rounds is $O(\varepsilon^{-1}n)$
  - in each round at least one fresh active node is rejected by a passive node
    → num. of rounds is also bounded by |E|

Conclusion

- stable Marriage problem in a distributed model
- bounded length messages

- any alg. requires $\Omega(\sqrt{n/B \log n})$ communication rounds in the worst case
- even for graphs of diameter O(log n)
- even if allowing O(√n) blocking pairs in the output
- our lower bound extends to ε-stability (ε ~ ½)

- our best upper bound is O(|E|) communication rounds, or $O(\varepsilon^{-1}n)$ for ε-DSM