CSCI 357: Algorithmic Game Theory

Lecture 11: Top Trading Cycles

Shikha Singh

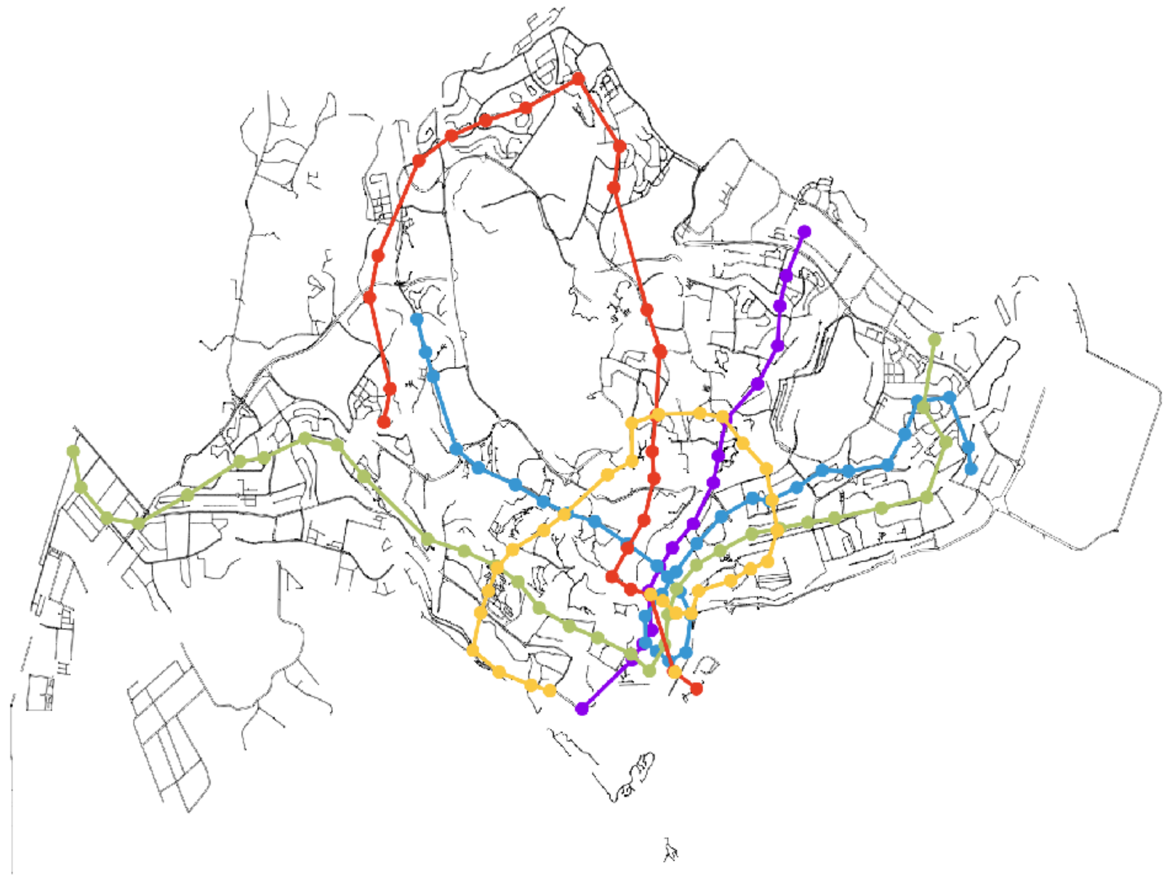

Announcements and Logistics

- **HW 5** is due April 8 in class
- **Assignment 3** will be on matching algorithms and due **April 11 at noon**
- **Hope you have a good break!!**
- Any questions?

Top Trading Cycle for Exchange Markets

Housing Exchange Market

- n agents and n houses, each agent has a strict preference over the n houses
- Suppose each agent already owns one of the n houses
- Agents are willing to exchange with others to get a better one
- **Goal.** A way to reassign items to agents (perform exchanges) st.:
  - No one gets a house they like worse than the one they started with
  - Outcome is Pareto optimal
  - Strategyproof: truthful reporting of preferences is a dominant strategy
  - Stable / core allocation: no subset of agents can exchange amongst themselves to get a better outcome
- Sometimes called the house allocation problem

Example Instance

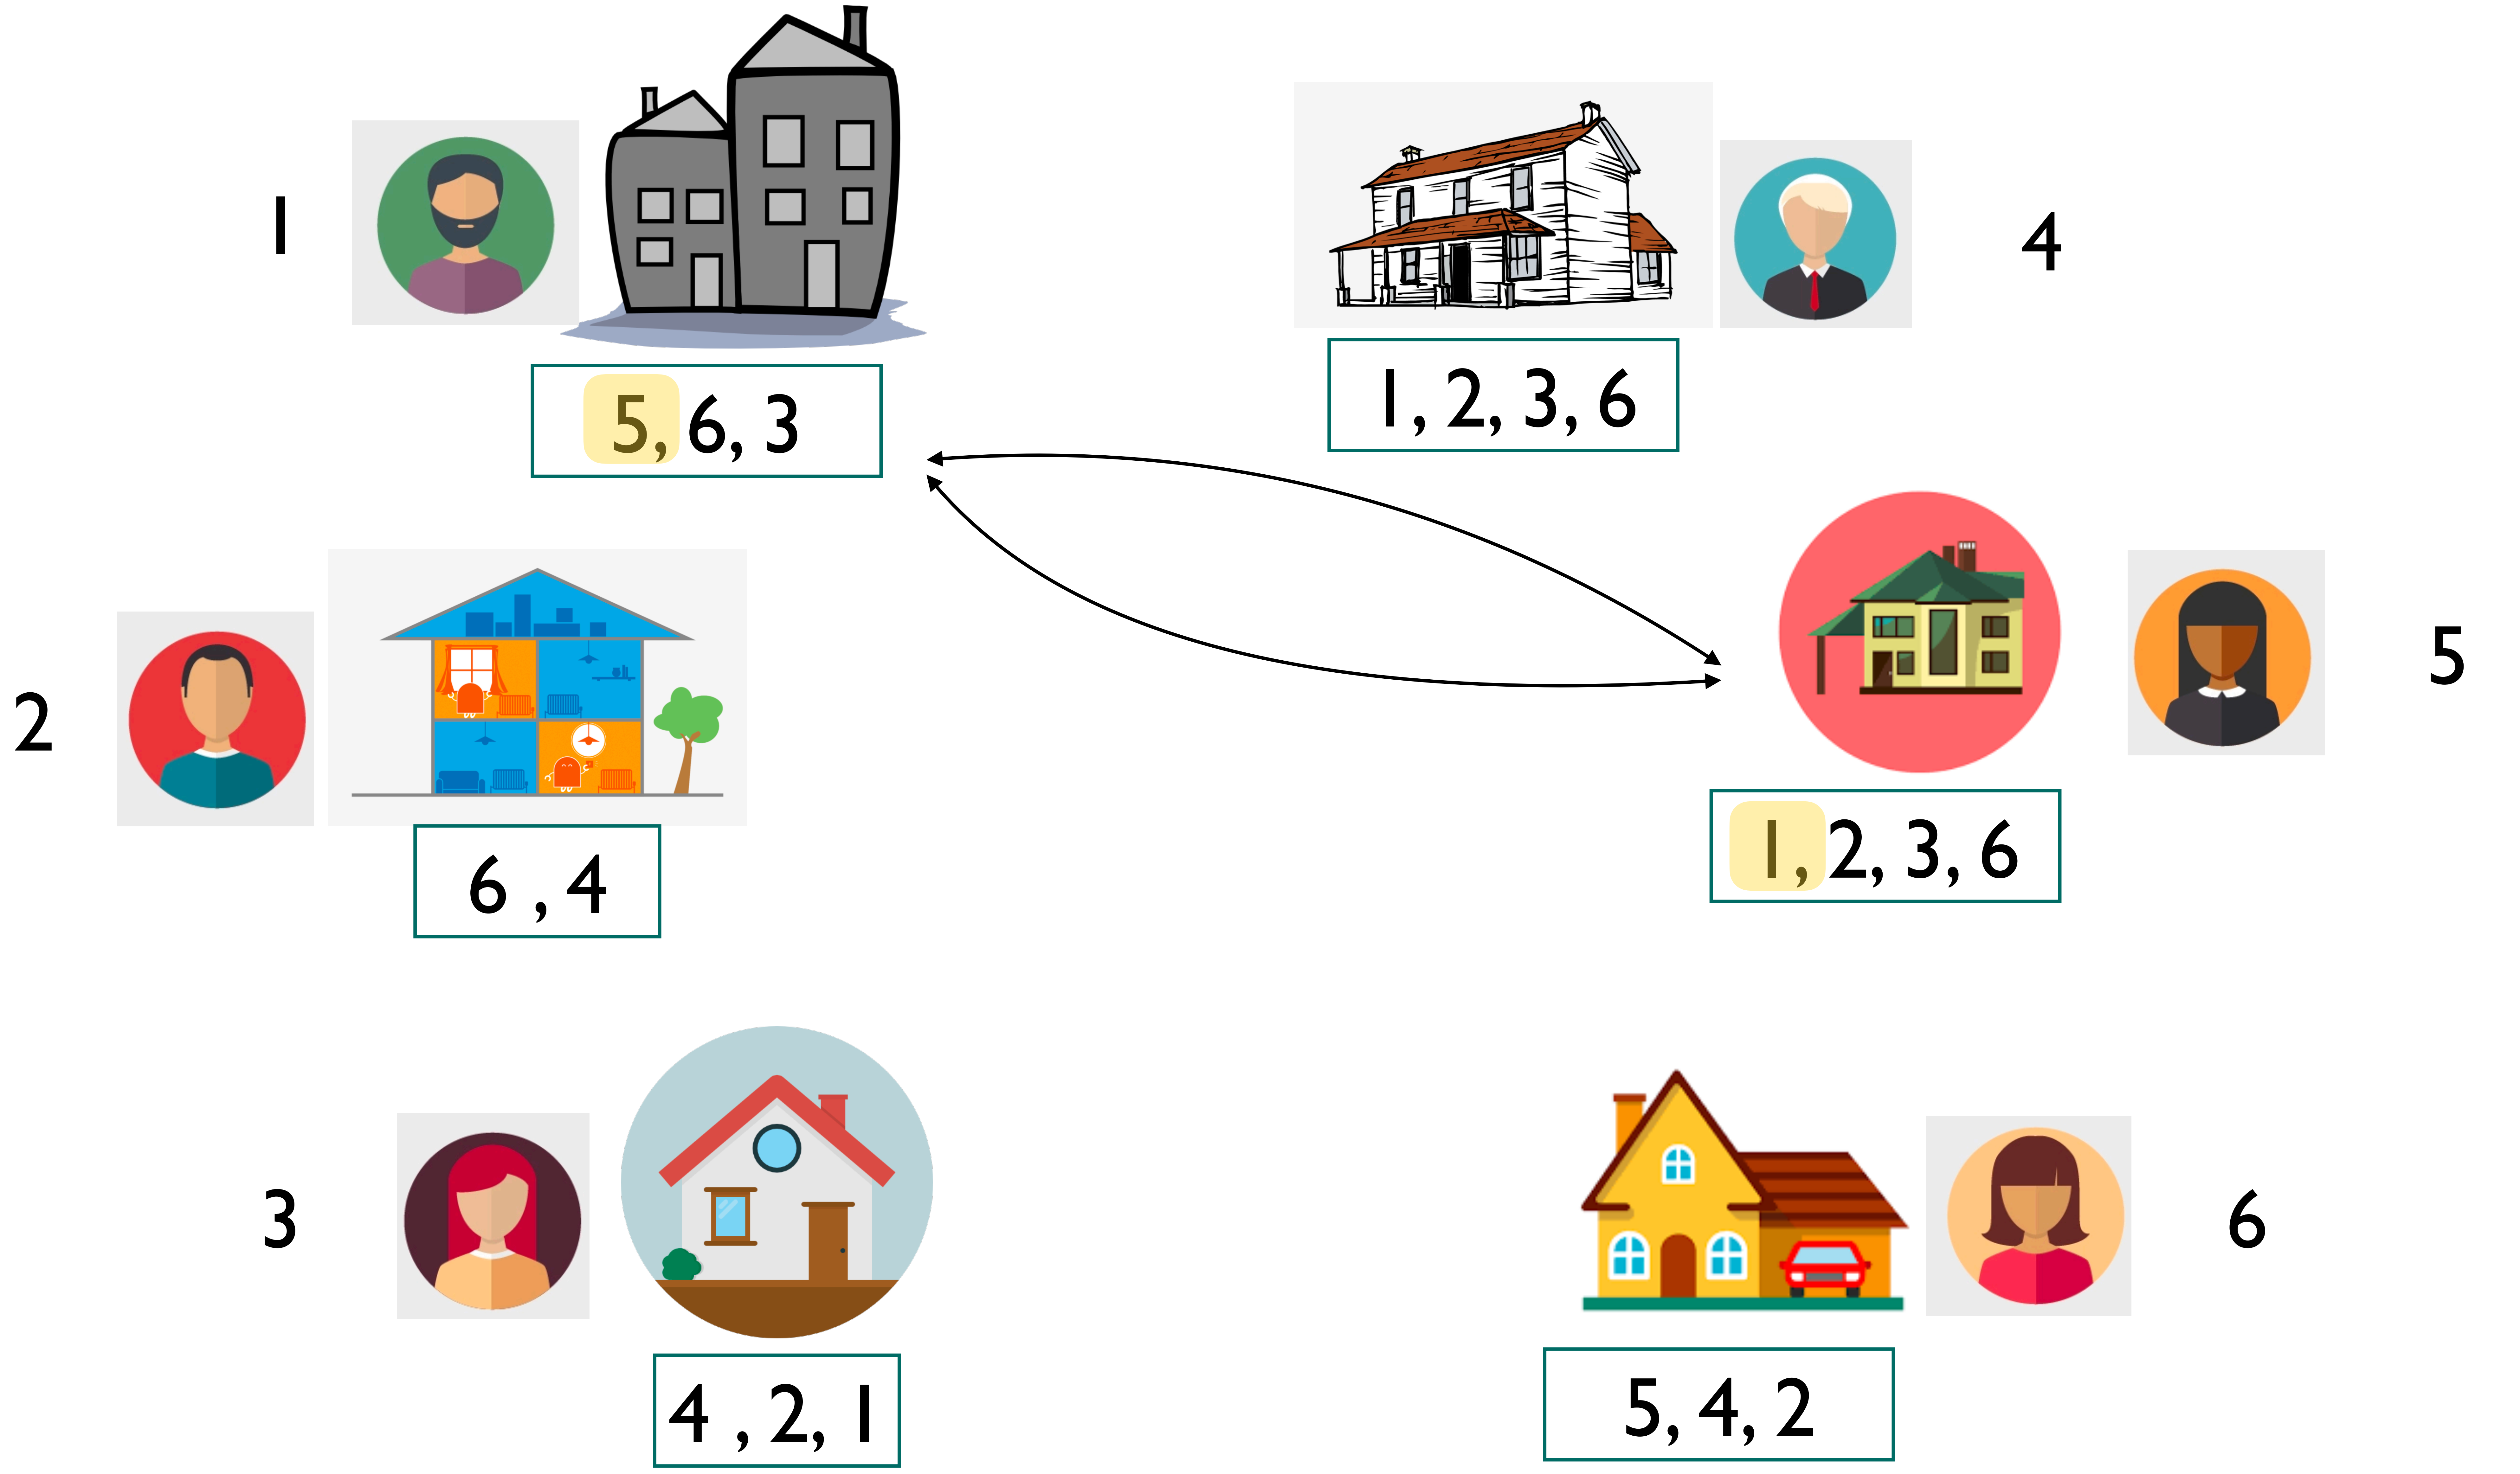

Top-Trading Cycle [Gale & Shapley]

- Each agent report their overall preferences in the beginning
- **Step 1.** Each agent (simultaneously) points to its favorite house (among houses remaining)
  - Induces a directed graph G in which every vertex has outdegree 1
  - G must have at least 1 directed cycle (self loops count)
  - Pick directed cycles and make all trades on it (each agent gives its house to the agent that points to it)
  - Delete all agents and houses that were traded in Step 1
- While agents remain, go back to Step 1.

Why is there at least one directed cycle?

Can an agent be involved in two directed cycles?

Example Instance

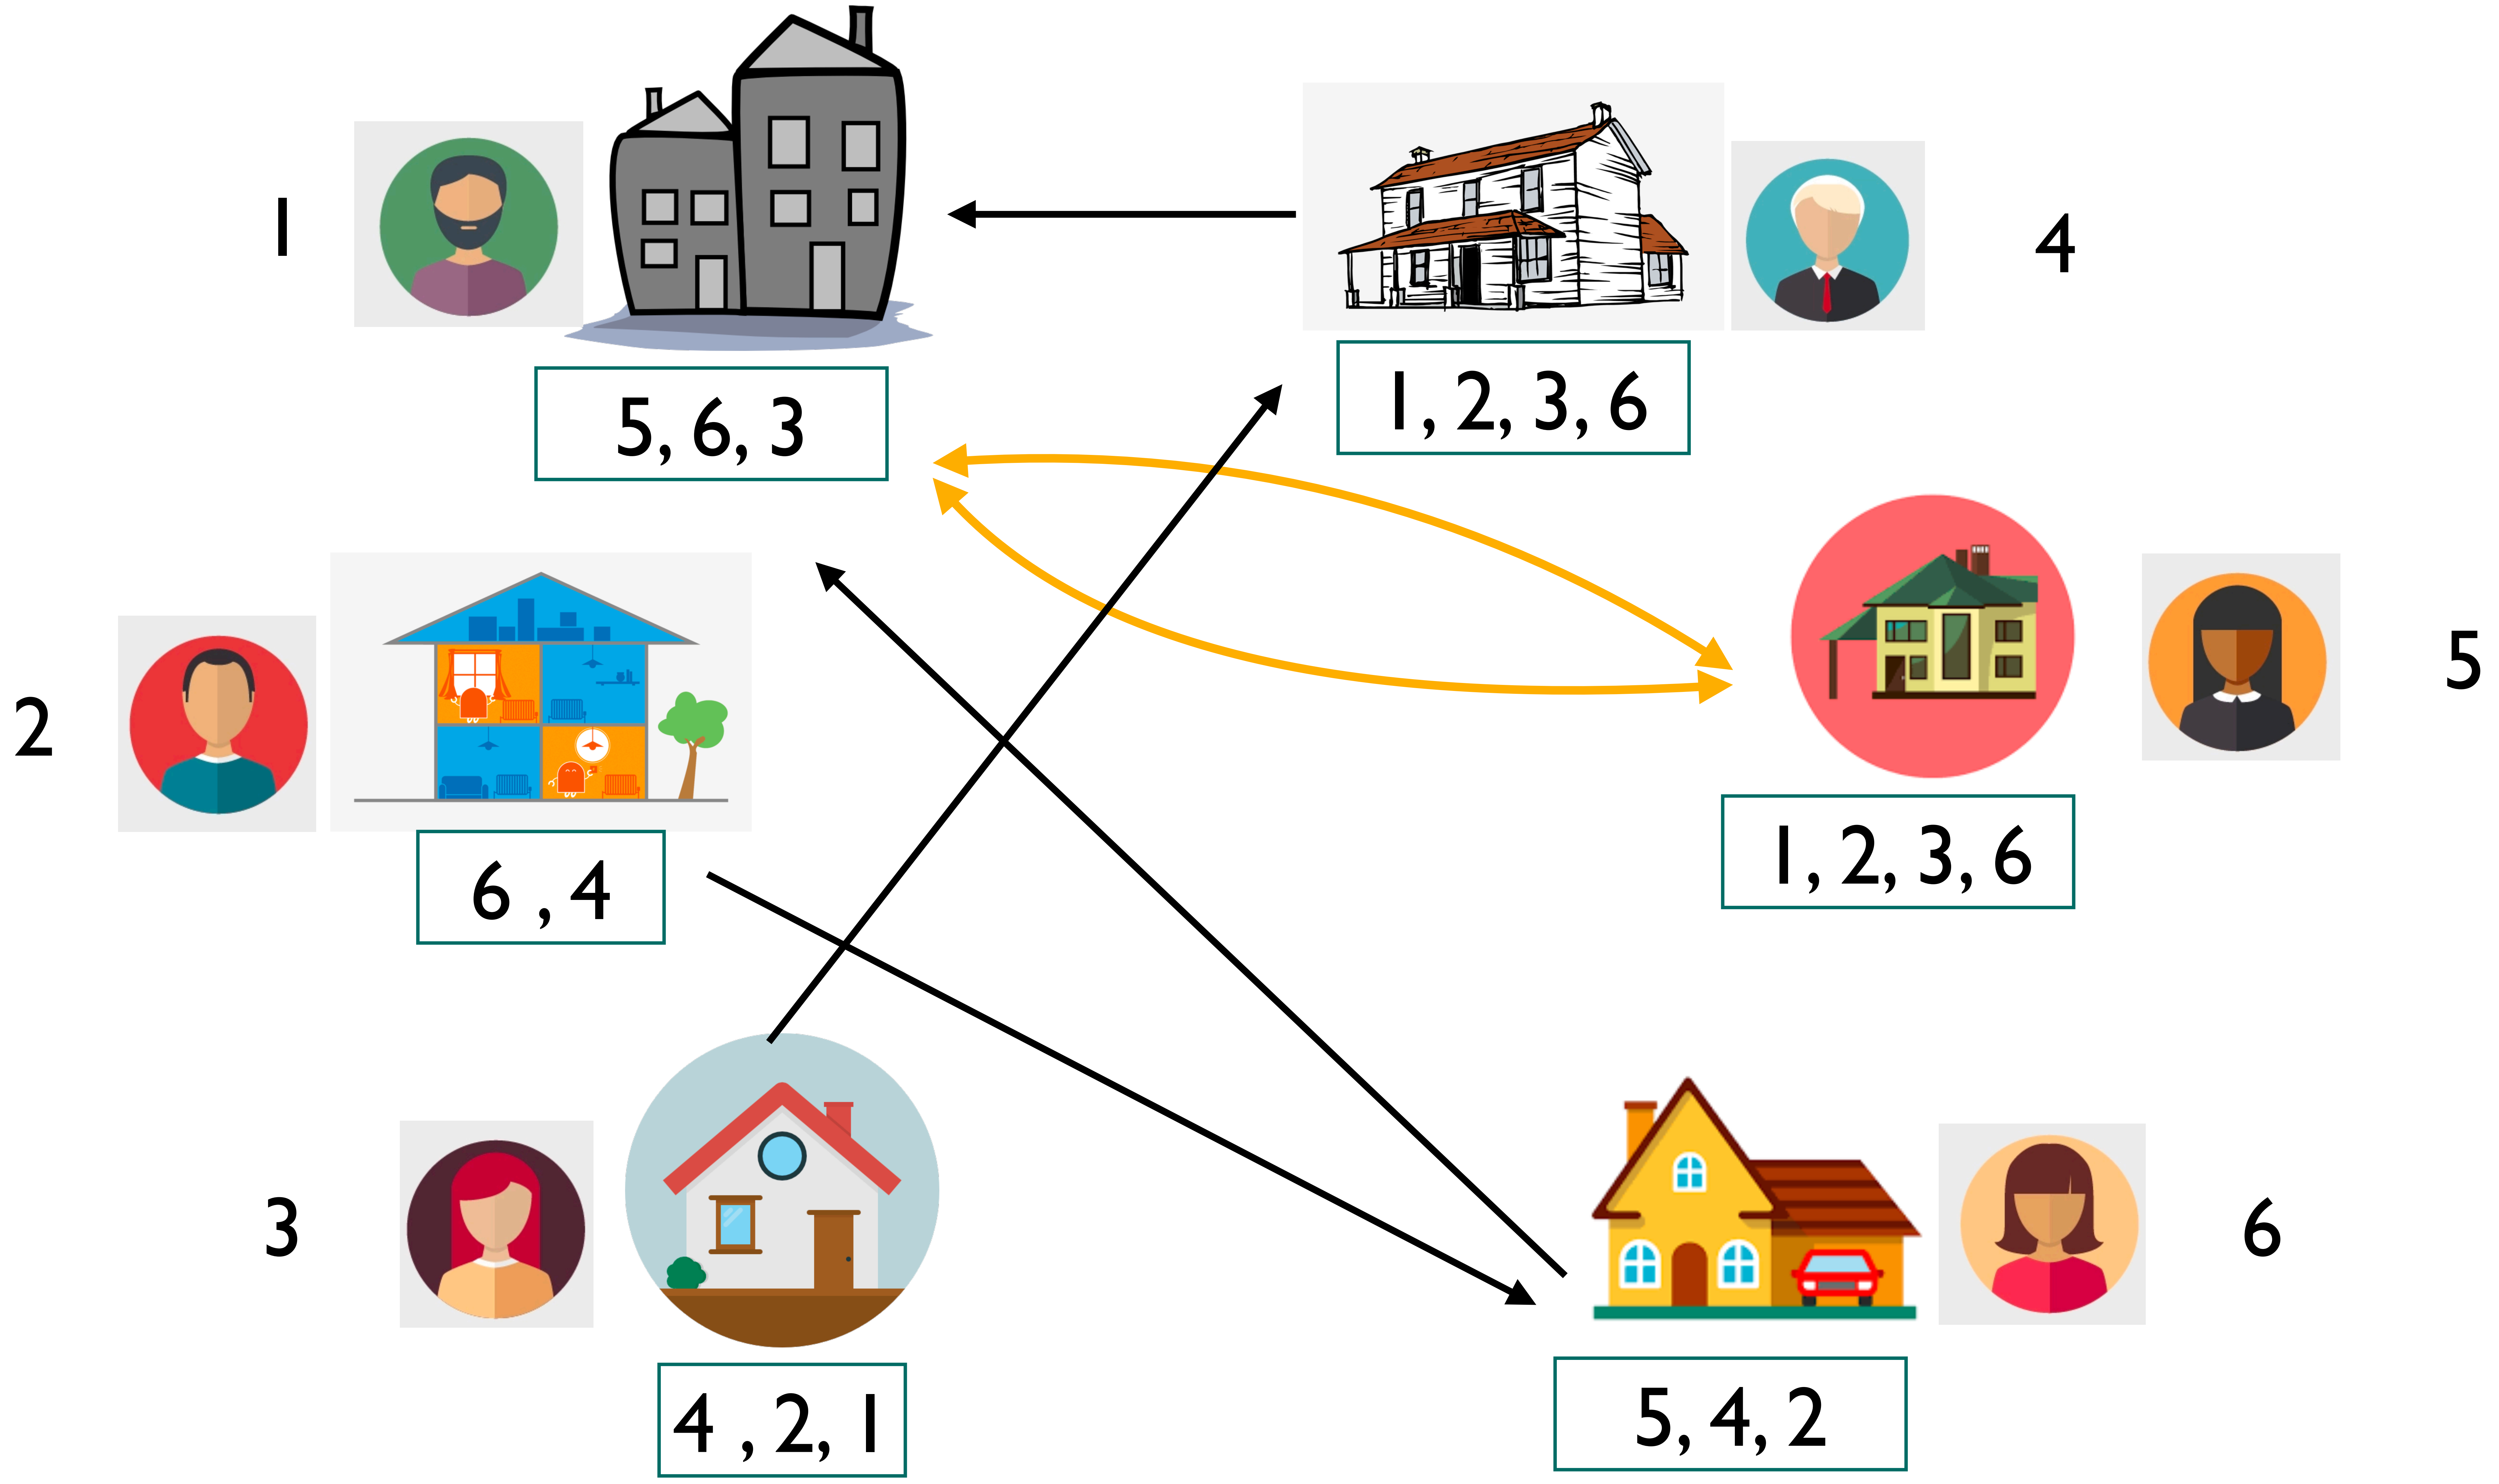

Example Instance

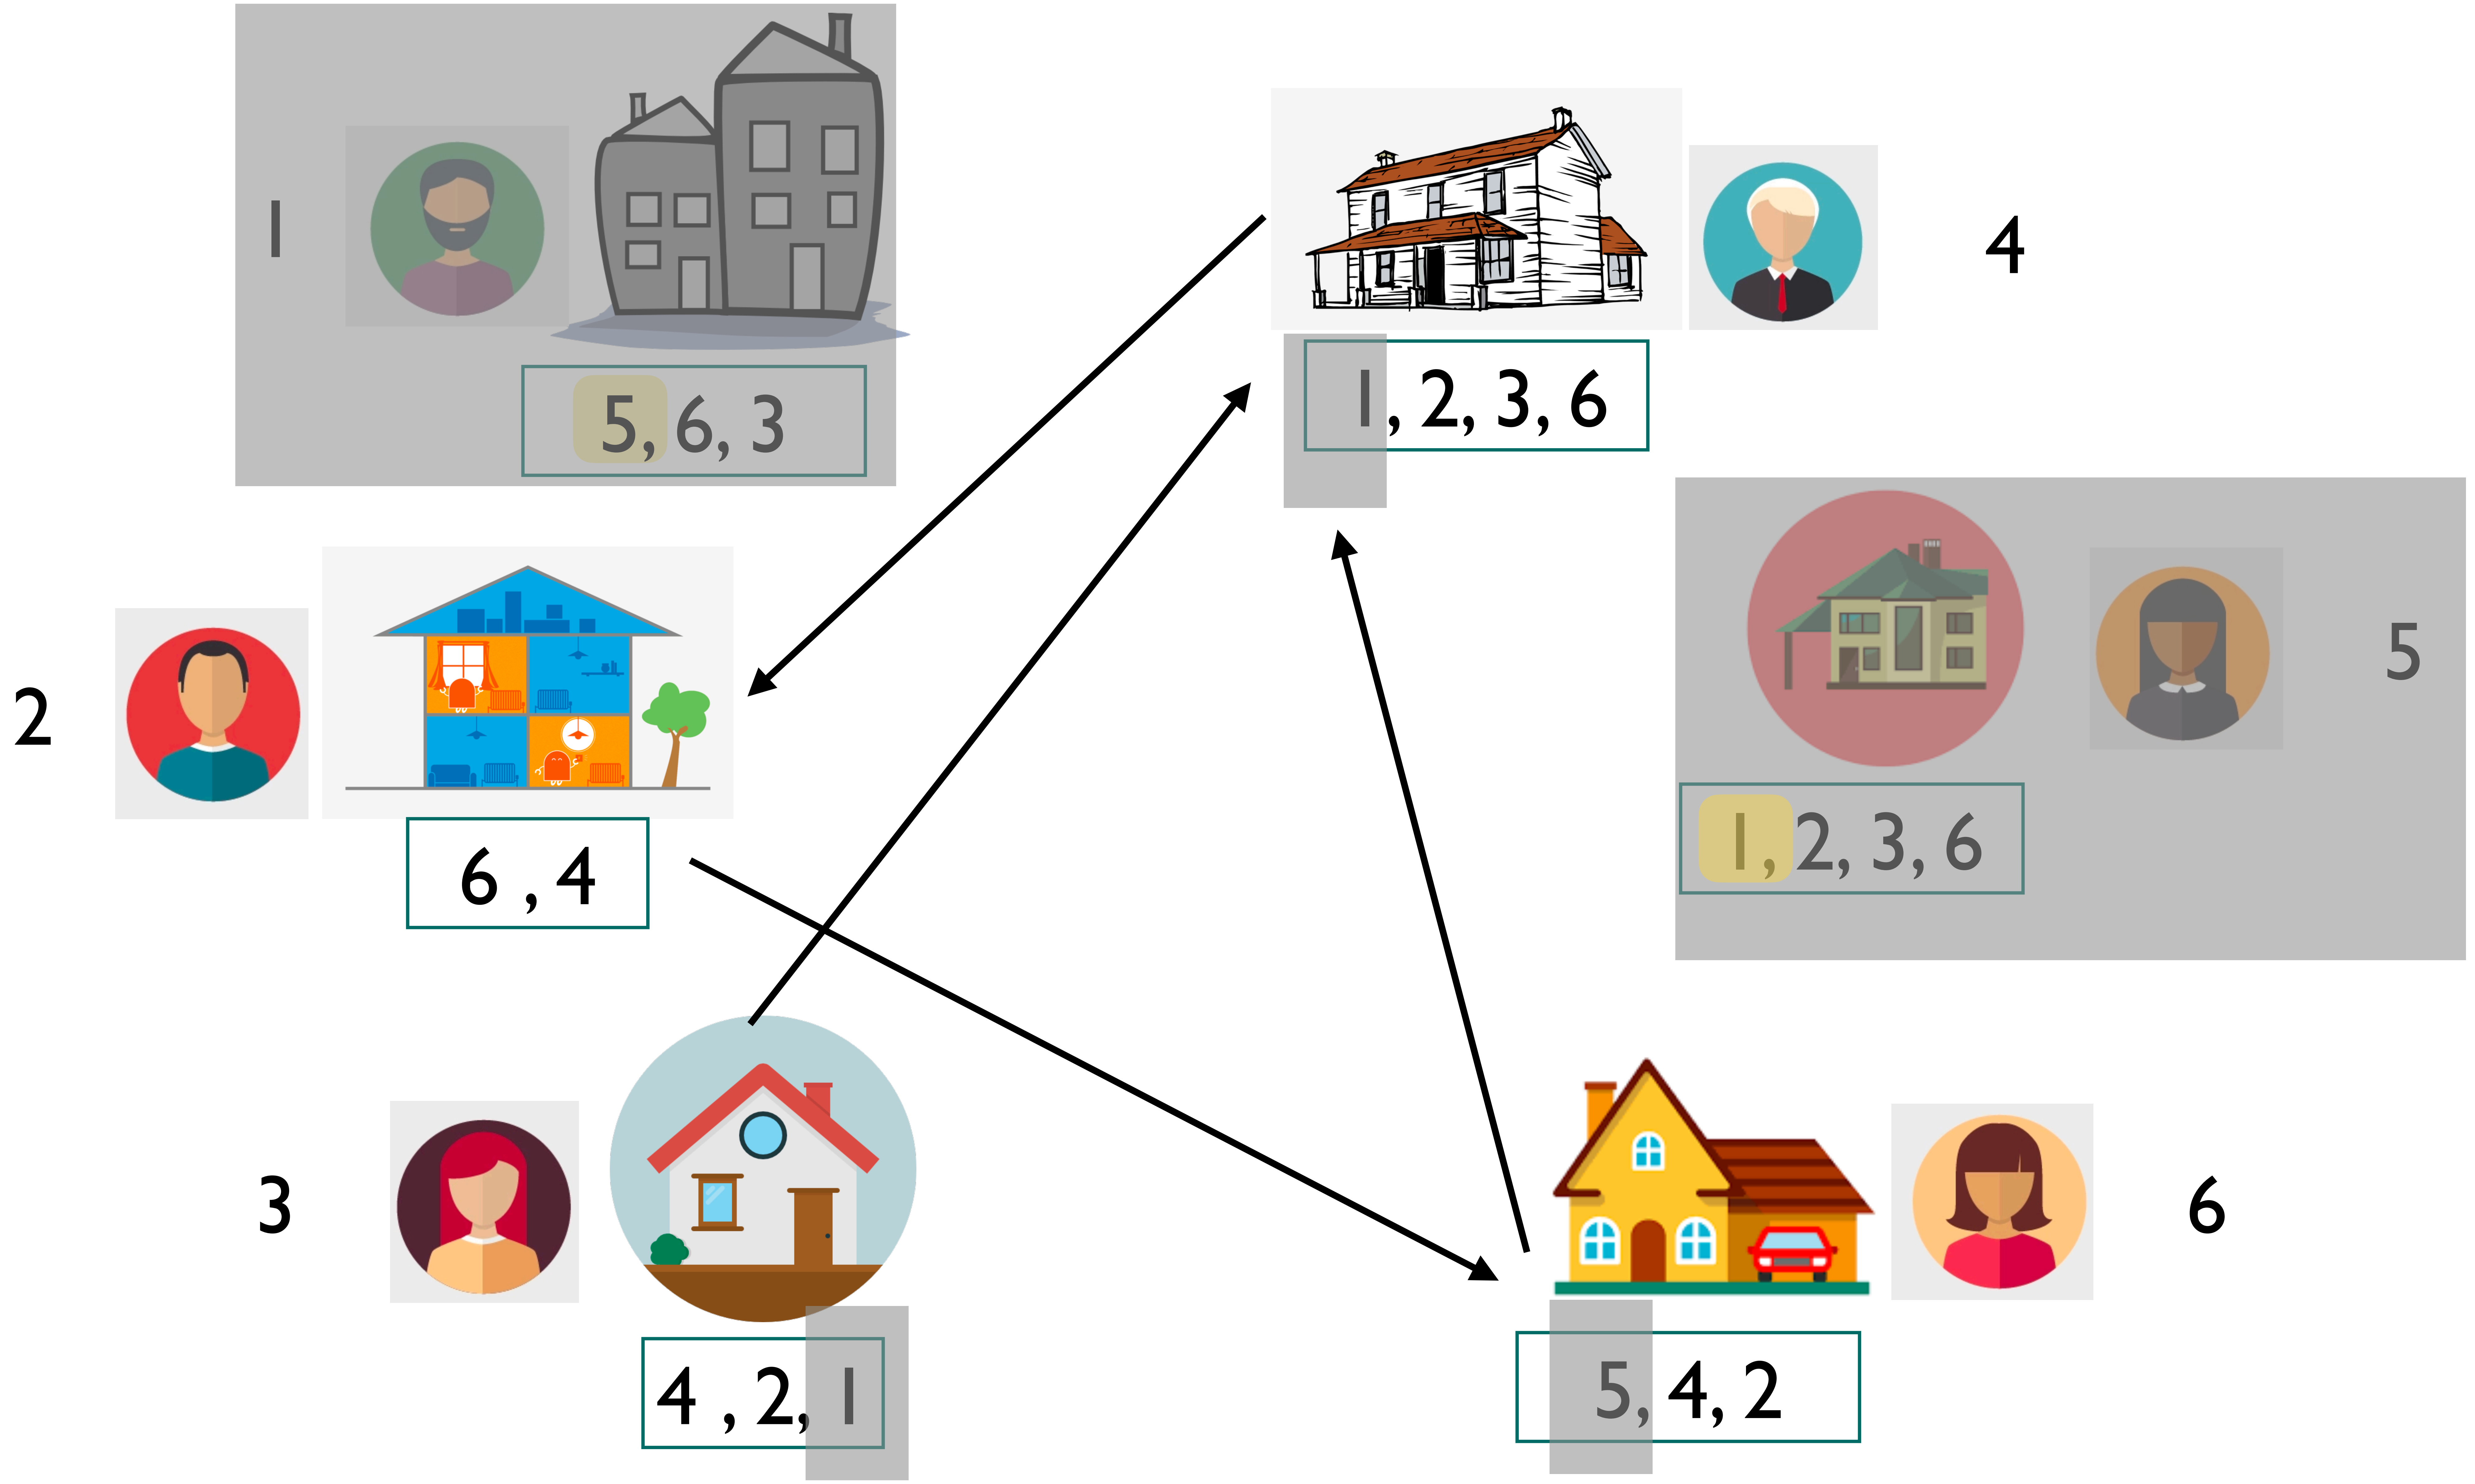

Example Instance

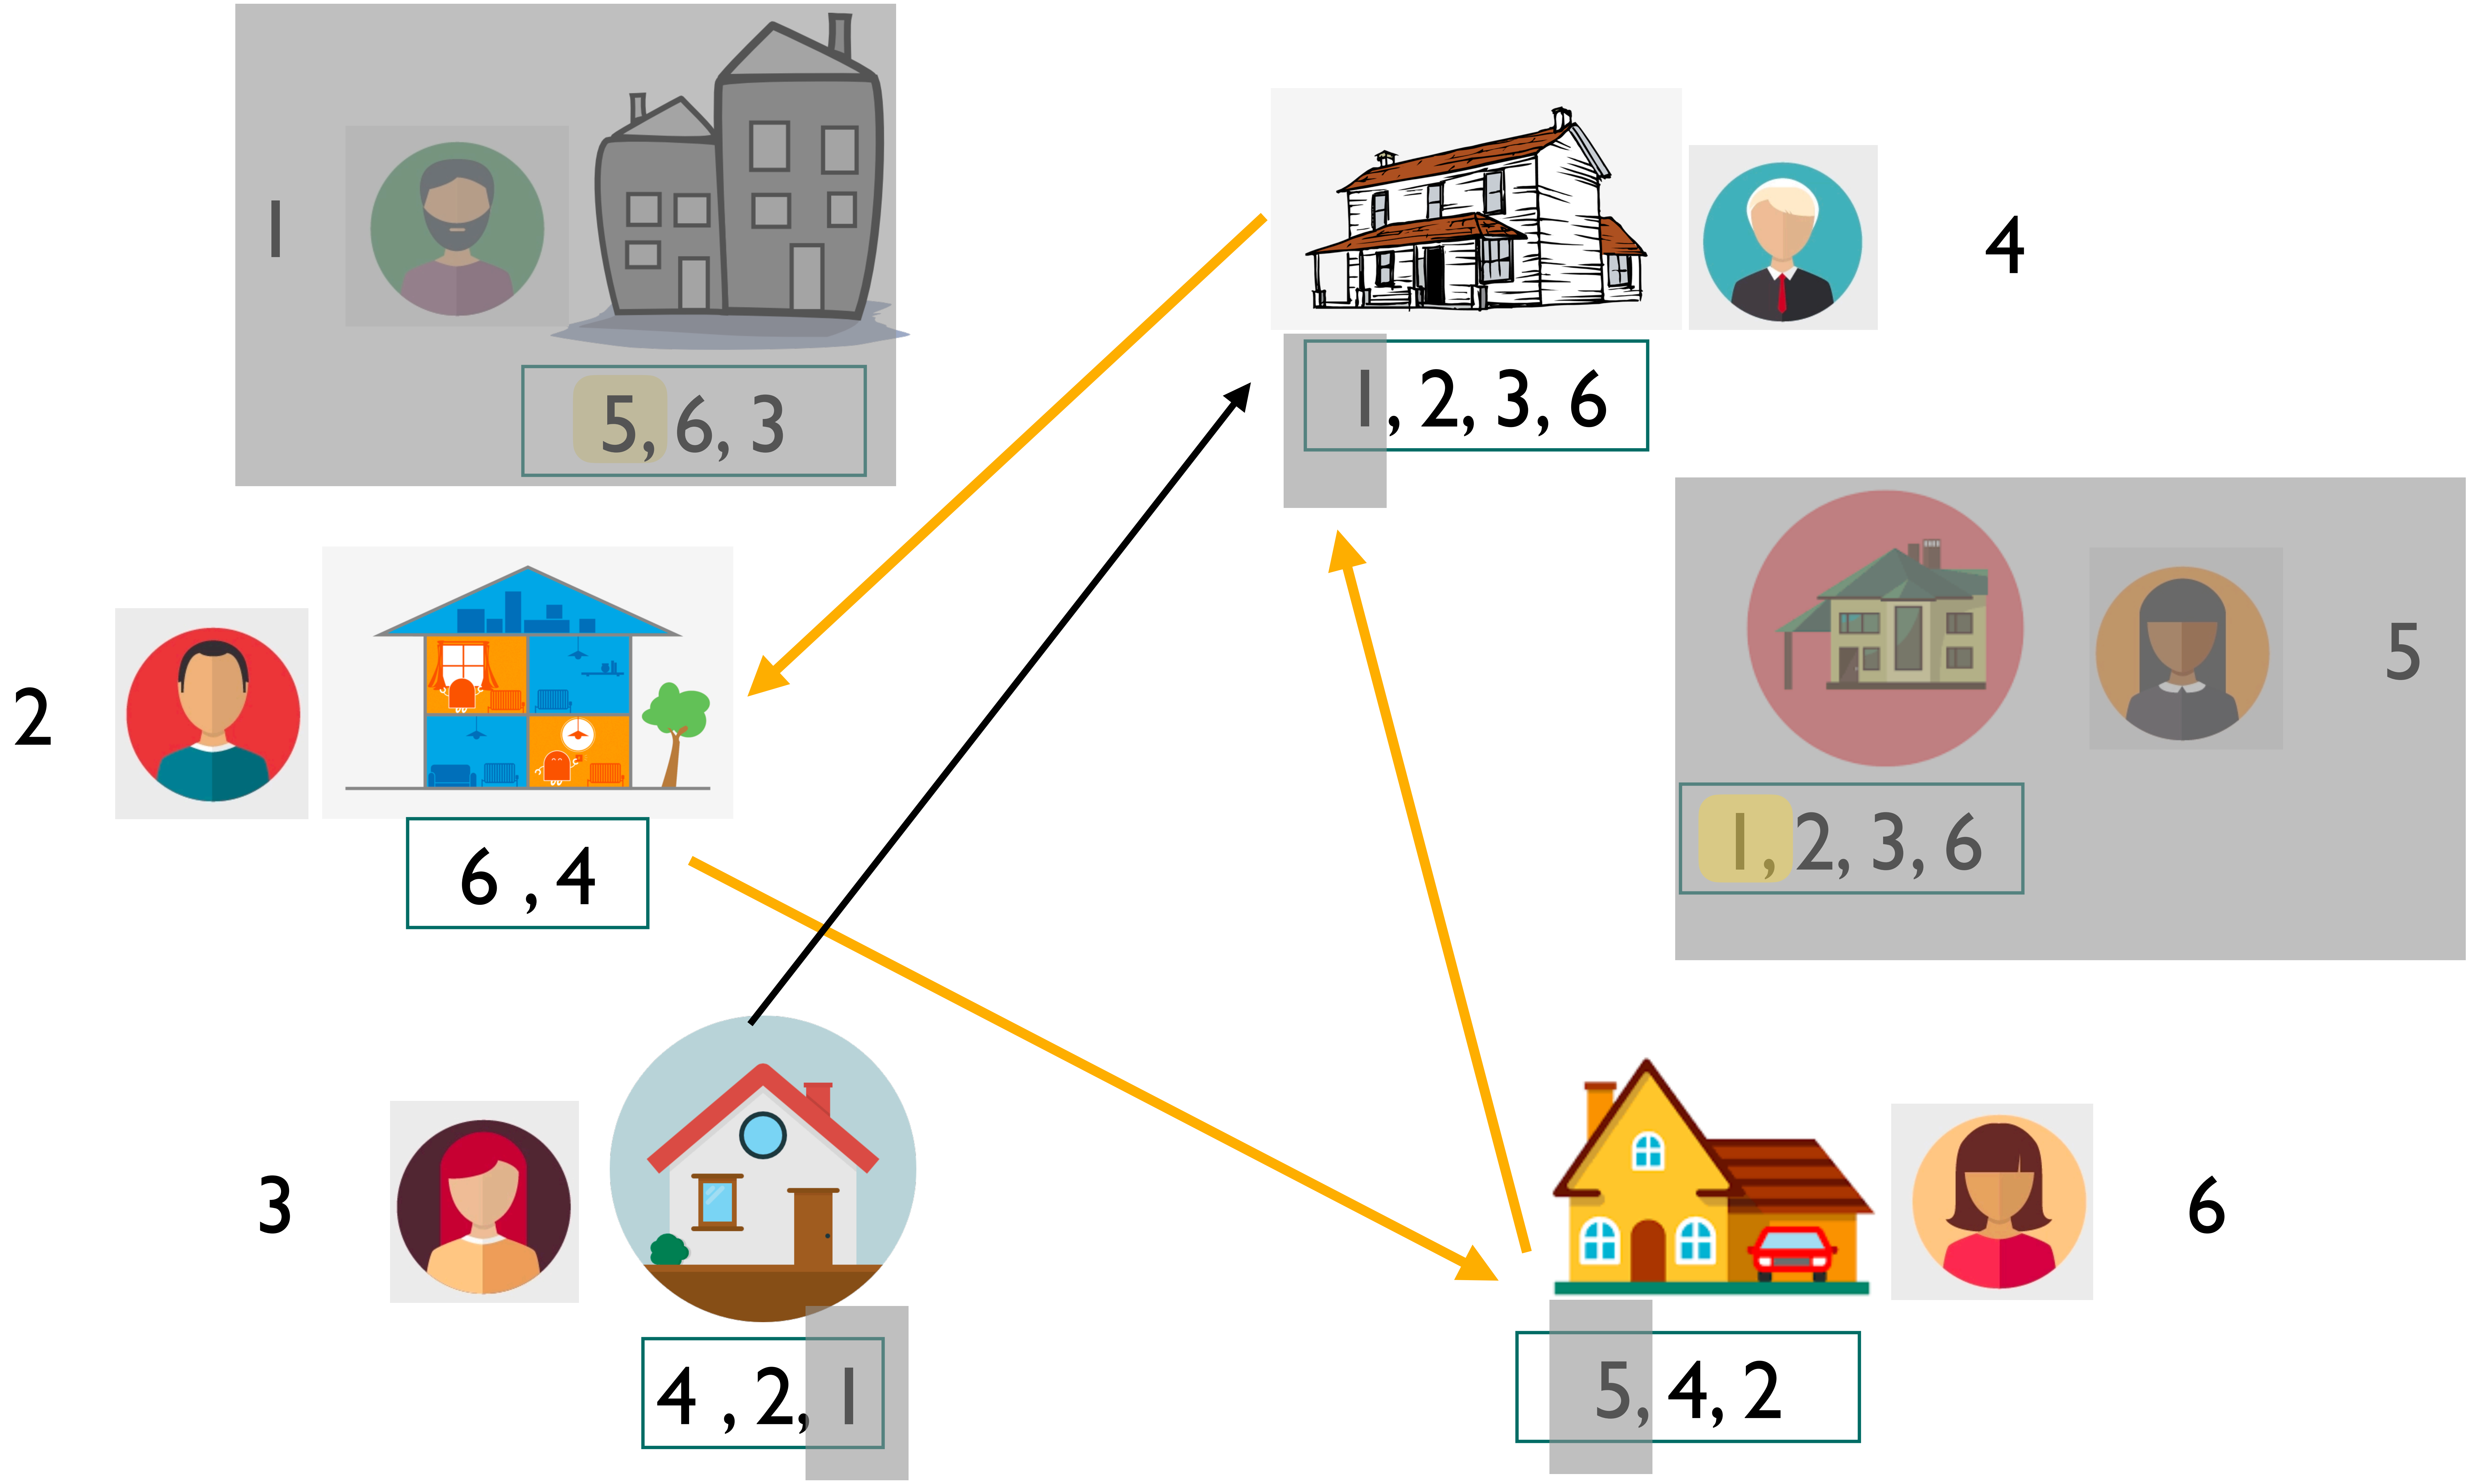

Example Instance

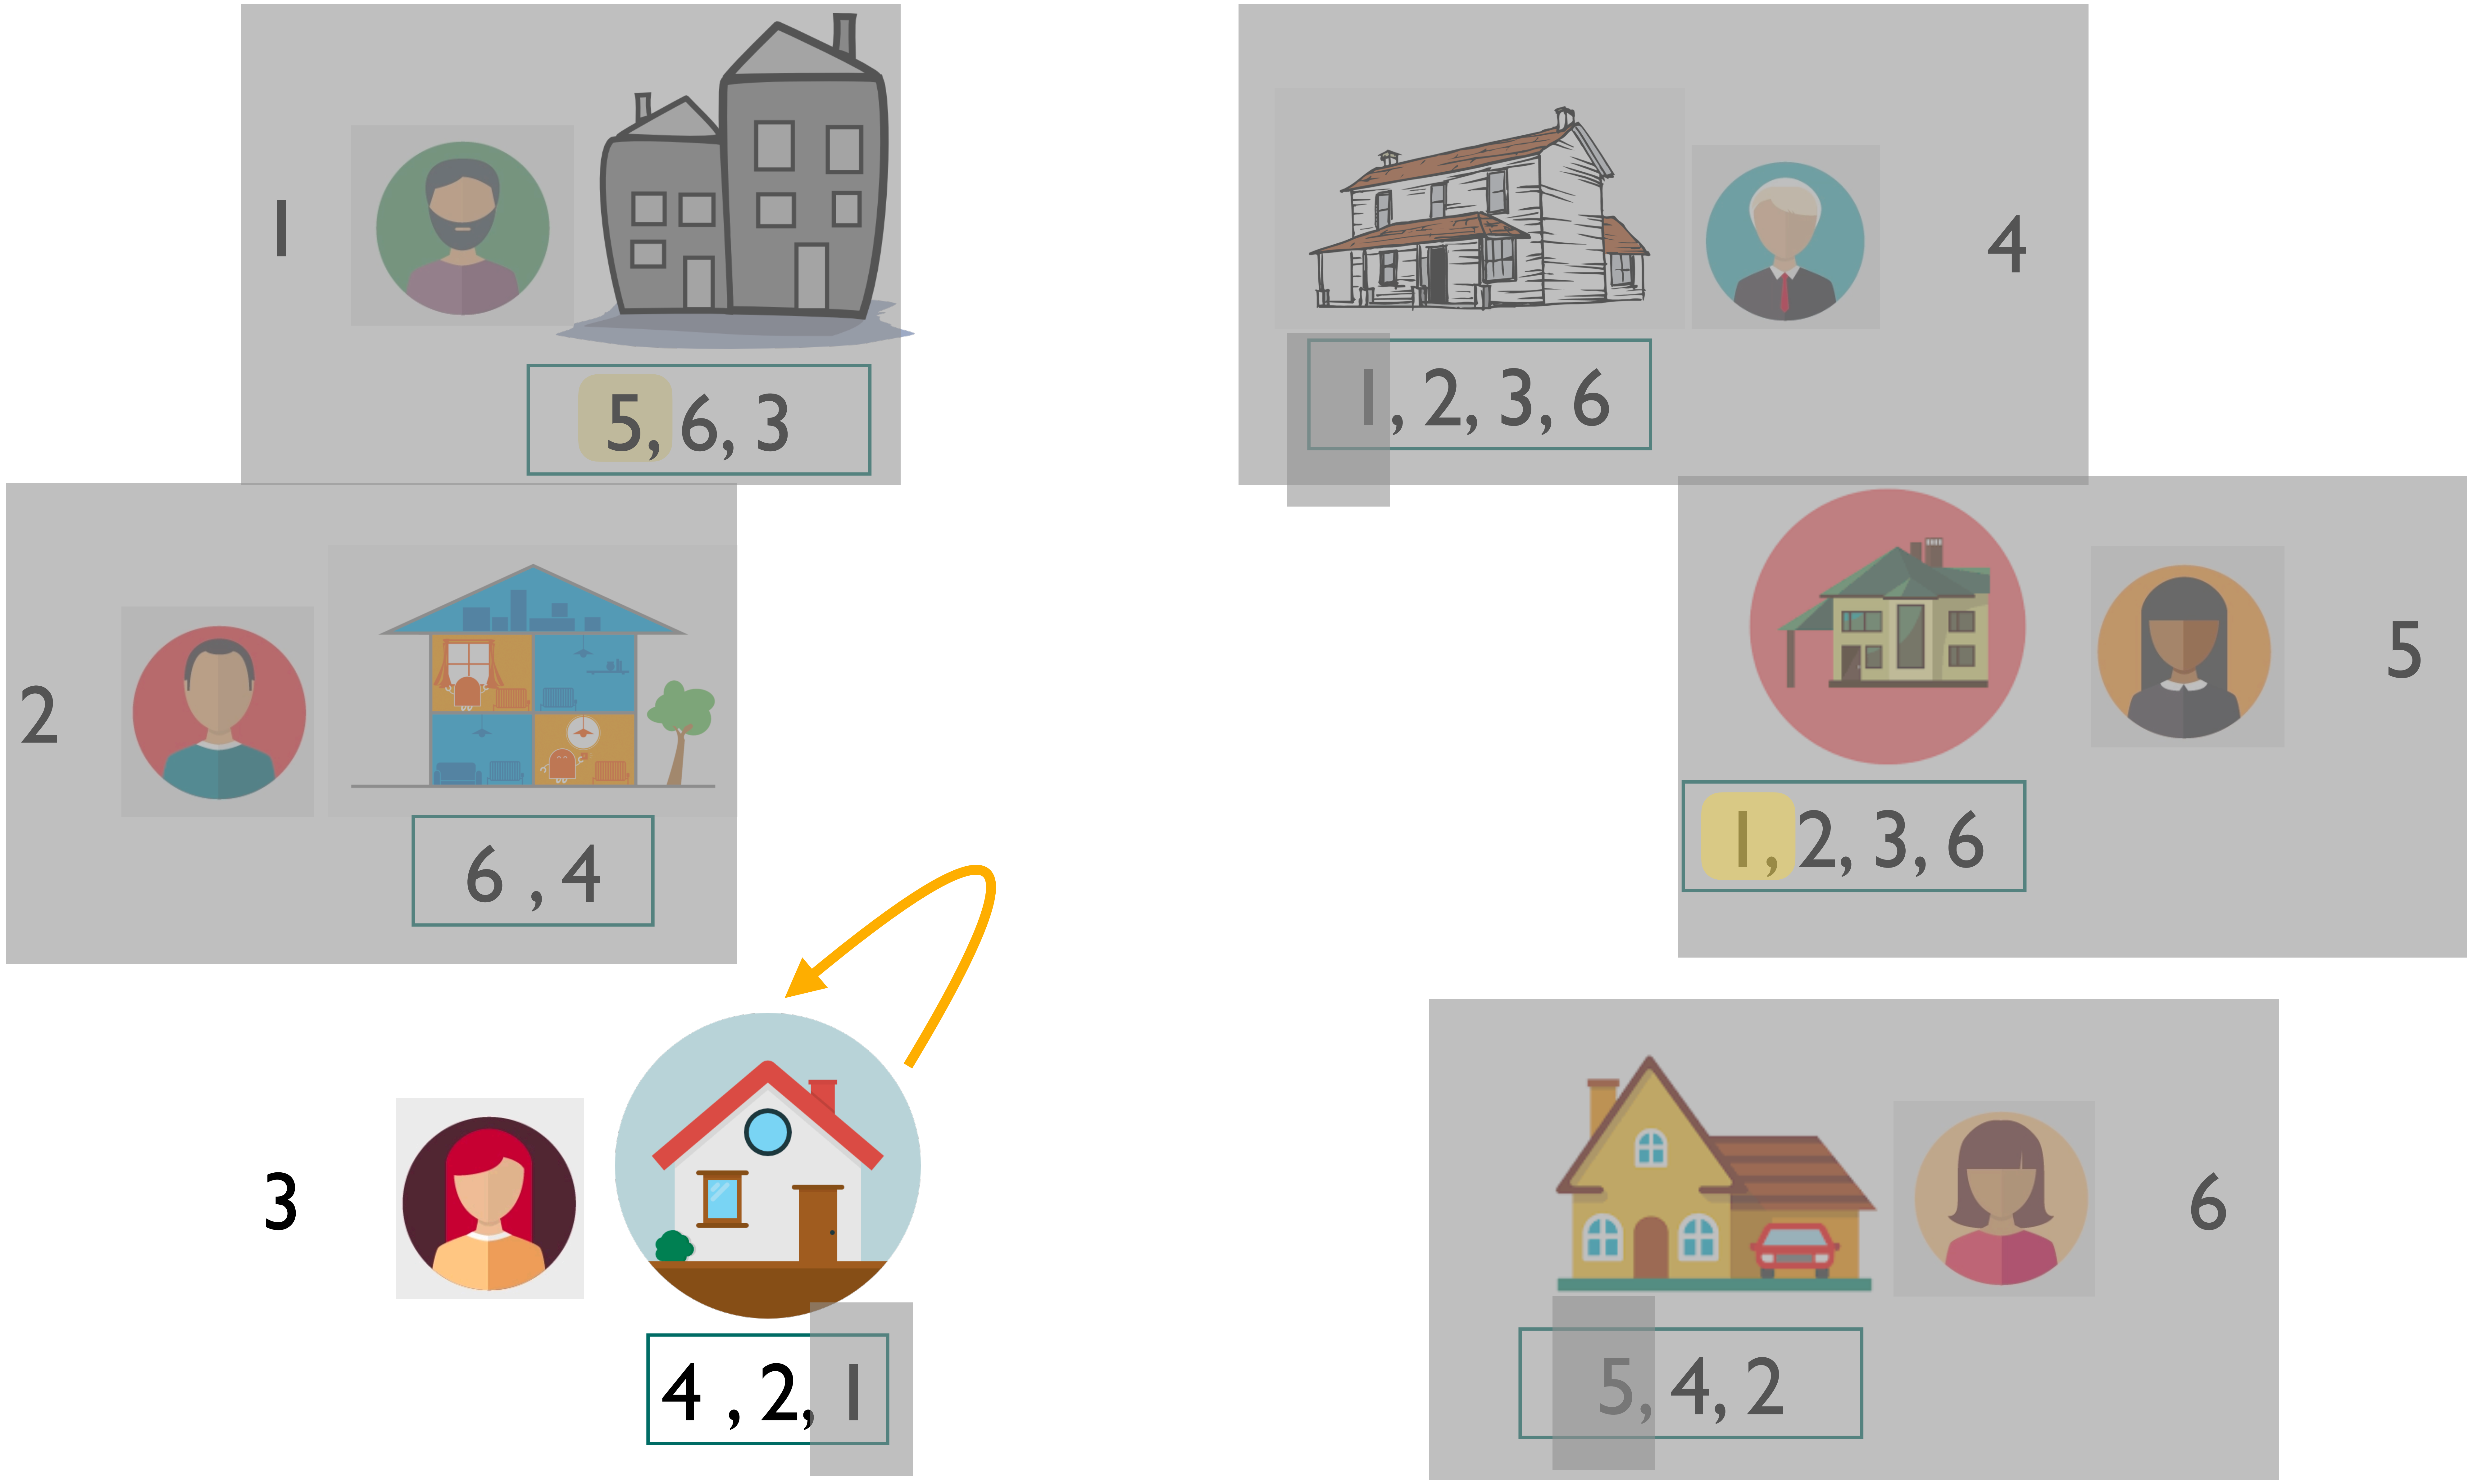

Final Output

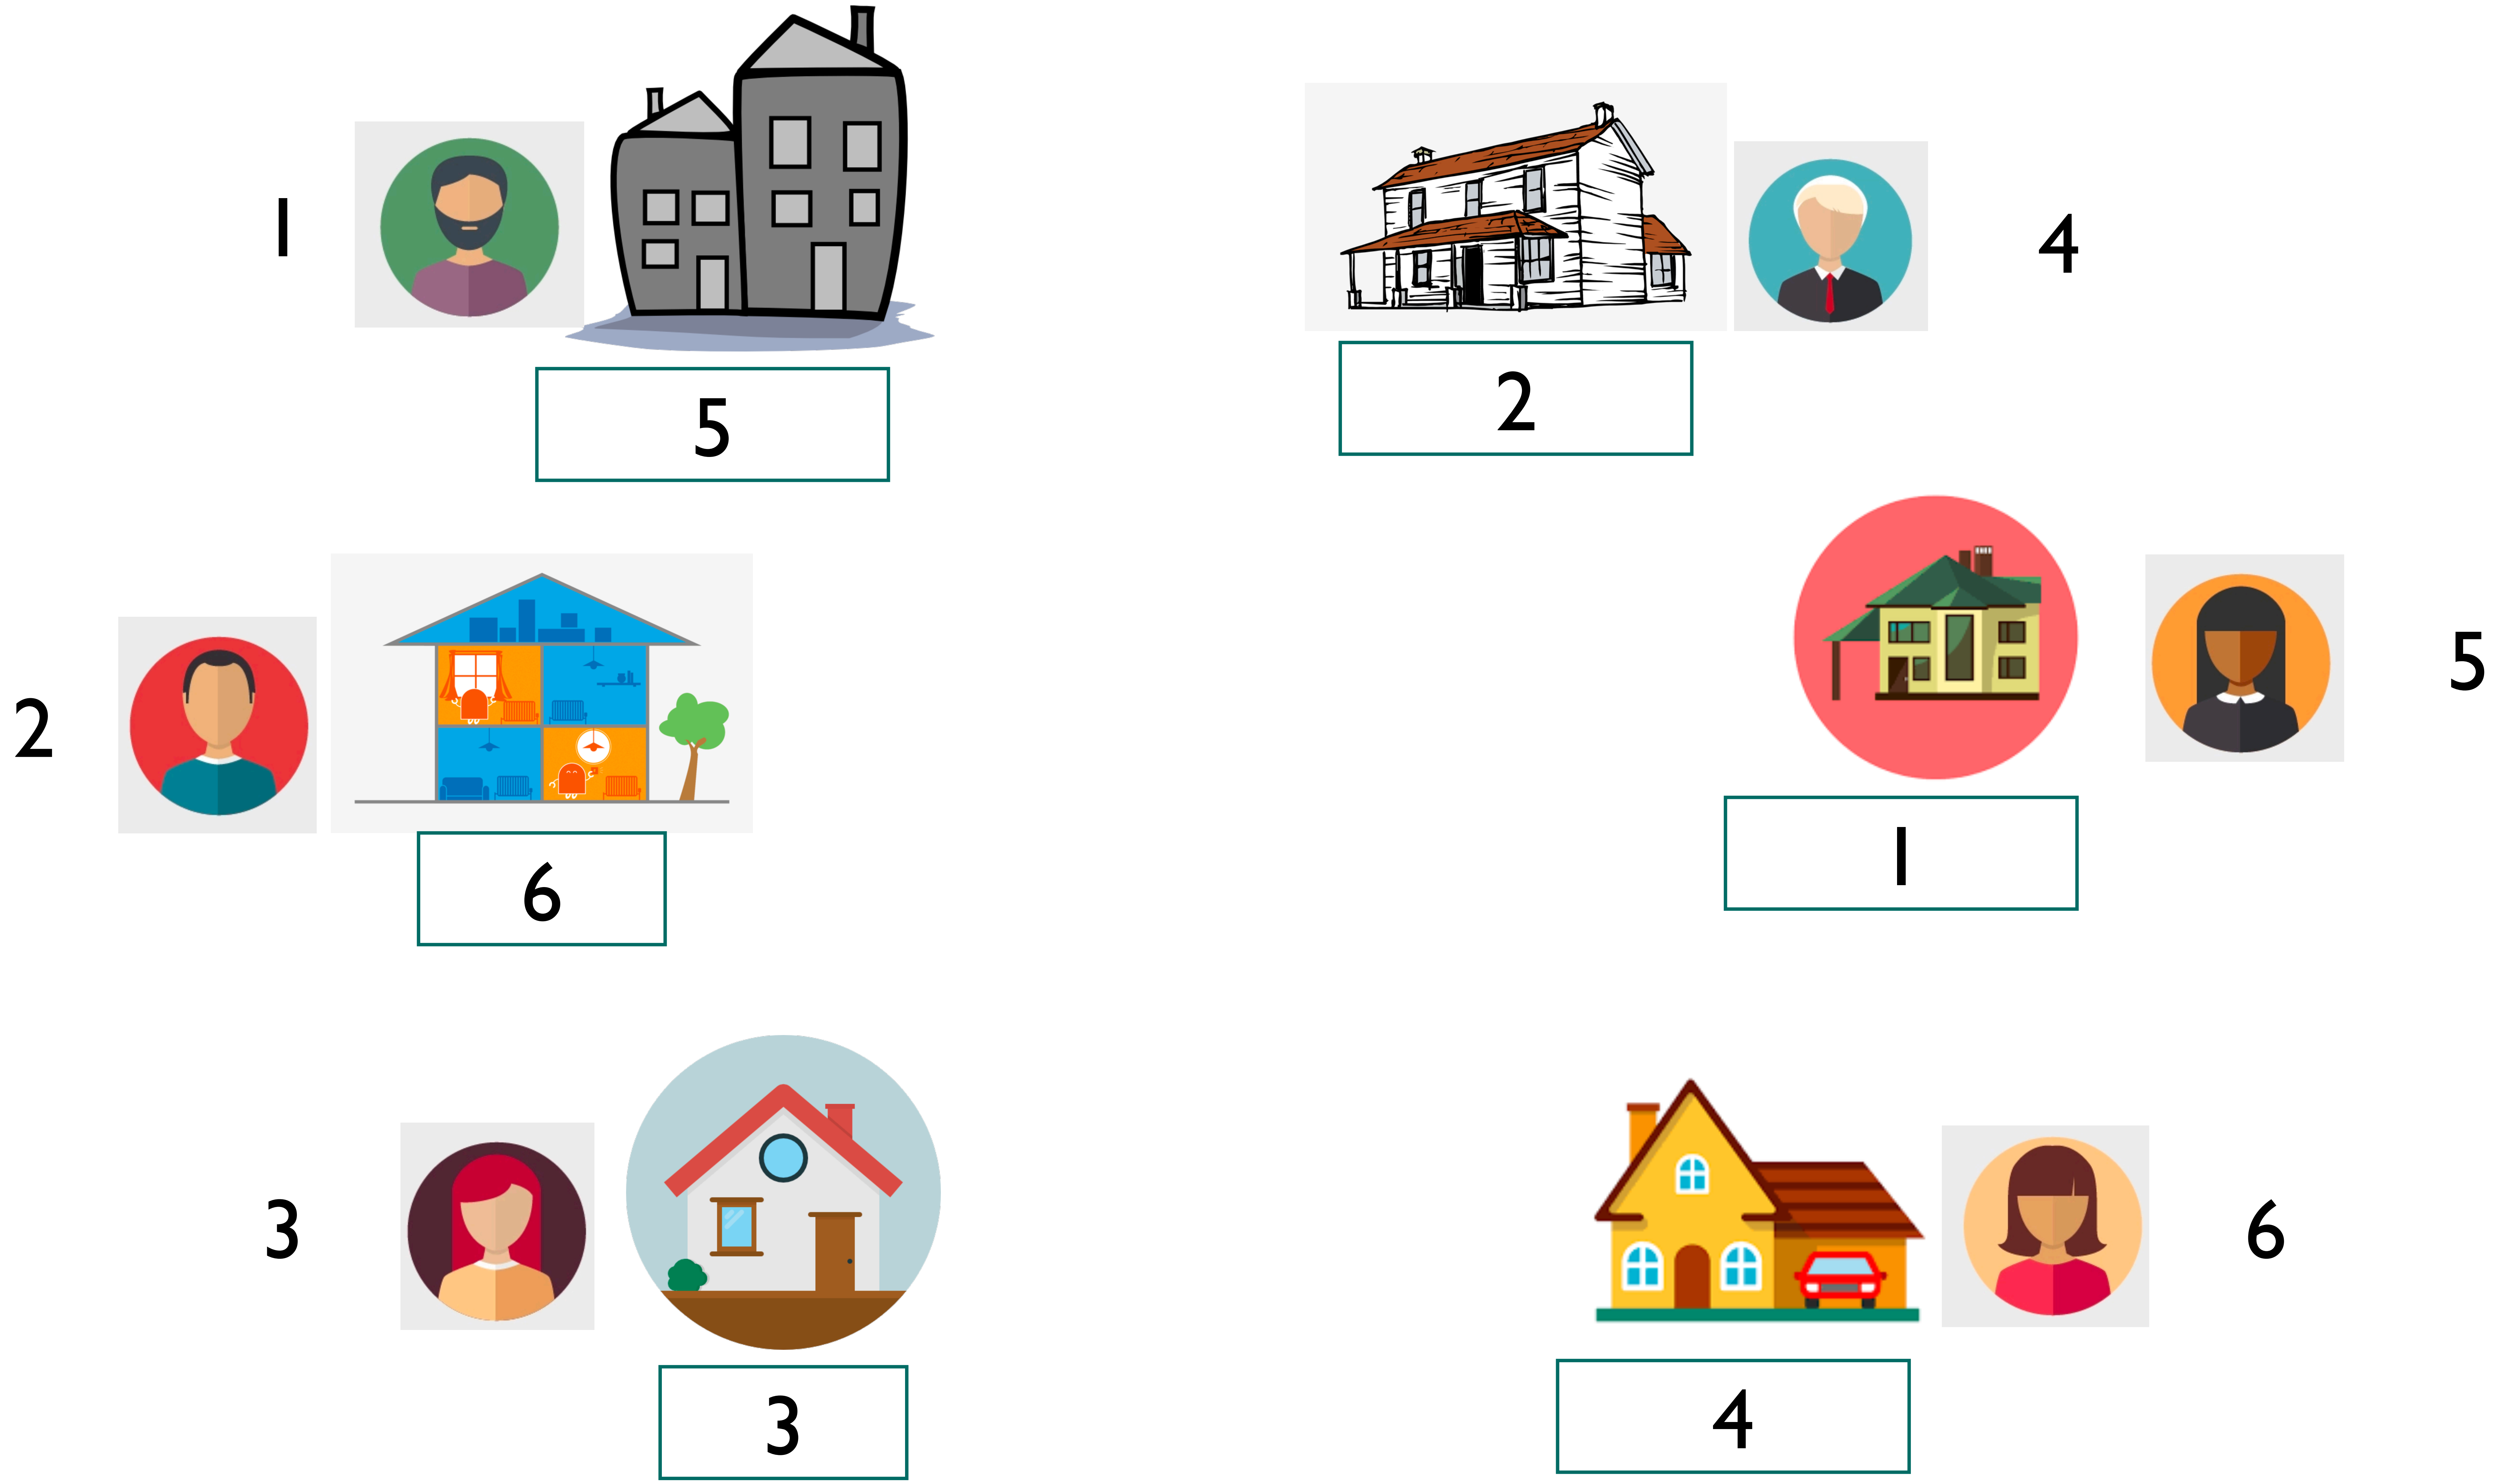

TTC Properties

- **Time Complexity.** How many rounds until the algorithm terminates?
  - At least one agent removed in each round, at most n rounds
  - Can show that each round can be implemented in $O(n)$ time
- **Incentive to participate**:
  - Allocation at least as good as the one they started with, why?
  - Everyone has their own house at the end of any preference ordering

TTC Invariant. Let N_k be the set of agents removed in the kth iteration of the TTC algorithm. Every agent of N_k receives their favorite house outside of the houses owned by $N_1 \cup N_2 \cup N_{k-1}$ and the original owner of any house allocated in this round is also in N_k.

TTC is Strategyproof

- **Proof.** Fix an agent i and preferences reported by others.
- Define the sets N_k as in the TTC invariant. Suppose $i \in N_j$ when i is truthful.
- **Lemma.** Regardless of i's preferences, i cannot get a house originally owned by $N_1 \cup \ldots \cup N_{j-1}$
  - Suppose i wants a house owned by $\ell \in N_k$, where where $k \in \{1,2,\ldots,j-1\}$
  - To get this house, ℓ must point to i in iteration k or earlier but this is not the case if $i \notin N_k$
- Truthful reporting gets i the best possible house they can achieve and thus is dominant strategy ∎

TTC Invariant. Let N_k be the set of agents removed in the kth iteration of the TTC algorithm. Every agent of N_k receives their favorite house outside of the houses owned by $N_1 \cup N_2 \cup N_{k-1}$ and the original owner of any house allocated in this round is also in N_k.

TTC: Unique Core Allocation

- Given a strict preference raking by n agents let $M(i)$ denote the house they receive by running TTC
- **Blocking coalition.** A subset $S \subseteq \{1,\ldots,n\}$ is a **blocking coalition** if members of S can trade houses amongst themselves such that at least one member is better off without making any member of S worse off.
  - If $S = N$, this property is the same as Pareto optimality of TTC
- **Core allocation.** An allocation is core is there is no such blocking coalition
- Stable allocations of DA are also called "core" allocations in the literature
- **Theorem.** For any house allocation instance, the output computed by the TTC algorithm is the unique core allocation.

TTC: Unique Core Allocation

- **Theorem.** For any house allocation instance, the output computed by the TTC algorithm is the unique core allocation.
- **Proof.** (Part 1 No other allocation can be core) Let N_j be defined by the TTC invariant.
- All agents of N_1 receive their first choice: this must be true in any core allocation
  - If not, the agents of N_1 can internally reallocate and can make everyone strictly better off
- Similarly, all agents of N_2 receive their top choice outside N_1
  - Given that every core allocation agrees with TTC for agents in N_1, such an allocation must also agree for agents in N_2
- Inductively, any core allocation must agree with TTC

TTC: Unique Core Allocation

- **Theorem.** For any house allocation instance, the output computed by the TTC algorithm is the unique core allocation.
- **Proof.** (Part 2 TTC allocation is core) Consider an arbitrary subset S
- Let $\ell = \min\{j \mid S \cap N_j \neq \emptyset\}$ (earliest round in which a member of S receives their house
- Consider $i \in N_\ell \cap S$, then i gets their favorite house among those not obtained by $N_1, \ldots, N_{\ell-1}$
- No member of S among these, that is,
  - $N_j \cap S = \emptyset$ for $j = 1, \ldots, \ell - 1$
  - Because ℓ is the first round where anyone in S gets their house
- No reallocation within S can make i better off.

Summary

- **TTC is awesome.** Computationally efficient, strategyproof, Pareto optimal and unique core allocation algorithm for exchange markets
- Given all its nice properties, we don't hear of it as much as lotteries, why??
  - Harder to explain what it does to a lay person
  - Harder for individuals to predict what outcome they will get

Leftovers on Stable Matching

Stability and Strategyproofness

- Lemma. Truthful reporting is a weakly dominant strategy for hospitals in the hospital-proposing deferred acceptance mechanism
  - While intuitive, this is surprisingly annoying to prove
  - See Theorem 10.6.18 in http://www.masfoundations.org/mas.pdf
  - Stability is only wrt to reported preferences, if someone misreports then stability is defined with respect to reported preferences only
- Is truthful reporting a dominant strategy if you are on the other-side of the market: for students in a hospital-proposing DA?
  - Let's take an example

Misreports from Students

- Consider the following truthful preference profile

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Consider the following truthful preference profile

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Consider the following truthful preference profile

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Consider the following truthful preference profile

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Consider the following truthful preference profile
  - Produces the following stable matching:
  - (MA, Beth), (NH, Chris), (OH, Aamir)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Class exercise. Can one of the students misreport their preferences to end up with a better match?
  - Does it every make sense to misreport about your top choice?
  - What about lower order misreports?

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

DA is not strategyproof (the receiving side can misreport and achieve a better match)

- Suppose Aamir misreports (swaps NH and OH)
- New matching: (MA, Aamir), (NH, Chris), (OH, Beth)
- Aamir improved from NH to top choice MA!

	1st	**2nd**	**3rd**
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	**2nd**	**3rd**
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Can't Have Both

- Can there be a mechanism that is both strategy proof and stable?
  - Unfortunately, no
- **Theorem**. No mechanism for two-sided matching is both stable and strategyproof.
  - Proof developed in Homework 6
- Many interesting questions:
  - How much information is needed to find a useful manipulation?
  - What is the optimal manipulation cheating strategy
- Empirically manipulations do not play a large role
  - If not many stable partners, can't gain much

The Match and its Evolution

- **NRMP Revisited**. The original 1952 implementation of the DA algorithm was the **hospital-optimal version**
- Students protested that the match was favoring hospitals

MONDAY, OCTOBER 22, 1951. The New York Times MONDAY, OCTOBER 22, 1951. 25 L+

MEDICAL SENIORS HIT INTERNE PLAN

Delegates of 44 Schools Meet Here to Protest Selection by 'Matching Machine'

TEACHERS PRAISE SYSTEM

They Argue That It Bars Unfair Recruiting—Students Insist on Choosing Their Hospitals

Delegates representing seniors in nearly all of the country's leading medical schools met here yesterday to express overwhelming opposition to a proposed mathematically contrived plan to place medical students in hospitals as internes.

They indicated that a great majority of their classmates preferred the present system whereby the country's hospitals, which have 10,000 internships, scramble for the best of a year's 6,000 medical graduates.

The meeting was held at Bard Hall of the College of Physicians and Surgeons of Columbia University, but it was made clear that the university was not its sponsor.

Seventy students attended the meeting at which forty-four colleges were represented either by delegates or through communications giving the opinion of the medical school's seniors on "the matching plan for internship" organized by the National Interassociation Committee on Internship about two years ago.

A prospectus of the plan of leading hospitals calls it "the accepted procedure for 1951-52" in determining which medical graduate shall go to what hospital to complete his medical education. Medical men said yesterday that its success or failure would have much to do with the chances of a medical student properly to complete his education during a wartime emergency.

By Winston Churchill: *The Second World War*

Volume V—Closing the Ring

INSTALLMENT 15—TEHERAN: CONCLUSIONS

Book II—Teheran to Rome

Conversation at Luncheon, December 1—The Frontiers of Poland—The "Curzon Line", and the Line of the Oder — Finland — "No Annexations and No Indemnities"—The Question of Germany—Partition?—President Roosevelt's Suggestion—I Unfold a Personal View — Marshal Stalin's Standpoint — Broad Agreement on Military Policy — Political Aspects Remote and Speculative—Deep Fear of German Might at This War Climax—The Present Partition—"It Cannot Last".

ROOSEVELT, STALIN, CHURCHILL AGREE ON PLANS FOR WAR ON GERMANY IN TALKS AT TEHERAN; 1,500 MORE TONS OF BOMBS DROPPED ON BERLIN

AGREEMENT: The news on the morning of Dec. 4, 1943.

Anglo-American Discussions in Cairo —Andaman Islands Plan—No Agreement at Our First Plenary Meeting, December 4—The President Agrees to Abandon Andamans Plan, December 5—President Roosevelt Decides to Appoint General Eisenhower to Command "Overlord"—The President and I Visit the Sphinx.

NEED FOR TEACHERS EXPECTED TO GROW

1,200 More a Year Required in State, Board Officials Say —Triple Sessions Feared

LACK IS WORST IN GRADES

Special Subjects Also Suffer— Syracuse Parley Cites High Birth Rate, Low Salaries

By LEONARD BUDER

The Match and its Evolution

- A new algorithm was adopted in 1997
  - Primary motivated was to give couples the option to get placed in geographically nearby programs
  - But in addition was made student-proposing
- Changes incentives for hospitals, but did it make a difference?
- Empirically, at least for the datasets arising in NRMP, less than 1% of the hospitals could have benefited by misreporting

Stable Matching Summary

- Hospital-proposing DA is hospital-optimal and student pessimal, among all stable matchings (regardless of the order of proposals)
- Stability matchings are not Pareto optimal overall, but are Pareto optimal among the set of all stable matchings
- Stable matchings are only strategyproof for the proposing side and cannot be strategyproof for both sides
- Lots of generalizations:
  - Incomplete preferences and ties
  - Stable "roommates" problem
  - Many-to-one stable matchings
  - Approximately stable matchings

Stable Matching Research

Mar 2024

Structural Complexities of Matching Mechanisms*

Yannai A. Gonczarowski† Clayton Thomas‡

March 30, 2024

Abstract

We study various novel complexity measures for two-sided matching mechanisms, applied to the two canonical strategyproof matching mechanisms, Deferred Acceptance (DA) and Top Trading Cycles (TTC). Our metrics are designed to capture the complexity of various structural (rather than computational) concerns, in particular ones of recent interest within economics. We consider a unified, flexible approach to formalizing our questions: Define a protocol or data structure performing some task, and bound the number of bits that it requires. Our main results apply this approach to four questions of general interest; for mechanisms matching applicants to institutions, our questions are:

(1) How can one applicant affect the outcome matching?

(2) How can one applicant affect another applicant's set of options?

(3) How can the outcome matching be represented / communicated?

(4) How can the outcome matching be verified?

2021

Deferred Acceptance with Compensation Chains

PIOTR DWORCZAK, Stanford University, Graduate School of Business

I introduce a class of algorithms called Deferred Acceptance with Compensation Chains (DACC). DACC algorithms generalize the DA algorithms by Gale and Shapley [1962] by allowing both sides of the market to make offers. The main result is a characterization of the set of stable matchings: a matching is stable if and only if it is the outcome of a DACC algorithm.

Nov 2024

Stable Matching with Ties: Approximation Ratios and Learning

Shiyun Lin * Simon Mauras † Nadav Merlis ‡ Vianney Perchet §

November, 2024

Abstract

We study the problem of matching markets with ties, where one side of the market does not necessarily have strict preferences over members at its other side. For example, workers do not always have strict preferences over jobs, students can give the same ranking for different schools and more. In particular, assume w.l.o.g. that workers' preferences are determined by their utility from being matched to each job, which might admit ties. Notably, in contrast to classical two-sided markets with strict preferences, there is no longer a single stable matching that simultaneously maximizes the utility for all workers.

We aim to guarantee each worker the largest possible share from the utility in her best possible stable matching. We call the ratio between the worker's best possible stable utility and its assigned utility the *Optimal Stable Share* (OSS)-ratio. We first prove that distributions over stable matchings cannot guarantee an OSS-ratio that is sublinear in the number of workers. Instead, randomizing over possibly non-stable matchings, we show how to achieve a tight logarithmic OSS-ratio. Then, we analyze the case where the real utility is not necessarily known and can only be approximated. In particular, we provide an algorithm that guarantees a similar fraction of the utility compared to the best possible utility. Finally, we move to a bandit setting, where we select a matching at each round and only observe the utilities for matches we perform. We show how to utilize our results for approximate utilities to gracefully interpolate between problems without ties and problems with statistical ties (small suboptimality gaps).

Jan 2025

UNBALANCED RANDOM MATCHING MARKETS WITH PARTIAL PREFERENCES

ADITYA POTUKUCHI AND SHIKHA SINGH

ABSTRACT. Properties of stable matchings in the popular random-matching-market model have been studied for over 50 years. In a random matching market, each agent has complete preferences drawn uniformly and independently at random. Wilson (1972), Knuth (1976) and Pittel (1989) proved that in balanced random matching markets, the proposers are matched to their $\ln n$th choice on average. In this paper, we consider markets where agents have partial (truncated) preferences, that is, the proposers only rank their top d partners. Despite the long history of the problem, the following fundamental question remained unanswered: *what is the smallest value of d that results in a perfect stable matching with high probability?* In this paper, we answer this question exactly—we prove that a degree of $\ln^2 n$ is necessary and sufficient. That is, we show that if $d < (1-\varepsilon)\ln^2 n$ then no stable matching is perfect and if $d > (1+\varepsilon)\ln^2 n$, then every stable matching is perfect with high probability. This settles a recent conjecture by Kanoria, Min and Qian (2021).

We generalize this threshold for unbalanced markets: we consider a matching market with n agents on the shorter side and $n(\alpha+1)$ agents on the longer side. We show that for markets with $\alpha = o(1)$, the sharp threshold characterizing the existence of perfect stable matching occurs when d is $\ln n \cdot \ln\left(\frac{1+\alpha}{\alpha+(1/n(\alpha+1))}\right)$.

Finally, we extend the line of work studying the effect of imbalance on the expected rank of the proposers (termed the "stark effect of competition"). We establish the regime in unbalanced markets that forces this stark effect to take shape in markets with partial preferences.

2021

Tiered Random Matching Markets: Rank Is Proportional to Popularity

Itai Ashlagi
Department of Management Science and Engineering, Stanford University, CA, USA
iashlagi@stanford.edu

Mark Braverman
Department of Computer Science, Princeton University, NJ, USA
mbraverm@cs.princeton.edu

Amin Saberi
Department of Management Science and Engineering, Stanford University, CA, USA
saberi@stanford.edu

Clayton Thomas
Department of Computer Science, Princeton University, NJ, USA
claytont@cs.princeton.edu

Geng Zhao
Department of Computer Science, Stanford University, CA, USA
gengz@stanford.edu

2022

On Fairness and Stability in Two-Sided Matchings

Gili Karni ✉
Weizmann Institute of Science, Rehovot, Israel

Guy N. Rothblum ✉
Weizmann Institute of Science, Rehovot, Israel

Gal Yona ✉
Weizmann Institute of Science, Rehovot, Israel

Matching Application:
Kidney Exchange

Kidney Exchange

- Many people suffer from kidney failure and need a transplat
- In the US in 2013, around 100,000 people were on a waiting list to receive kidneys
- A third of kidney transplants come from living organ donors
- Unfortunately, having a kidney is not enough, sometimes a patient-donor pair is incompatible
- Two incompatible donor-patient pairs might be able to participate in an exchange
- National kidney exchanges have gain momentum
- Kidney exchange is legal but compensation for organ donation illegal in US (and every country except Iran)
    - Ideal application for mechanism design without money

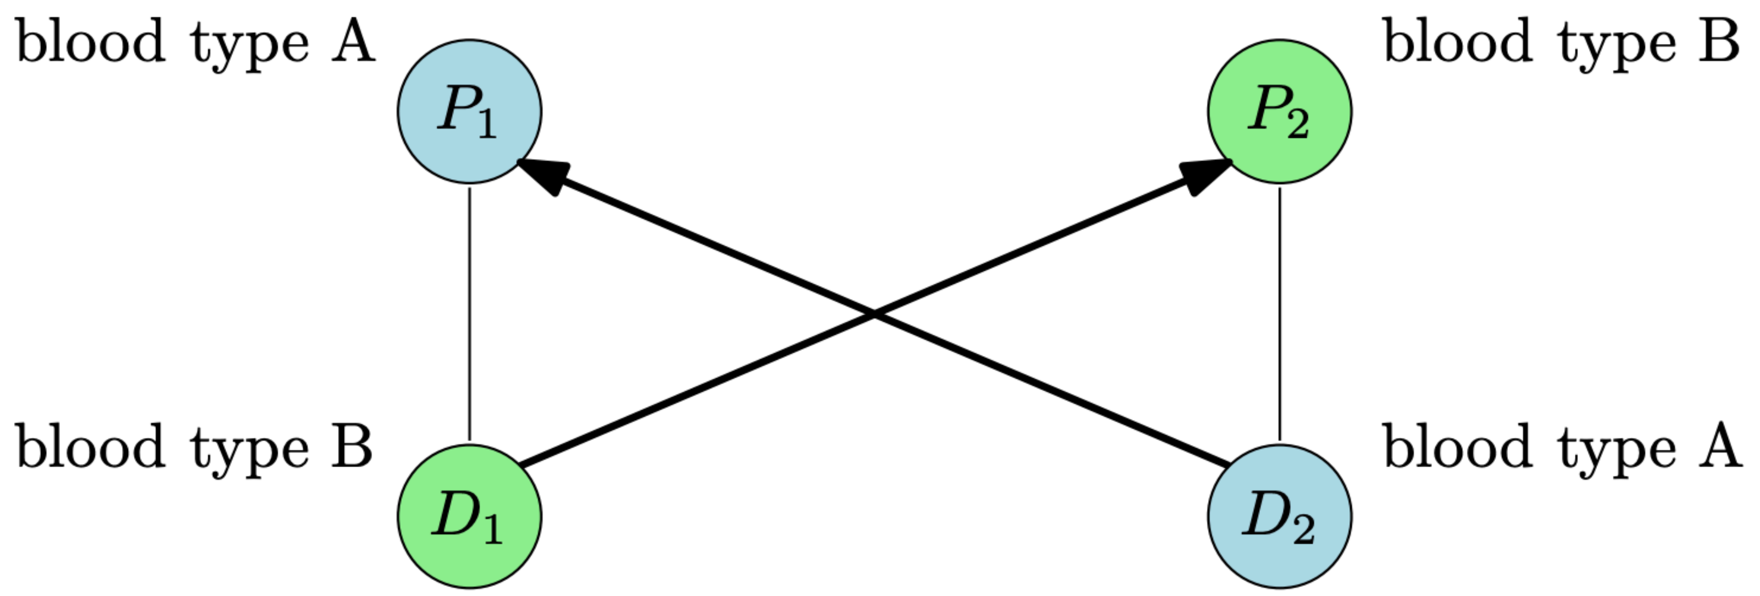

Using TTC: Challenges

- In an influential study in 2004, Roth Sonmez and Unver advocated for the TTC algorithm for kidney exchange
- Agent, house pairs are now patient, donor pairs
- A total ordering over kidneys can be determined by the likelihood of the transplant being successful
- The goal is to reallocate kidneys in way that everyone is collectively as better off as possible
- The actual problem is a bit more complicated and TTC extensions can handle some of them (e.g., accommodating patients without donors, and deceased donors)
- The biggest dealbreaker in TTC for kidney exchange is **long trading cycles**

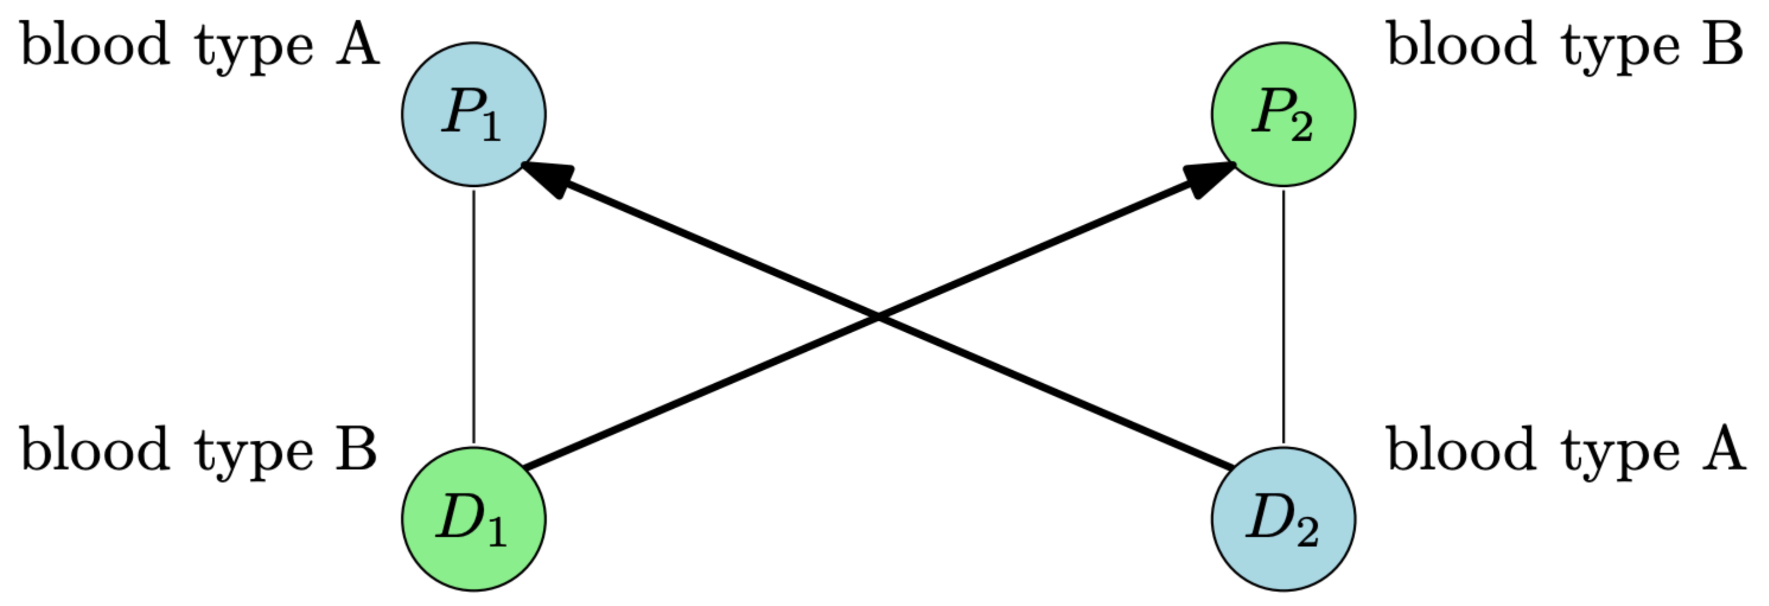

Using TTC: Challenges

- The biggest dealbreaker in TTC for kidney exchange is long trading cycles
    - Transplants must occur simultaneously due to incentive issues (if surgeries for P1 and D2 happen first, there is a risk that D1 will renege on its offer)
- TCC model requires a total ordering over kidneys
    - In reality patients don't care which kidney they get as long as it is compatible with them
    - Binary preferences are more appropriate
- These challenges triggered further research into a DSIC mechanism for kidney exchange

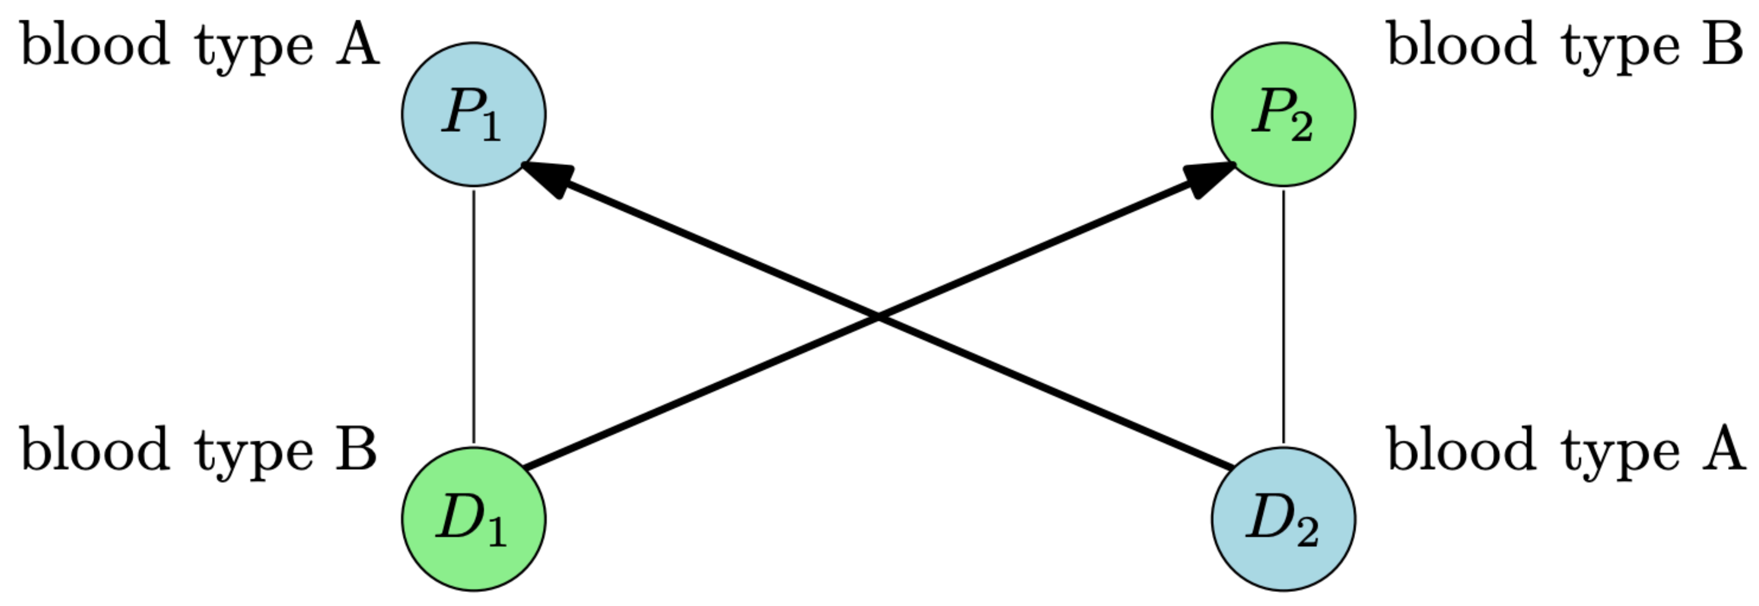

Max Cardinality Matching

- In a subsequent paper, Roth Sonmez and Unver propose using matchings
- The nodes of the graph are patient donor pairs and edges are between compatible pairs that can lead to an exchange
    - Matchings lead to 2-way swaps
- **Model**. Each agent i has a true edge set E_i and can report any subset $F_i \subseteq E_i$ to a mechanism (patients can refuse exchanges $E_i \backslash F_i$ for any reason)
- **Goal**. Compute a maximum-cardinality matching and to be DSIC (for each agent, truthfully reporting its full edge set is a dominant strategy.)

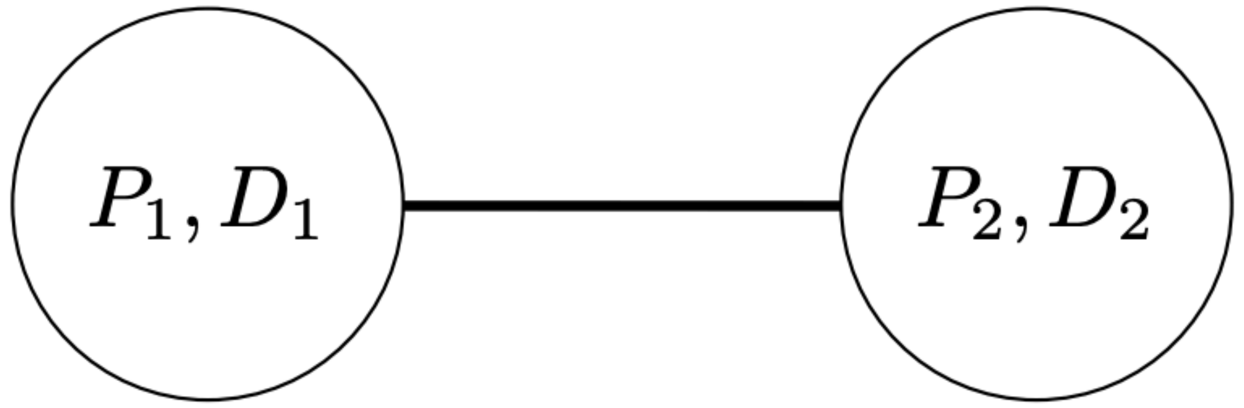

Multiple Matchings

- Even if we collect preferences, create a graph and find a maximum-cardinality matching, there is still a wrinkle
- A graph can have many matchings of the same cardinality
    - How do we handle tie breaks?

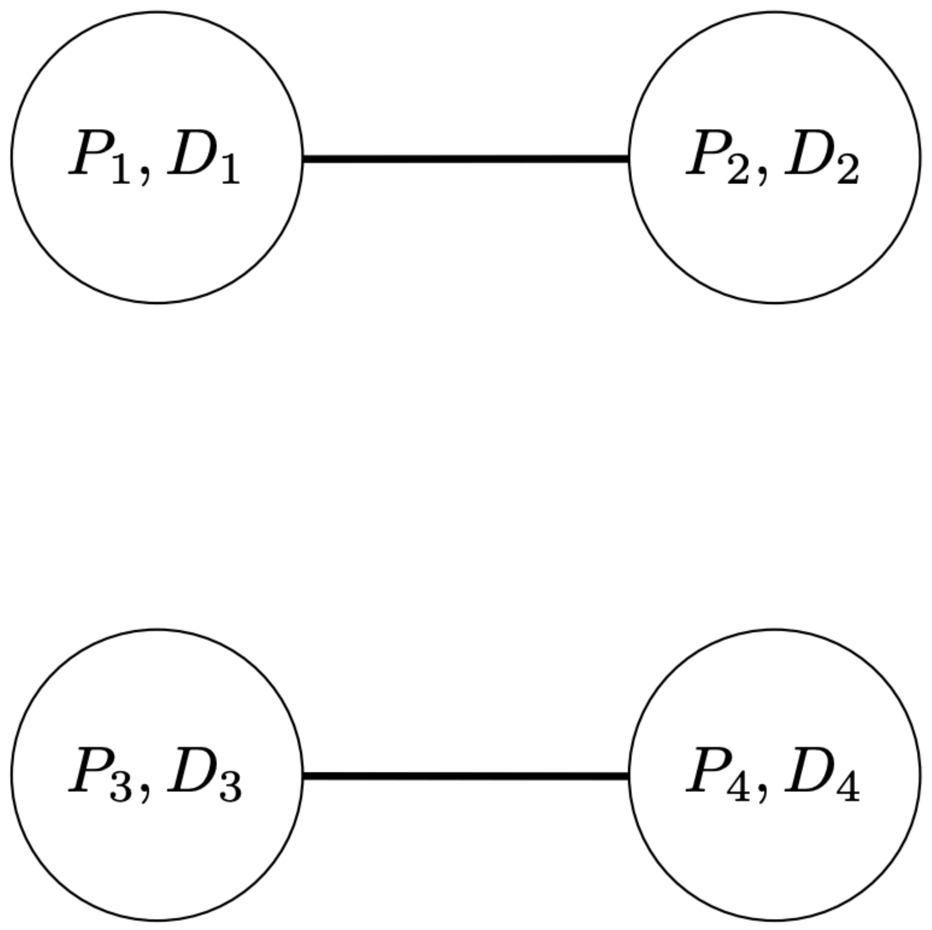

vs

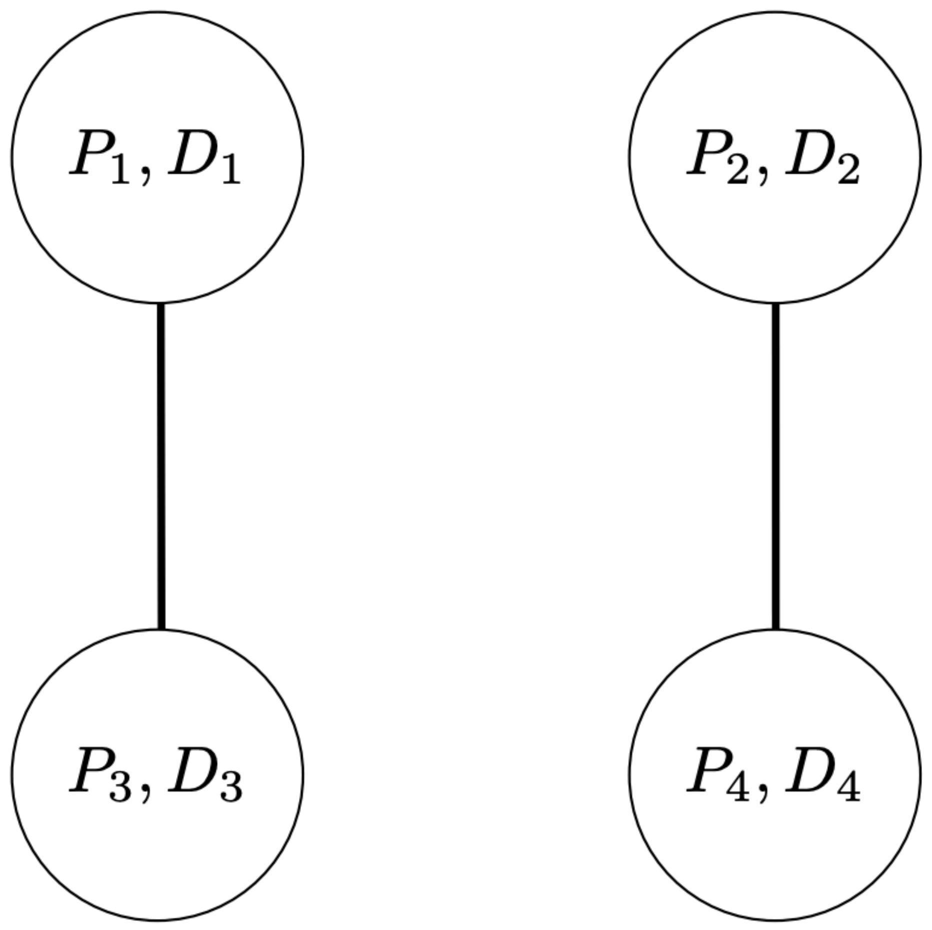

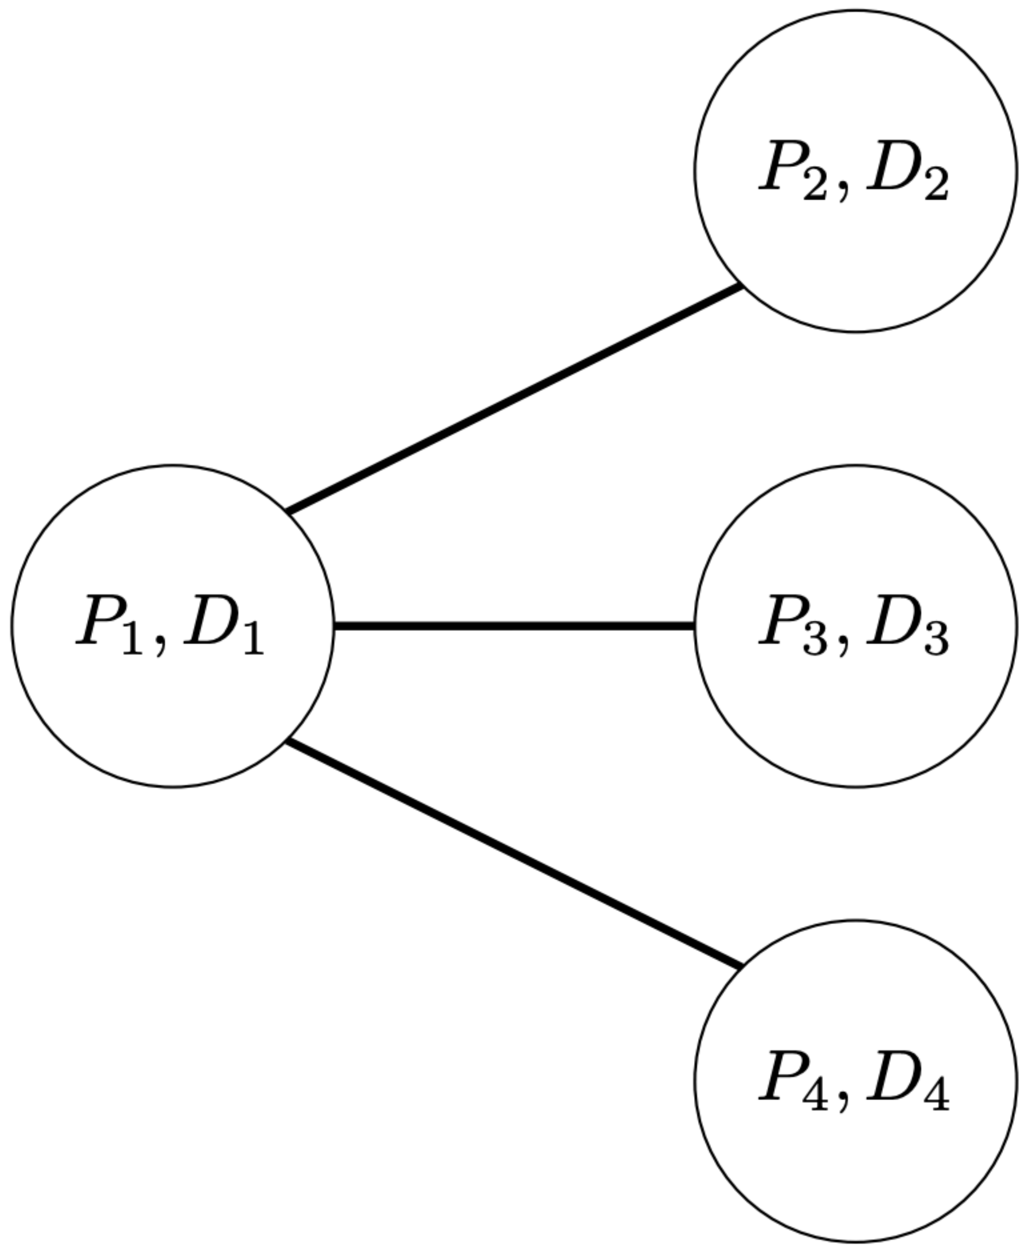

Priority Order Over Nodes

- One way this is resolved through a priority order over nodes
- A priority maximum matching mechanism turns out to be DSIC: no agent can go from unmatched to match by reporting a subset of its edges

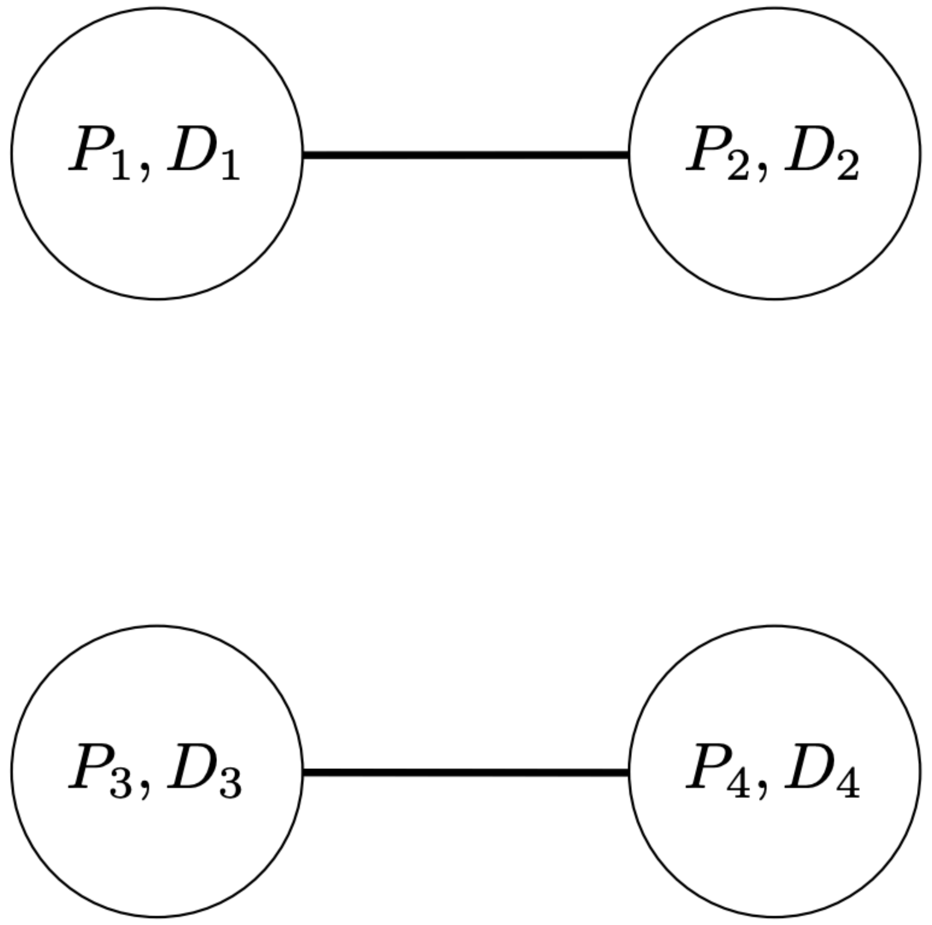

vs

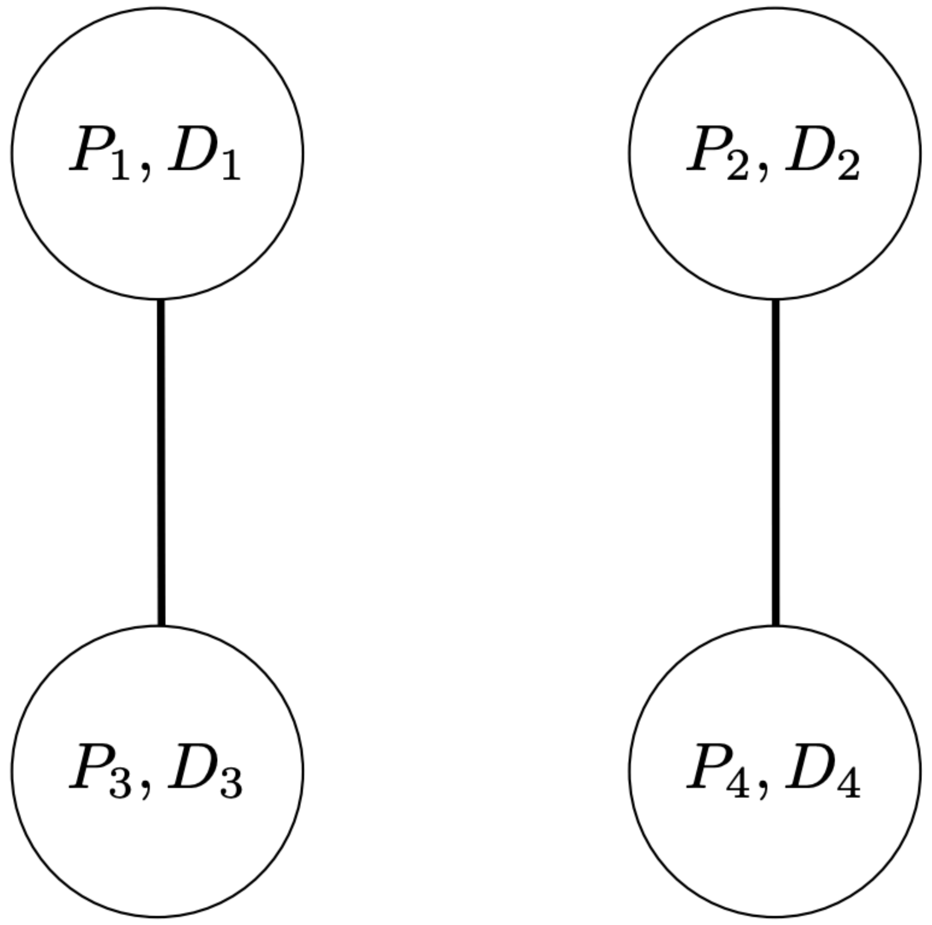

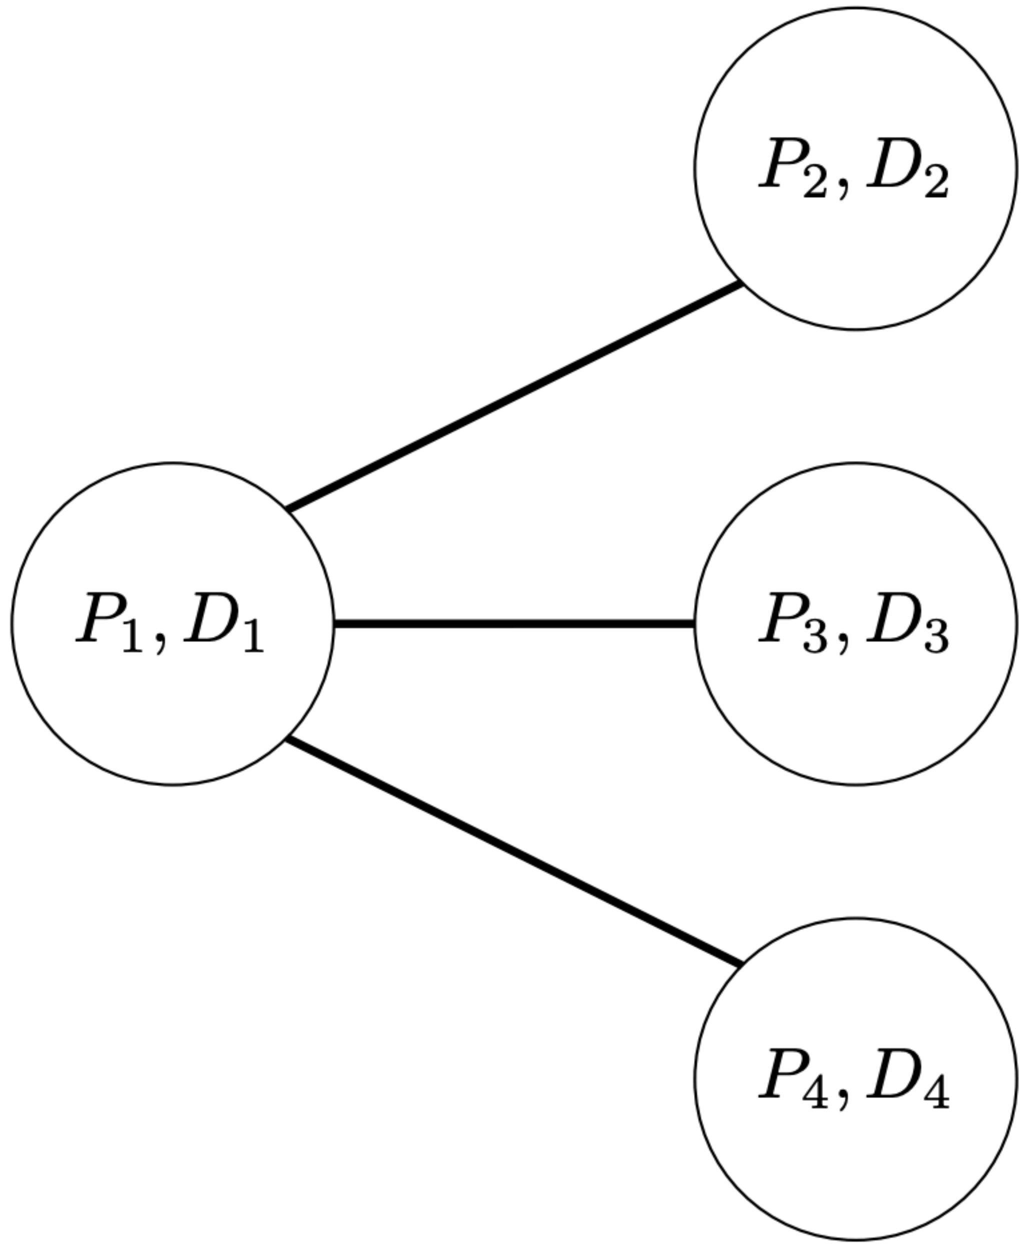

Challenges

- Need for full reporting at the hospital level
- Objective of individual hospitals: match as many of their patients as possible
- Objective of society: match as many patients as possible
- Need for approximately optimal DSIC mechanisms

Incentives of H1 and H2 are at odds: no DSIC mechanism that maximizes cardinality of matching