CS4800: Algorithms & Data
Jonathan Ullman

Lecture 23:
- Stable Matching: the Gale-Shapley Algorithm

Apr 13, 2018

Final Exam is comprehensive up through greedy algorithms:

morning of F 4/27

- Asymptotic notation
- Induction
- Recurrences
- Divide and conquer
- DP

} MT I

- Graph algs } MT II
- Reductions / Network Flow
- Greedy

Labor Markets

- Every year about 40k doctors graduate from medical school and need a residency…
- … somehow they get assigned to the roughly 40k residencies available at US hospitals

David Gale (1921-2008)
PROFESSOR, UC BERKELEY

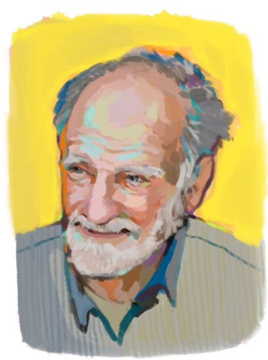

Lloyd Shapley
PROFESSOR EMERITUS, UCLA

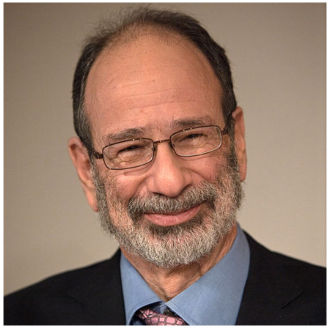

Alvin Roth
PROFESSOR, STANFORD

Why are labor markets a pain?

- Not everyone can get their favorite job ☹

- Nobody has the full picture!
    - How do you decide whether to accept an exploding offer now or keep looking for a better job?

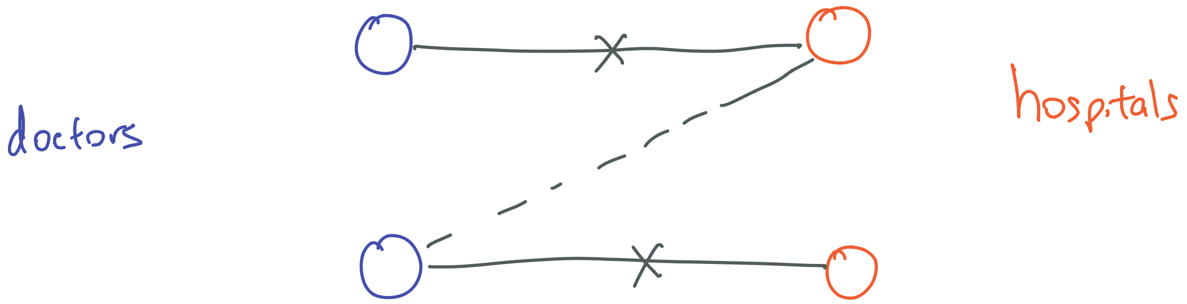

What can we do about it?

- What if we could just assign jobs?
  - What information would we need?
  - How would we choose the assignment?

Some kind of preference information

Numerical preferences are difficult:

- People dont really know them?
- They are to compare

Ranking are easier to deal with

Stable Matching

doctor d

$h_{37} \succ h_{26} \succ h_1 \succ \dots \succ h_{92}$

hospital h

$d_2 \succ d_6 \succ d_{111} \succ \dots \succ d_4$

- n doctors and n hospitals
  - each doctor wants one job
  - each hospital wants one resident
  - every doctor has a ranking of hospitals
  - every hospital has a ranking of doctors

- want a perfect matching μ of doctors to hospitals
  - $\mu(d) = h$ means "d is matched to h"
  - $\mu(d) = \emptyset$ means "d is unmatched"
  - perfect means there are n matches

always possible

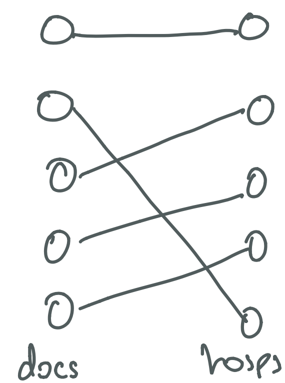

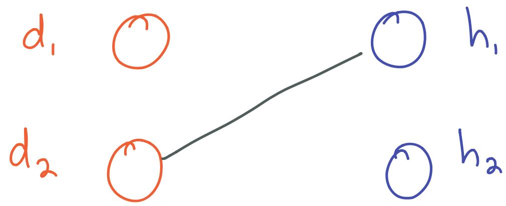

$$\mu(d_1) = \emptyset$$
$$\mu(d_2) = h_1$$
$$\mu(h_1) = d_2$$
$$\mu(h_2) = \emptyset$$

Stable Matching

- We want our matching μ to be stable
  - No doctor and hospital want to break their assignment and match themselves

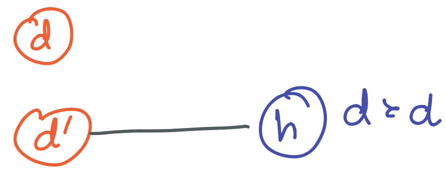

- Instabilities
  - d, h such that $\mu(d) = \emptyset$ and $d \succ_h \mu(h)$
  - h, d such that $\mu(h) = \emptyset$ and $h \succ_d \mu(d)$
  - d, d', h, h' such that $h \succ_d \mu(d)$ and $d \succ_h \mu(h)$

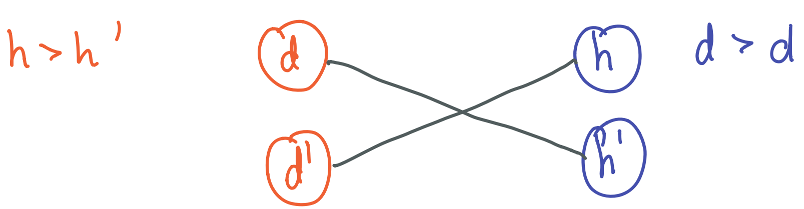

Gale-Shapley Algorithm

- Start with empty matching μ
- While (some hospital h is unmatched): every hospital has offered a job to every doc
  - Let d be the h's highest ranked doctor that it hasn't already offered a job to (h offers a job to d)
  - If ($\mu(d) = \emptyset$): If d has no job, d tentatively accepts
    - Set $\mu(h) \leftarrow d$ and $\mu(d) \leftarrow h$
  - ElseIf ($\mu(d) = h'$ and $h \succ_d h'$): d likes h better than its current offer
    - Set $\mu(h) \leftarrow d$, $\mu(d) \leftarrow h$, and $\mu(h') \leftarrow \emptyset$
  - Else:
    - Do nothing
- Output the matching

Always pick the highest unmatched

Gale-Shapley Demo

	1st	2nd	3rd	4th	5th
MGH	~~Bob~~	Alice	Dorit	Ernie	Clara
BW	~~Dorit~~	~~Bob~~	~~Alice~~	Clara	Ernie
BID	Bob	Ernie	Clara	Dorit	Alice
MTA	~~Alice~~	~~Dorit~~	~~Clara~~	~~Bob~~	Ernie
CH	~~Bob~~	Dorit	Alice	Ernie	Clara

	1st	2nd	3rd	4th	5th
Alice	CH	MGH	~~BW~~	~~MTA~~	BID
Bob	BID	~~BW~~	~~MTA~~	~~MGH~~	~~CH~~
Clara	BW	BID	~~MTA~~	CH	MGH
Dorit	MGH	CH	~~MTA~~	BID	~~BW~~
Ernie	MTA	BW	CH	BID	MGH

produce the reverse lookup in $O(n^2)$

	Alice	Bob	Clare
MGH	2	4	5
BW	3	2	1
BID			
MTA			
CH			

Observations

1. Hospitals go down their rankings

   If h is matched to d then h has not offered a job to a lower ranked d'

2. Doctors who get jobs always have jobs

   If d has ever been matched then d is still matched

3. Doctors go up their rankings

   If d is matched to h at some point then d is always matched to someone at least as good.

Gale-Shapley Algorithm

- Questions
  - Does there even exist a stable matching?
  - Will this algorithm terminate?
  - Does it output a perfect matching? → all n docs have jobs
  - Does it output a stable matching? → no instabilities
  - How do we implement this algorithm efficiently?

Gale-Shapley Algorithm

- Termination: GS algorithm terminates after n^2 iterations of the main loop

  - Never have h offer a job to d twice
  - Only n^2 pairs (h, d).

Gale-Shapley Algorithm

- Matching: GS algorithm outputs a perfect matching
  - If there is ever an unmatched hospital, then there is a doctor it has not yet offered a job to

    Suppose h is unmatched. There $\exists$ d that is unmatched.
    h must not have made an offer to d, because if d ever rec'd an offer he would be matched.
  - Therefore, GS returns a perfect matching

    If we terminate it must be b/c all h are matched

Gale-Shapley Algorithm

- Stability: GS algorithm outputs a stable matching
- Proof by contradiction:
  - Suppose there is an instability d, d', h, h'

$d > d'$

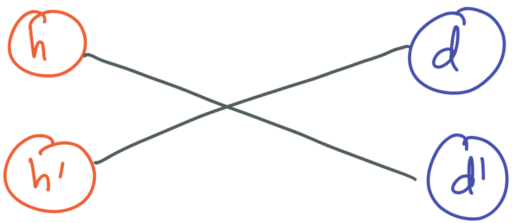

$h > h'$

case 1:
- h never offered a job to d
- h has offered a job to d'
- but then $d' >_h d$

Gale-Shapley Algorithm

- Stability: GS algorithm outputs a stable matching
- Proof by contradiction:
  - Suppose there is an instability d, d', h, h'

$d > d'$

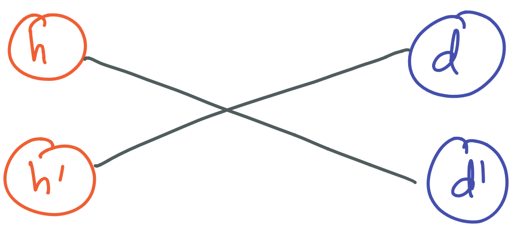

$h > h'$

case 2: (h offered a job to d)

either d rejected, meaning d was already matched to some h'' that $h'' \succ_d h$, now d ended up with h', so $h' \succ_d h'' \succ_d h$.

Gale-Shapley Algorithm

- But wait, what if there is no stable matching?
- Cool example of an algorithmic proof of a purely mathematical fact.

Gale-Shapley Algorithm

- Running Time: GS runs in $O(n^2)$ time
  - Needs to be implemented properly
  - Note that the size of the input is $2n$ arrays of len n, so the size of the input itself is $\Omega(n^2)$

Real World Impact

TABLE I

STABLE AND UNSTABLE (CENTRALIZED) MECHANISMS

Market	Stable	Still in use (halted unraveling)
American medical markets		
NRMP	yes	yes (new design in '98)
Medical Specialties	yes	yes (about 30 markets)
British Regional Medical Markets		
Edinburgh ('69)	yes	yes
Cardiff	yes	yes
Birmingham	no	no
Edinburgh ('67)	no	no
Newcastle	no	no
Sheffield	no	no
Cambridge	no	yes
London Hospital	no	yes
Other healthcare markets		
Dental Residencies	yes	yes
Osteopaths (<'94)	no	no
Osteopaths (≥'94)	yes	yes
Pharmacists	yes	yes
Other markets and matching processes		
Canadian Lawyers	yes	yes (except in British Columbia since 1996)
Sororities	yes (at equilibrium)	yes

Table 1. Reproduced from Roth (2002, Table 1).

Real World Impact

- Doctors to Hospitals
  - Have to deal with two-body problems
  - Have to make sure doctors do not game the system
- Kidneys to Patients
  - Not all matches are feasible (blood types, histones)
  - Certain pairs must be matched
- Students to Public Schools
  - Siblings, walking zones, diversity
- Reform Rabbis to Synagogues
  - No idea, just a fun example

!!!NOBEL PRIZE IN 2012!!!