CS3000: Algorithms & Data
Jonathan Ullman

Lecture 2:

- Stable Matching: the Gale-Shapley Algorithm

Sep 11, 2018

National Residency Matching Program

- National system for matching US medical school graduates to medical residencies
  - Roughly 40,000 doctors per year
  - Assignment is almost entirely algorithmic

David Gale (1921-2008)
PROFESSOR, UC BERKELEY

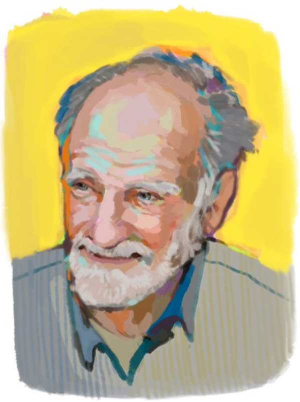

Lloyd Shapley
PROFESSOR EMERITUS, UCLA

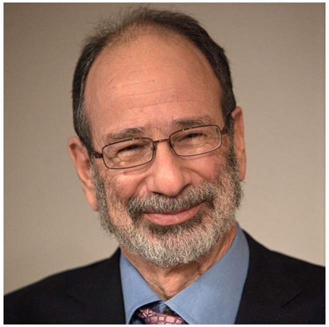

Alvin Roth
PROFESSOR, STANFORD

Labor Markets

- Most labor markets are frustrating
  - Not everyone can get their favorite job
  - The market is decentralized

- Decentralized labor markets are confusing

Nobody has all the information

Whatever you do could lead to an untenable

Centralized Labor Markets

- What if we could just assign jobs?
  - What information would we want?

    List of doctors and hospitals

    Preferences (ranking, ordinal preferences)

    ↳ from each doctor and each hospital
  - How would we choose the assignment?

    Stable

Matchings

In the real world, doctors only rank ≤ 15 hospitals

- We are given the following information
  - n doctors $d_1 \dots d_n$
  - n hospitals $h_1 \dots h_n$
  - each doctor's ranking of hospitals $d_1 : h_2 \succ h_3 \succ h_1$
  - each hospital's ranking of doctors $h_1 : d_1 \succ d_3 \succ d_2$

simplifying assumption

	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

Matchings

- A matching M is a set of doctor-hospital pairs
  - $M = \{ (d_1, h_2), (d_2, h_3) \}$
  - matching: no doctor/hospital appears twice
  - perfect matching: every doctor/hospital appears once
  - "d is matched to h": $(d, h) \in M$

"d is matched": $\exists\, h$ s.t. $(d, h) \in M$

"d is unmatched"

Stable Matchings

- A matching M is unstable if some doctor-hospital pair prefer one another to their mate in M

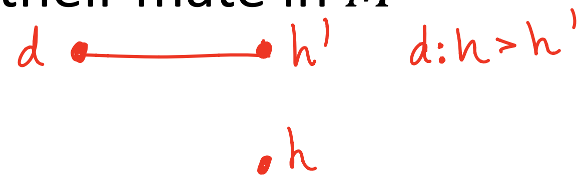

- Instabilities
  1. d, h such that d is matched to h', h is unmatched, but $d : h \succ h'$
  2. d, h such that h is matched to d', d is unmatched, but $h : d \succ d'$
  3. d, h such that d is matched to h', h is matched to d', but $d : h \succ h'$ and $h : d \succ d'$

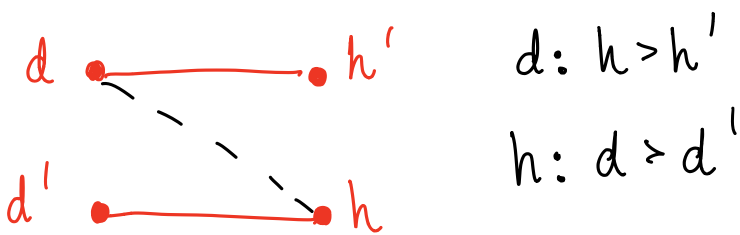

Ask the Audience

- Either find a stable matching or convince yourself that there is no stable matching

	1st	2nd	3rd
MGH	Alice	Bob	Clara
BW	Bob	Clara	Alice
BID	Alice	Clara	Bob

	1st	2nd	3rd
Alice	BW	BID	MGH
Bob	BID	MGH	BW
Clara	MGH	BID	BW

M = { (Alice, BW), (Bob, MGH), (Clara, BID) }

M' = { (Alice, BID), (Bob, MGH), (Clara, BW) }

M'' = { (Alice, BW), (Bob, BID), (Clara, MGH) }

Gale-Shapley Algorithm

```
• Let M be empty
• While (some hospital h is unmatched):
    • If (h has offered a job to everyone): break
    • Else: let d be the highest-ranked doctor to
      which h has not yet offered a job
    • h makes an offer to d:
        • If (d is unmatched):
            • d accepts, add (d,h) to M
        • ElseIf (d is matched to h′ & d: h′ > h):
            • d rejects, do nothing
        • ElseIf (d is matched to h′ & d: h > h′):
            • d accepts, remove (d,h′) from M and
              add (d,h) to M
• Output M
```

A job offer

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

Observations

- Hospitals make offers in descending order

  If h made offers to d, d' and d got an offer first, then $h: d > d'$

- Doctors that get a job never become unemployed

  If a doctor has ever had a job, they will always have a job.

- Doctors accept offers in ascending order

  If a doctor was ever matched to h, then d is never matched to a lower ranked hospital than h.

Gale-Shapley Algorithm

- Questions about the Gale-Shapley Algorithm:
  - Will this algorithm terminate?
  - Does it output a perfect matching?
  - Does it output a stable matching? (Does one even exist?)
  - How do we implement this algorithm efficiently?

GS Algorithm: Termination

- **Claim:** The GS algorithm terminates after n^2 iterations of the main loop

  - There are only n^2 doctor-hospital pairs
  - Never make the same offer twice
  - Alg halts if all offers are made

GS Algorithm: Perfect Matching

- **Claim:** The GS algorithm returns a perfect matching (all doctors/hospitals are matched)

Proof by Contradiction:

- Suppose some h is unmatched at the end.
- ⇒ there is some d that is unmatched
- B/c the alg terminated, h has made an offer to d

⇙ d accepted

d was matched and stays matched ∴ contradiction

⇓ d rejected

d was matched and stays matched ∴ contradiction

GS Algorithm: Stable Matching

only type of instability b/c M is a perfect matching

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

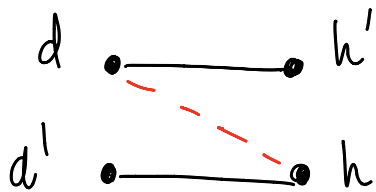

$(d, h') \in M$ $\quad d: h > h'$

$(d', h) \in M$ $\quad h: d > d'$

We'll derive the contradiction $d: h' > h$

- Because h prefers d, h made an offer to d before d'
- Case 1: d accepted

  Case 2: d rejected

GS Algorithm: Stable Matching

only type of instability b/c M is a perfect matching

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

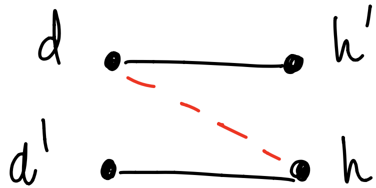

$(d, h') \in M$ $\quad$ $d: h > h'$

$(d', h) \in M$ $\quad$ $h: d > d'$

Case 1: d accepted

- At some point d broke off the match with h
- Because "doctors go up" $\quad$ $d: h' > h$

GS Algorithm: Stable Matching

only type of instability b/c M is a perfect matching

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

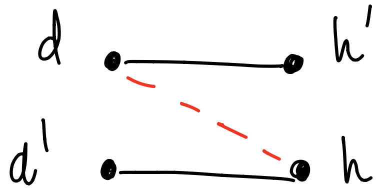

$(d, h') \in M$

$(d', h) \in M$

$d: h > h'$

$h: d > d'$

Case 2: d rejected

- The d was matched to some h'' s.t. $d: h'' > h$
- Because "doctors go up" $d: h' \geq h'' \geq h$

Contradiction.

GS Algorithm: Running Time

- Let M be empty
- While (some hospital h is unmatched):

  ≤ n^2 iterations

  - If (h has offered a job to everyone): break
  - Else: let d be the highest-ranked doctor to which h has not yet offered a job
  - h makes an offer to d:
    - If (d is unmatched):
      - d accepts, add (d,h) to M
    - ElseIf (d is matched to h′ & d: h′ > h):
      - d rejects, do nothing
    - ElseIf (d is matched to h′ & d: h > h′):
      - d accepts, remove (d,h′) from M and add (d,h) to M
- Output M

GS Algorithm: Running Time

(n² offers) × (n per offer)

- **Running Time:**
  - A straightforward implementation requires $\approx n^3$ operations, $\approx n^2$ space

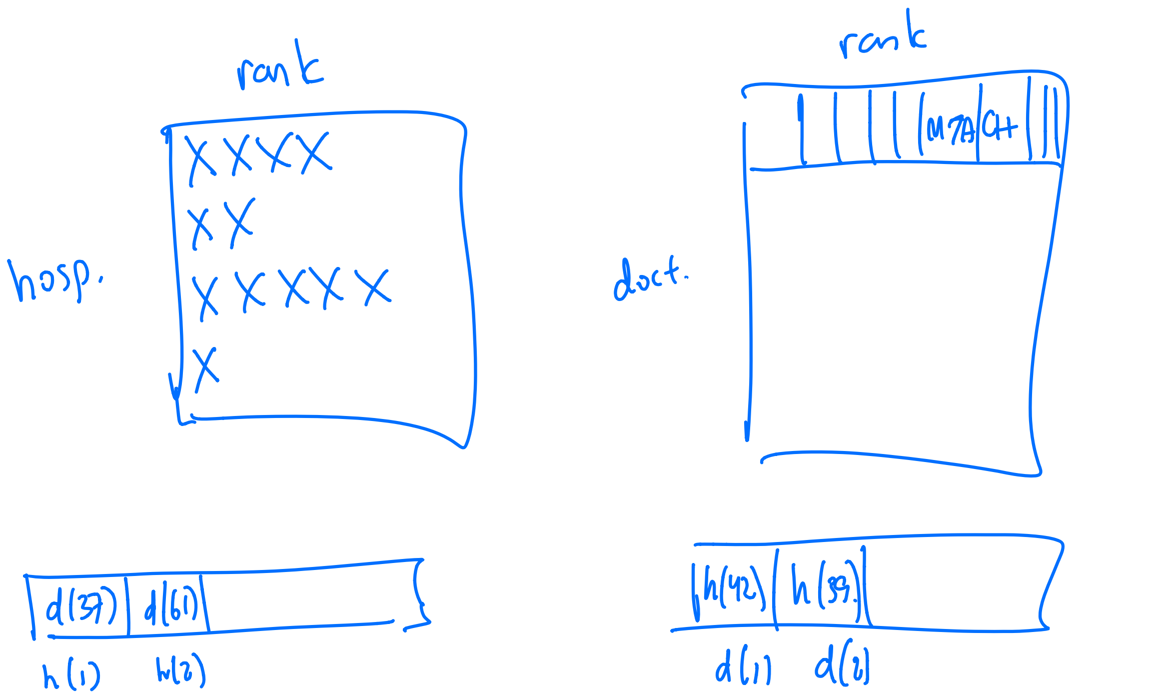

GS Algorithm: Running Time

- **Running Time:**
  - A careful implementation requires just $\approx n^2$ time and $\approx n^2$ space

GS Algorithm: Running Time

- **Running Time:**
  - A careful implementation requires just $\approx n^2$ time and $\approx n^2$ space

Can convert from doc × rank → doc × hosp in n^2 ops

	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

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

GS Algorithm: Running Time

- **Running Time:**
  - A careful implementation requires just $\approx n^2$ time and $\approx n^2$ space

① Convert the doctors' preferences n^2 ops

② Run GS $(n^2$ offers$) \times (1$ operation$)$ n^2 ops

$\approx n^2$ operations

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

2012 Nobel Prize to Lloyd Shapley and Al Roth!