Mechanisms for Large-Scale Kidney Exchanges

(And More)

Felix Fischer
Harvard University
fischerf@seas.harvard.edu

(joint work with Noga Alon, Itai Ashlagi, Ian Kash, Ariel Procaccia, and Moshe Tennenholtz)

October 13, 2010

Chronic Kidney Disease

- Kidneys: filter wastes from the blood, among other things
- Chronic renal failure
  - life-threatening
  - usually irreversible
  - only treatment is replacement of kidney function
- Short run: dialysis
  - filtering outside the body, 3 to 7 times a week
  - low quality of life, low life expectancy
- Long run: transplantation of a healthy kidney

Kidney Transplantation

data and more from a very nice talk by Al Roth, see http://crcs.seas.harvard.edu/videos/#roth

- Situation in the US:
  - more than 86,000 patients on the waiting list
  - 33,678 patients added in 2009
  - 10,441 transplants from deceased donors
  - 4,456 patients died while on the waiting list, 1,941 more removed as "too sick to transplant"

Kidney Transplantation

data and more from a very nice talk by Al Roth, see http://crcs.seas.harvard.edu/videos/#roth

- Situation in the US:
  - more than 86,000 patients on the waiting list
  - 33,678 patients added in 2009
  - 10,441 transplants from deceased donors
  - 4,456 patients died while on the waiting list, 1,941 more removed as “too sick to transplant”
- Good news:
  - live donation possible, without long-term effect on donor
  - many patients have somebody (usually a relative or friend) willing to donate them one of their kidneys

Outline

Kidney Exchanges

Incentives in Large-Scale Kidney Exchanges

Results: Mechanisms for Large-Scale Kidney Exchanges

A Larger Agenda: Approximate Mechanisms without Payments

Kidney Exchange

- Kidney exchange enables transplantations in cases where patient and potential donor are not compatible

patient 1
blood type A

donor 1
blood type B

Kidney Exchange

- Kidney exchange enables transplantations in cases where patient and potential donor are not compatible

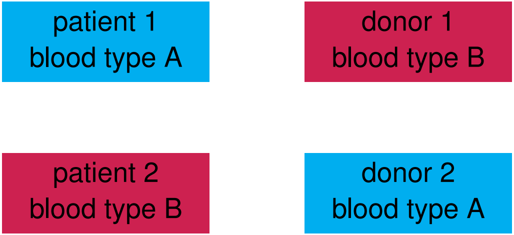

Kidney Exchange

- Kidney exchange enables transplantations in cases where patient and potential donor are not compatible

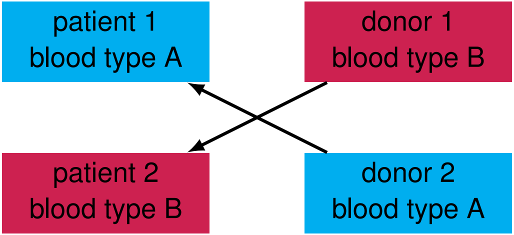

Kidney Exchange

- Kidney exchange enables transplantations in cases where patient and potential donor are not compatible

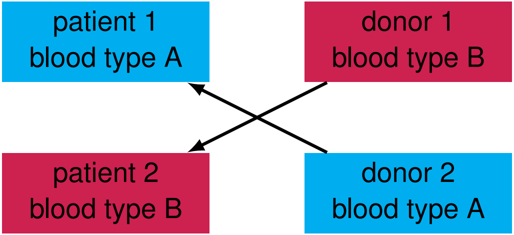

- 6,387 transplants from living donors in 2009 (first year with more living than deceased donors)

Kidney Exchange

- One-to-one barter (item-to-item exchange) is hard
- Money would be one way to solve this, but buying and selling organs is illegal (National Organ Transplant Act, 1984)
- Amended in 2007 to allow kidney exchange, but we still need to bring (compatible) donors and patients together

Kidney Exchange Institutions

- One-to-one barter (item-to-item exchange) is hard
- Money would be one way to solve this, but buying and selling organs is illegal (National Organ Transplant Act, 1984)
- Amended in 2007 to allow kidney exchange, but we still need to bring (compatible) donors and patients together

- New England Program for Kidney Exchange: since 2005, 14 transplant centers in New England
- Alliance for Paired Donation: since 2006, 81 transplant centers in 27 states

- A national exchange would clearly be desirable

Incentives, Old and New

- Transplantations conducted simultaneously to control incentives (and other problems) within a single paired exchange
- In large-scale exchanges, incentives of hospitals become an issue (first pointed out by Roth, can be observed in practice)

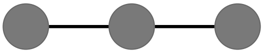

(from now on, vertices represent donor-patient pairs)

Incentives, Old and New

- Transplantations conducted simultaneously to control incentives (and other problems) within a single paired exchange
- In large-scale exchanges, incentives of hospitals become an issue (first pointed out by Roth, can be observed in practice)

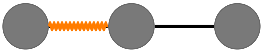

(from now on, vertices represent donor-patient pairs)

Incentives, Old and New

- Transplantations conducted simultaneously to control incentives (and other problems) within a single paired exchange
- In large-scale exchanges, incentives of hospitals become an issue (first pointed out by Roth, can be observed in practice)

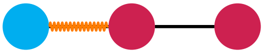

(from now on, vertices represent donor-patient pairs)

Incentives, Old and New

- Transplantations conducted simultaneously to control incentives (and other problems) within a single paired exchange
- In large-scale exchanges, incentives of hospitals become an issue (first pointed out by Roth, can be observed in practice)

(from now on, vertices represent donor-patient pairs)

Incentives, Old and New

- Transplantations conducted simultaneously to control incentives (and other problems) within a single paired exchange
- In large-scale exchanges, incentives of hospitals become an issue (first pointed out by Roth, can be observed in practice)

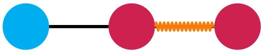

(from now on, vertices represent donor-patient pairs)

Incentives, Old and New

- Transplantations conducted simultaneously to control incentives (and other problems) within a single paired exchange
- In large-scale exchanges, incentives of hospitals become an issue (first pointed out by Roth, can be observed in practice)

(from now on, vertices represent donor-patient pairs)

The Formal Model

- Set N of agents, corresponding to hospitals
- Graph $G = (V, E)$ with $V = \biguplus_i V_i$
  - Each vertex $v \in V$ corresponds to a donor-patient pair
  - Edge $(u, v) \in E$ means donor of u is compatible with patient of v, and donor of v with patient of u

- Agents report subsets $V_i' \subseteq V_i$
- Mechanism produces matching M' of subgraph induced by $\biguplus_i V_i'$
- Agent i adds matching $\widehat{M}_i$ for hidden and unmatched vertices
- Utility of agent i is number of vertices in V_i matched in M' or $\widehat{M}_i$

Goals

- No payments
- Strategyproofness: agents cannot gain by hiding vertices
- Approximate optimality: call a mechanism α-approximate if returned matching is at most by factor α smaller than maximum cardinality matching ($\alpha = 1$ is optimal, smaller is better)
- We are going to prove lower bounds (impossibility results) and upper bounds

- Complementary results: optimality on average with a weaker notion of stability (Ashlagi and Roth, 2010)

Lower Bounds

Theorem: If there are at least two agents,

no deterministic strategyproof mechanism can be α-approximate for $\alpha < 2$, and

no randomized strategyproof mechanism can be α-approximate for $\alpha < 4/3$.

Lower Bounds

Theorem: If there are at least two agents,

no deterministic strategyproof mechanism can be α-approximate for $\alpha < 2$, and

no randomized strategyproof mechanism can be α-approximate for $\alpha < 4/3$.

Lower Bounds

Theorem: If there are at least two agents,

no deterministic strategyproof mechanism can be α-approximate for $\alpha < 2$, and

no randomized strategyproof mechanism can be α-approximate for $\alpha < 4/3$.

Lower Bounds

Theorem: If there are at least two agents,

no deterministic strategyproof mechanism can be α-approximate for $\alpha < 2$, and

no randomized strategyproof mechanism can be α-approximate for $\alpha < 4/3$.

Lower Bounds

Theorem: If there are at least two agents,

no deterministic strategyproof mechanism can be α-approximate for $\alpha < 2$, and

no randomized strategyproof mechanism can be α-approximate for $\alpha < 4/3$.

A Deterministic Mechanism for Two Agents

- Choose matching that has
  - (i) maximum cardinality on both V_1 and V_2 and
  - (ii) maximum cardinality among all matchings satisfying (i)

A Deterministic Mechanism for Two Agents

- Choose matching that has
  (i) maximum cardinality on both V_1 and V_2 and
  (ii) maximum cardinality among all matchings satisfying (i)

A Deterministic Mechanism for Two Agents

- Choose matching that has
  - (i) maximum cardinality on both V_1 and V_2 and
  - (ii) maximum cardinality among all matchings satisfying (i)

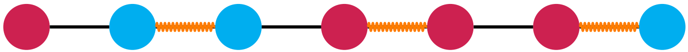

A Deterministic Mechanism for Two Agents

- Choose matching that has
  - (i) maximum cardinality on both V_1 and V_2 and
  - (ii) maximum cardinality among all matchings satisfying (i)

A Deterministic Mechanism for Two Agents

- Choose matching that has
  (i) maximum cardinality on both V_1 and V_2 and
  (ii) maximum cardinality among all matchings satisfying (i)

- Strategyproof (more on the next slide)
- 2-approximate, as returned matching is inclusion-maximal

A Generalization

- Fix a bipartition $\Pi = (\Pi_1, \Pi_2)$ of N
- MATCH_Π: choose matching that has
  (i) maximum cardinality on V_i for all $i \in N$
  (ii) no edges between V_i and V_j if $i, j \in \Pi_\ell$ for $\ell \in \{1, 2\}$
  (iii) maximum cardinality among all matchings satisfying (i) and (ii), breaking ties serially

A Generalization

- Fix a bipartition $\Pi = (\Pi_1, \Pi_2)$ of N
- $\textsc{Match}_\Pi$: choose matching that has
  - (i) maximum cardinality on V_i for all $i \in N$
  - (ii) no edges between V_i and V_j if $i, j \in \Pi_\ell$ for $\ell \in \{1, 2\}$
  - (iii) maximum cardinality among all matchings satisfying (i) and (ii), breaking ties serially
- $\Pi = (\{1\}, \{2\})$ yields the two-agent mechanism we just saw

A Generalization

- Fix a bipartition $\Pi = (\Pi_1, \Pi_2)$ of N
- $\textsc{Match}_\Pi$: choose matching that has
  (i) maximum cardinality on V_i for all $i \in N$
  (ii) no edges between V_i and V_j if $i, j \in \Pi_\ell$ for $\ell \in \{1, 2\}$
  (iii) maximum cardinality among all matchings satisfying (i) and (ii), breaking ties serially
- $\Pi = (\{1\}, \{2\})$ yields the two-agent mechanism we just saw

Theorem: For any number of agents and any bipartition Π, $\textsc{Match}_\Pi$ is strategyproof.

A Generalization

- Fix a bipartition $\Pi = (\Pi_1, \Pi_2)$ of N
- Match_Π: choose matching that has
  - (i) maximum cardinality on V_i for all $i \in N$
  - (ii) no edges between V_i and V_j if $i, j \in \Pi_\ell$ for $\ell \in \{1, 2\}$
  - (iii) maximum cardinality among all matchings satisfying (i) and (ii), breaking ties serially
- $\Pi = (\{1\}, \{2\})$ yields the two-agent mechanism we just saw

Theorem: For any number of agents and any bipartition Π, Match_Π can be executed in polynomial time.

Proof idea: reduction to weighted matching

A Generalization

- Fix a bipartition $\Pi = (\Pi_1, \Pi_2)$ of N
- $\textsc{Match}_\Pi$: choose matching that has
  (i) maximum cardinality on V_i for all $i \in N$
  (ii) no edges between V_i and V_j if $i, j \in \Pi_\ell$ for $\ell \in \{1, 2\}$
  (iii) maximum cardinality among all matchings satisfying (i) and (ii), breaking ties serially
- $\Pi = (\{1\}, \{2\})$ yields the two-agent mechanism we just saw

- Strategyproof and 2-approximate mechanism for two agents
- No finite approximation ratio for more than two agents

Mix and Match

- MIX-AND-MATCH:
  1. Construct random bipartition $\Pi = (\Pi_1, \Pi_2)$: for each agent flip a fair coin to determine whether she goes to Π_1 or Π_2
  2. Execute MATCH_Π

Mix and Match

- MIX-AND-MATCH:
  1. Construct random bipartition $\Pi = (\Pi_1, \Pi_2)$: for each agent flip a fair coin to determine whether she goes to Π_1 or Π_2
  2. Execute MATCH_Π

Theorem: For any number of agents, MIX-AND-MATCH is (universally) strategyproof and 2-approximate in expectation.

What We (Don't) Know

	deterministic		randomized	
	lower bound	upper bound	lower bound	upper bound
two agents	2	2	$4/3$	2
n agents	2	$O(k)$	$4/3$	2

What We (Don't) Know

	deterministic		randomized	
	lower bound	upper bound	lower bound	upper bound
two agents	2	2	$4/3$	2
n agents	2	$O(k)$	$4/3$	2

What We (Don't) Know

	deterministic		randomized	
	lower bound	upper bound	lower bound	upper bound
two agents	2	2	$4/3$	2
n agents	2	$O(k)$	$4/3$	2

FLIP-AND-MATCH: with probability $1/2$ each

- execute MIX-AND-MATCH
- return a (specific) maximum cardinality matching

What We (Don't) Know

	deterministic		randomized	
	lower bound	upper bound	lower bound	upper bound
two agents	2	2	$4/3$	2
n agents	2	$O(k)$	$4/3$	2

- Ongoing work
  - simulations with true distribution of blood types
  - longer exchange sequences (three or four might be enough)

Approximate Mechanisms without Payments

- VCG mechanism: socially optimal outcome, payments according to damage caused to respective other agents
- Legal, ethical, or technical reasons may prohibit payments
- Examples: kidney exchanges, voting

- Dekel, F, Procaccia (2008); Procaccia and Tennenholtz (2009): in the absence of payments, select outcome that is as close as possible to the optimal one in the worst case
- Same objective as in algorithmic mechanism design, but used due to incentive rather than computational constraints

Strategyproof Selection from the Selectors

- Approval voting
  - each voter approves of set of candidates (of any size)
  - choose candidate (or committee of desired size) with largest number of votes
- Strategyproof (assuming dichotomous preferences)

- No longer the case when sets of candidates and voters coincide
  - scientific organizations (GTS, AMS, IEEE, IFAAMAS)
  - web graph, (directed) social networks, reputation systems

The Model

- Set $N = [n]$ of agents
- Directed graph $G = (N, E) \in \mathcal{G}$, no self-loops
- Ideally: select $S \in \mathcal{S}_k = \{T \subseteq N : |T| = k\}$ to maximize
  $\sum_{i \in S} \deg(i) = \sum_{i \in S} |\{j \in N : (j, i) \in E\}|$
- Mechanism $M : \mathcal{G} \to \Delta(\mathcal{S}_k)$
- Strategyproofness: probability of selecting i independent of edges (i, j) for $j \in N$

The Model

- Set $N = [n]$ of agents
- Directed graph $G = (N, E) \in \mathcal{G}$, no self-loops
- Ideally: select $S \in \mathcal{S}_k = \{T \subseteq N : |T| = k\}$ to maximize
  $\sum_{i \in S} \deg(i) = \sum_{i \in S} |\{j \in N : (j, i) \in E\}|$
- Mechanism $M : \mathcal{G} \to \Delta(\mathcal{S}_k)$
- Strategyproofness: probability of selecting i independent of edges (i, j) for $j \in N$

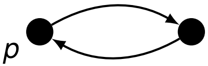

The Model

- Set $N = [n]$ of agents
- Directed graph $G = (N, E) \in \mathcal{G}$, no self-loops
- Ideally: select $S \in \mathcal{S}_k = \{T \subseteq N : |T| = k\}$ to maximize
  $\sum_{i \in S} \deg(i) = \sum_{i \in S} |\{j \in N : (j, i) \in E\}|$
- Mechanism $M : \mathcal{G} \to \Delta(\mathcal{S}_k)$
- Strategyproofness: probability of selecting i independent of edges (i, j) for $j \in N$

$\leq \frac{1}{2}$ 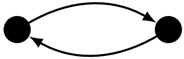

$\leq \frac{1}{2}$ 

The Model

- Set $N = [n]$ of agents
- Directed graph $G = (N, E) \in \mathcal{G}$, no self-loops
- Ideally: select $S \in \mathcal{S}_k = \{T \subseteq N : |T| = k\}$ to maximize
  $\sum_{i \in S} \deg(i) = \sum_{i \in S} |\{j \in N : (j, i) \in E\}|$
- Mechanism $M : \mathcal{G} \to \Delta(\mathcal{S}_k)$
- Strategyproofness: probability of selecting i independent of edges (i, j) for $j \in N$
- α-efficiency: for every graph,

$$\frac{\max_{S \in \mathcal{S}_k} \sum_{i \in S} \deg(i)}{\mathbb{E}_{S \sim M}[\sum_{i \in S} \deg(i)]} \leq \alpha$$

$\leq \frac{1}{2}$ 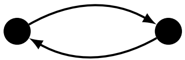

$\leq \frac{1}{2}$ 

The Model

- Set $N = [n]$ of agents
- Directed graph $G = (N, E) \in \mathcal{G}$, no self-loops
- Ideally: select $S \in \mathcal{S}_k = \{T \subseteq N : |T| = k\}$ to maximize
  $\sum_{i \in S} \deg(i) = \sum_{i \in S} |\{j \in N : (j, i) \in E\}|$
- Mechanism $M : \mathcal{G} \to \Delta(\mathcal{S}_k)$
- Strategyproofness: probability of selecting i independent of edges (i, j) for $j \in N$
- α-efficiency: for every graph,

$$\frac{\max_{S \in \mathcal{S}_k} \sum_{i \in S} \deg(i)}{\mathbb{E}_{S \sim M}\left[\sum_{i \in S} \deg(i)\right]} \leq \alpha$$

upper bounds: mechanisms

lower bounds: impossibility results

Bad News

Theorem: Let $n \geq 2$, $k \leq n - 1$. Then there is no strategyproof and α-approximate deterministic mechanism for any finite α.

Proof: very cute

Bad News

Theorem: Let $n \geq 2$, $k \leq n - 1$. Then there is no strategyproof and α-approximate deterministic mechanism for any finite α.

Proof: very cute

Particularly surprising for $k = n - 1$: cannot guarantee to select unique agent receiving any votes

Bounds for Randomized Mechanisms

Theorem: There exists a class of mechanisms, parameterized by $m \geq 2$, that is (universally) strategyproof for all n, k, m,

- 4-approximate for $m = 2$, and
- $1 + O(1/k^{\frac{1}{3}})$-approximate for $m \approx k^{\frac{1}{3}}$.

Theorem: Let $n \geq 2$, $k \leq n - 1$. Then there is no strategyproof and α-approximate mechanism for $\alpha < 1 + \Omega(1/k^2)$.

Bounds for Randomized Mechanisms

Theorem: There exists a class of mechanisms, parameterized by $m \geq 2$, that is (universally) strategyproof for all n, k, m,

- 4-approximate for $m = 2$, and
- $1 + O(1/k^{\frac{1}{3}})$-approximate for $m \approx k^{\frac{1}{3}}$.

Theorem: Let $n \geq 2$, $k \leq n - 1$. Then there is no strategyproof and α-approximate mechanism for $\alpha < 1 + \Omega(1/k^2)$. 2 for $k = 1$

Thank you!