DS 574 Algorithmic Mechanism Design	**Lecture #11 Worksheet**
Fall 2023	Prof. Kira Goldner

Kidney Exchange

- There are more than 92,000 on the waitlist for a kidney transplant in the US; this makes up 87% of the organ transplant list [1].
- Healthy people have two kidneys and can survive fine with only one.
- A donor and recipient must be "compatible" (blood and tissue types).
- Two incompatible patient-donor pairs can agree to a kidney exchange. This is legal. (Compensation for kidneys is not, except in Iran.)

Question: How would one design a centralized mechanism for kidney exchange, where incompatible patient-donor pairs can register and be matched with others?

Idea #1: Use the Top Trading Cycle Algorithm

Vanilla Top Trading Cycles

Consider the housing allocation problem defined by Shapley and Scarf [5]: There are n agents, and each initially owns one house. Each agent has a total ordering over the n houses, and need not prefer their own over the others. How can we reallocate the houses to make the agents better off?

The Top Trading Cycle Algorithm [Gale [5]].

While agents remain:

- Each remaining agent points to its favorite remaining house. This induces a directed graph G on the remaining agents in which every vertex has out-degree 1 (Figure 1).
- The graph G has at least one directed cycle. Self-loops count as directed cycles.
- Reallocate as suggested by the directed cycles, with each agent on a directed cycle C giving its house to the agent that points to it, that is, to its predecessor on C.
- Delete the agents and the houses that were reallocated in the previous step.

Observations:

•
•
•

Theorem 1. *The TTCA induces a DSIC mechanism.*

[*Hint:* Divide agents into those allocated to in the jth iteration.]

Proof.

Definition 1. A *core allocation* is an allocation such that no coalition of agents can make all of its members better off via internal reallocations.

Theorem 2. *For every house allocation problem, the allocation computed by the TTCA is the unique core allocation.*

Modifications for Kidney Exchange

The first attempt was via the TTCA by Roth, Sönmez, and Ünver [3] before the authors talked extensively to doctors. People's "preferences" over kidneys would just be via decreasing probability of success of the transplant.

But kidney exchange is more complicated:

(1) There are patients without living donors, and deceased donors.

(2) The cycles along which reallocations are made can be arbitrarily long.

(3) Modeling preferences as a total ordering over the set of living donors is overkill: empirically, patients don't really care which kidney they get as long as it is compatible with them.

    Instead:

Idea #2: Use a Matching Algorithm

(2) Short reallocation cycles and (3) binary preferences motivate looking for *matchings*, as done in [4].

What's the relevant graph for kidney exchange? Describe the vertices, edges, and what a matching would look like.

How do incentives work here? What should the mechanism look like?

(1)

(2)

(3)

But how do we tie-break between maximum-cardinality matchings?

Theorem 3. *For every collection $\{E_i\}_{i=1}^n$ of edge sets and every ordering of the vertices, the priority matching mechanism above is DSIC: no agent can go from unmatched to matched by reporting a strict subset F_i of E_i rather than E_i.*

Hospital Incentives

Current research is focused on incentive problems at the *hospital* level, rather than at the level of individual patient-donor pairs. Hospitals are the ones who actually report the pairs to the national kidney exchange, but the objectives of a hospital (to match as many of its patients as possible) and of society (to match as many patients overall as possible) are not perfectly aligned.

The Need for Full Reporting. Only reporting pairs who the hospital can't match internally can result in fewer exchanges.

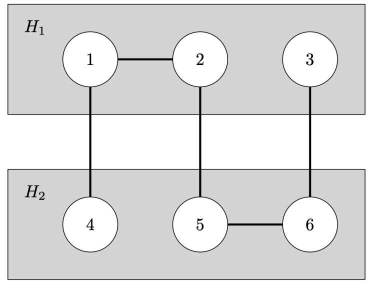

Figure 1: Full reporting by hospitals leads to more matches than with only internal matches.

Hiding patients. If H_1 hides patients 2 and 3 from the exchange (while H_2 reports truthfully), then H_1 guarantees that all of its patients are matched. The unique maximum matching in the report graph matches patient 6 with 7 (and 4 with 5), and H_1 can match 2 and 3 internally. On the other hand, if H_2 hides patients 5 and 6 while H_1 reports truthfully, then all of H_2's patients are matched. In this case, the unique maximum matching in the graph of report matches patient 1 with 2 and 4 with 3, while H_2 can match patients 5 and 6 internally.

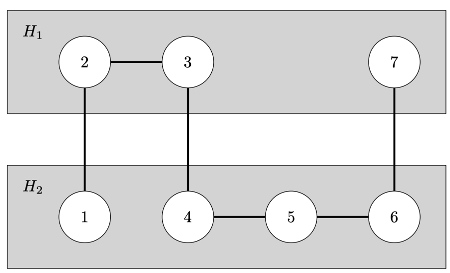

Figure 2: Hospitals can have an incentive to hide patient-donor pairs.

It turns out there cannot be a DSIC mechanism that always computes a maximum-cardinality matching in the full graph.

In light of this example, the revised goal should be to compute an approximately maximum-cardinality matching so that, for each participating hospital, the number of its patients that get matched is approximately as large as in any matching, maximum-cardinality or otherwise. Understanding the extent to which this is possible, in both theory and practice, is an active research topic [2, 6].

References

[1] American Kidney Fund, Jun 2022.

[2] Itai Ashlagi, Felix Fischer, Ian A Kash, and Ariel D Procaccia. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*, 91:284–296, 2015.

[3] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Kidney exchange. *The Quarterly journal of economics*, 119(2):457–488, 2004.

[4] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.

[5] Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.

[6] Panagiotis Toulis and David C Parkes. A random graph model of kidney exchanges: efficiency, individual-rationality and incentives. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 323–332, 2011.