Online Optimization for Dynamic Matching Markets

Patrick Jaillet

Department of Electrical Engineering and Computer Science
Laboratory for Information and Decision Systems
Operations Research Center
Massachusetts Institute of Technology

November 8, 2016

Research funded in part by the U.S. National Science Foundation (NSF) and Office of Naval Research (ONR)
Collaborators: Itai Ashlagi (Stanford), Maximilien Burq (MIT), Xin Lu (Amazon), Vahideh Manshadi (Yale)

outline

1. dynamic matching: some contexts and background
   - sponsored search and ad displays
   - kidney exchange programs
   - online optimization framework

2. online bipartite matching problems

3. matching markets for kidney exchange

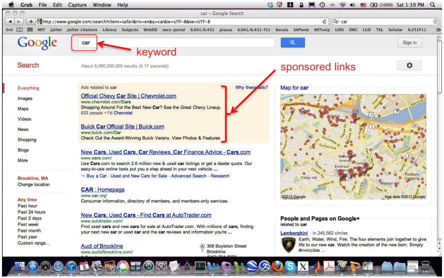

	sponsored search	display
marketer	get clicks and conversions	build brand awareness
publisher	search engine	(almost) any website
matching	publisher	third party

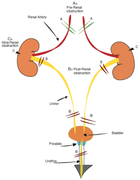

kidney transplant: best treatment for end stage renal disease.

- 99,379 patients on the waiting list in the US as of yesterday[1] [France: 21,000 people waiting for a kidney in 2015 (up from 12,000 in 2006)]
- over 30,000 new patients each year in the US
- only 18,000 transplants per year [France: 3,500 in 2015]
- about 12,000 from deceased donors [France: 3,000 in 2015]

[1] United Network for Organ Sharing (UNOS) http://www.unos.org/

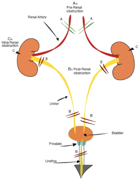

kidney transplant: best treatment for end stage renal disease.

- 99,379 patients on the waiting list in the US as of yesterday[1] [France: 21,000 people waiting for a kidney in 2015 (up from 12,000 in 2006)]
- over 30,000 new patients each year in the US
- only 18,000 transplants per year [France: 3,500 in 2015]
- about 12,000 from deceased donors [France: 3,000 in 2015]
- alternative: live donation

[1] United Network for Organ Sharing (UNOS) http://www.unos.org/

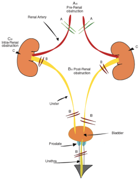

kidney transplant: best treatment for end stage renal disease.

- 99,379 patients on the waiting list in the US as of yesterday[1] [France: 21,000 people waiting for a kidney in 2015 (up from 12,000 in 2006)]
- over 30,000 new patients each year in the US
- only 18,000 transplants per year [France: 3,500 in 2015]
- about 12,000 from deceased donors [France: 3,000 in 2015]
- alternative: live donation

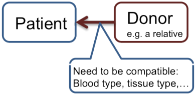

[1] United Network for Organ Sharing (UNOS) http://www.unos.org/

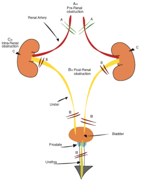

kidney transplant: best treatment for end stage renal disease.

- 99,379 patients on the waiting list in the US as of yesterday[1] [France: 21,000 people waiting for a kidney in 2015 (up from 12,000 in 2006)]
- over 30,000 new patients each year in the US
- only 18,000 transplants per year [France: 3,500 in 2015]
- about 12,000 from deceased donors [France: 3,000 in 2015]
- alternative: live donation

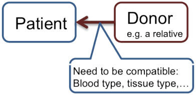

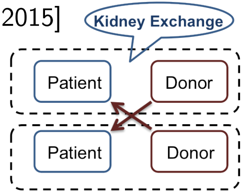

[1] United Network for Organ Sharing (UNOS) http://www.unos.org/

online concepts

- online optimization problem = instance incrementally revealed over time; need to make decisions "online" without knowledge about what comes next.
- online algorithm:
    - the "quality" of an online algorithm ALG is its competitive ratio c, measured against an optimal clairvoyant algorithm OPT with full knowledge of the instance; for a maximization problem:

$$c = \sup \{ r \mid \mathrm{alg(i)/opt(i)} \geq r, \ \forall \text{ instances i} \}$$

    - it is said to be best possible if there is no other such algorithm with a larger competitive ratio.

... online concepts ...

- randomized online algorithm: online algorithm with random steps designed to solve an online problem.
- competitive analysis depends on the "power of the offline adversary":
  - oblivious adversary: doesn't know the realizations of the random steps.
  - adaptive-online adversary: generate the instance adaptively based on past realizations of the random steps.
  - adaptive-offline adversary: perfect knowledge of all the realizations (past and future) of the random steps.
- we consider only oblivious adversary.

... online concepts ...

- competitive ratio of a randomized online algorithm against an oblivious adversary:

$$c = \sup \{r \mid \mathbb{E}[\mathrm{ALG(i)}]/\mathrm{opt(i)} \geq r, \ \forall \text{ instances i}\}$$

- a randomized online algorithm is said to be best possible if there is no other such algorithm with a larger competitive ratio

... online concepts

- in some cases uncertainty about the future input streams can be modeled using probabilistic concepts.
- how to properly include this "stochastic information" in order to design better algorithms? can this be quantified?
- look at the competitive ratio of an online algorithm defined as

$$c = \sup \{ r \mid \mathbb{E}_{\mathbb{P}}[\mathrm{ALG(I)}/\mathrm{OPT(I)}] \geq r \}$$

- sometimes can show results hold "with high probability" as opposed to simply "in expectation".

First Part: Online Bipartite Matching Problems

bipartite matching: classical and online version

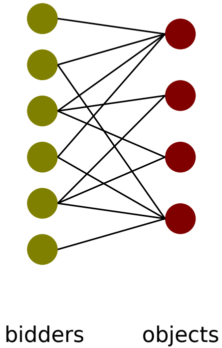

bipartite matching: classical and online version

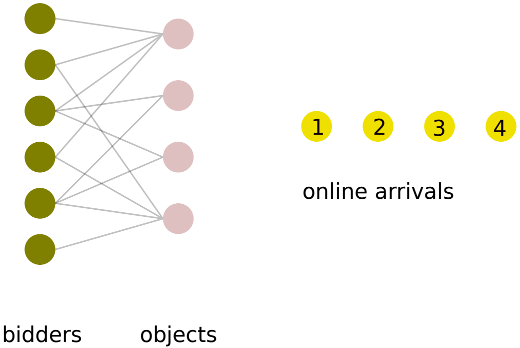

bipartite matching: classical and online version

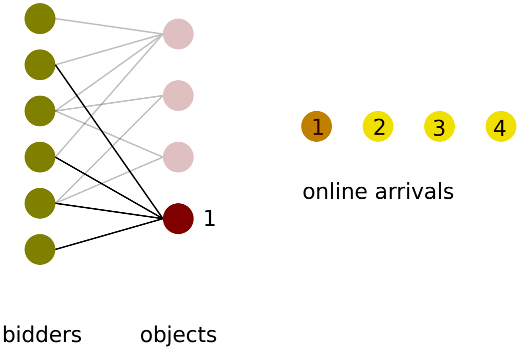

bipartite matching: classical and online version

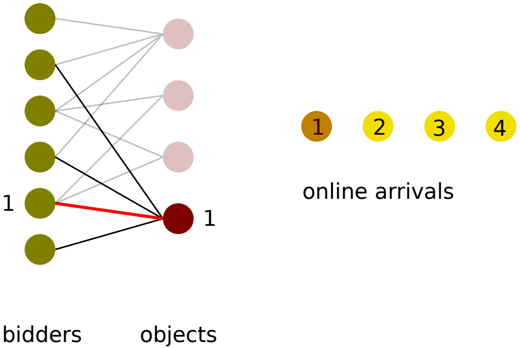

bipartite matching: classical and online version

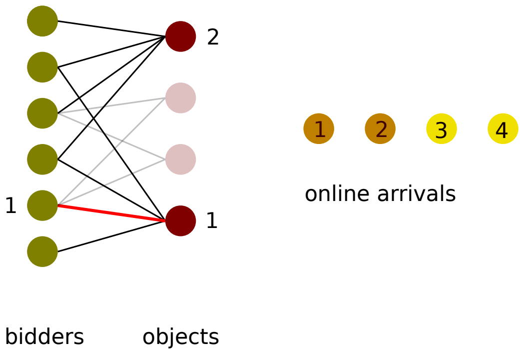

bipartite matching: classical and online version

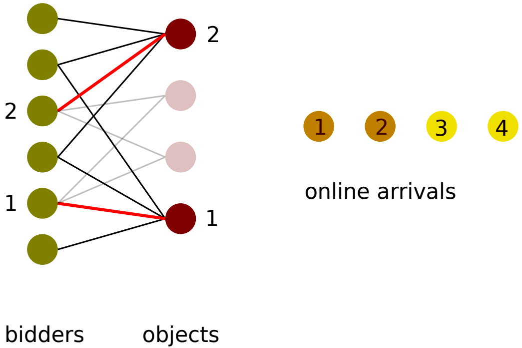

bipartite matching: classical and online version

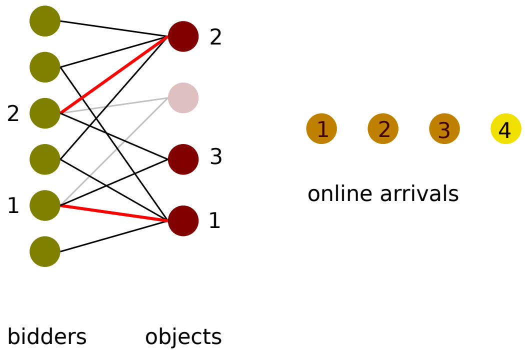

bipartite matching: classical and online version

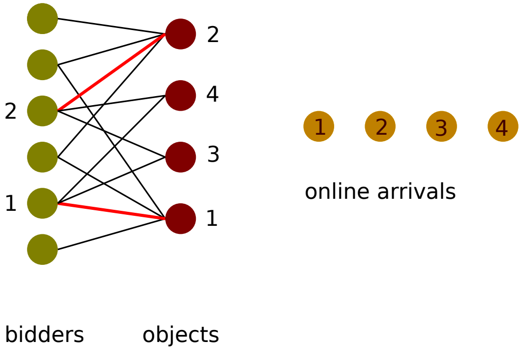

selected work

- KVV90: Karp, Vazirani, Vazirani, [STOC 1990] [online bipartite matching]
- FMMM09: Feldman, Mehta, Mirrokni, Muthukrishnan, [FOCS 2009] [online stochastic matching under i.i.d. integral arrival rates]
- MY11: Mahdian, Yan, [STOC 2011] [online stochastic matching under unknown distribution]
- KMT11: Karande, Mehta, Tripathi, [STOC 2011] [online stochastic matching under unknown distribution]
- MOS13: Manshadi, Oveis-Gharan, Saberi, [MathOR 2013] [online stochastic matching under i.i.d. general arrival rates]
- JL14: J., Lu, [MathOR 2014] [online stochastic matching under various random models]

online bipartite matching: summary of main results

- for the adversarial model:
  - no deterministic online algorithms can do better than 0.5
  - the *ranking* algorithm of KVV90 provides a best possible randomized algorithm with competitive ratio $1 - 1/e \approx 0.632$

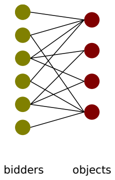

online bipartite matching: summary of main results

- for the adversarial model:
  - no deterministic online algorithms can do better than 0.5
  - the *ranking* algorithm of KVV90 provides a best possible randomized algorithm with competitive ratio $1 - 1/e \approx 0.632$

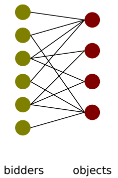

- what about the case of random inputs, i.i.d. with known distribution?

online bipartite matching: summary of main results

- for the adversarial model:
  - no deterministic online algorithms can do better than 0.5
  - the *ranking* algorithm of KVV90 provides a best possible randomized algorithm with competitive ratio $1 - 1/e \approx 0.632$

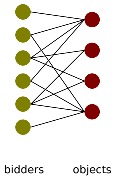

- what about the case of random inputs, i.i.d. with known distribution?
  - FMMM09 prove for the first time that one can do better under an i.i.d. stochastic model and give a 0.670-competitive algorithm under this scenario.
  - MOS13 provide a 0.702-competitive algorithm
  - in JL14, we obtain a $(1 - 2/e^2 \approx 0.729)$-competitive algorithm, best-known so far

an overall approach

- given an i.i.d. model with known distribution, by solving an optimal maximum cardinality matching for each possible i.i.d. draws, one could "calculate" p^*_{ik} the probability that edge (i, k) is part of an optimal solution in any given random realization ...

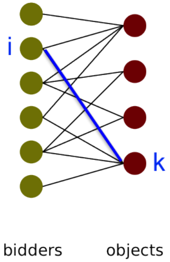

an overall approach

- given an i.i.d. model with known distribution, by solving an optimal maximum cardinality matching for each possible i.i.d. draws, one could "calculate" p^*_{ik} the probability that edge (i, k) is part of an optimal solution in any given random realization ...
- ... instead of computing p^*_{ik}, we formulate special maximum flow problems whose optimal solutions provide the input for the design of good online algorithms.

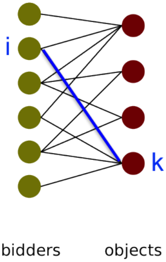

a maximum flow problem for the case of i.i.d. uniform

under the i.i.d. uniform assumption one can show

$$p^*_{ik} \leq 1 - 1/e \approx 0.632 \; \forall (i,k) \in E.$$

$$\begin{aligned}
\max \quad & \sum_{(i,k)\in E} x_{ik} \\
s.t. \quad & \sum_{k:(i,k)\in E} x_{ik} \leq 1 \quad \forall i \in B \\
& \sum_{i:(i,k)\in E} x_{ik} \leq 1 \quad \forall k \in O \\
& x_{ik} \in [0, 2/3] \qquad \forall (i.k) \in E
\end{aligned}$$

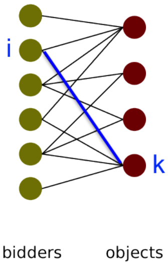

a "light" version of the 0.729-competitive algorithm

our randomized online algorithms (a bit simplified)

1. let $\mathbf{x}^*$ be an optimal solution to the previous LP.
2. When an object of type k arrives,
   - if all bidders i such that $x^*_{ik} > 0$ have already been matched, then drop the request;
   - otherwise, randomly assign the request to one of these unmatched bidders, proportionally to x^*_{ik}

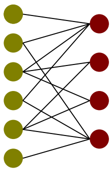

a "light" version of the 0.729-competitive algorithm

our randomized online algorithms (a bit simplified)

1. let $\mathbf{x}^*$ be an optimal solution to the previous LP.
2. When an object of type k arrives,
   - if all bidders i such that $x^*_{ik} > 0$ have already been matched, then drop the request;
   - otherwise, randomly assign the request to one of these unmatched bidders, proportionally to x^*_{ik}

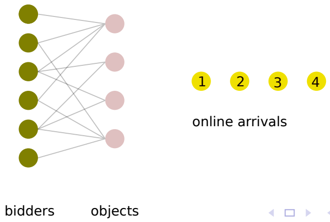

a "light" version of the 0.729-competitive algorithm

our randomized online algorithms (a bit simplified)

1. let $\mathbf{x}^*$ be an optimal solution to the previous LP.
2. When an object of type k arrives,
   - if all bidders i such that $x^*_{ik} > 0$ have already been matched, then drop the request;
   - otherwise, randomly assign the request to one of these unmatched bidders, proportionally to x^*_{ik}

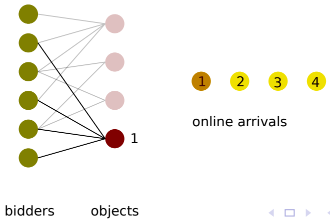

a "light" version of the 0.729-competitive algorithm

our randomized online algorithms (a bit simplified)

1. let $\mathbf{x}^*$ be an optimal solution to the previous LP.
2. When an object of type k arrives,
   - if all bidders i such that $x^*_{ik} > 0$ have already been matched, then drop the request;
   - otherwise, randomly assign the request to one of these unmatched bidders, proportionally to x^*_{ik}

$\Rightarrow$ competitive ratio $= 1 - 2/e^2$

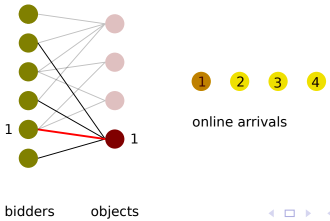

Second Part: Matching Markets for Kidney Exchange

kidney exchange

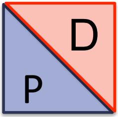

kidney exchange

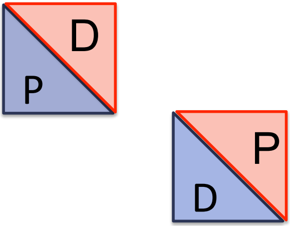

kidney exchange

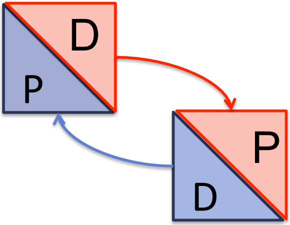

kidney exchange

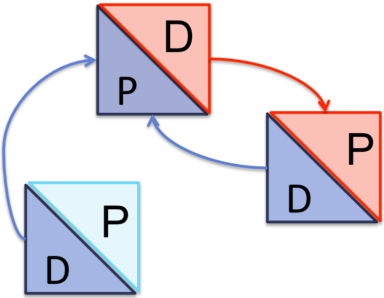

kidney exchange

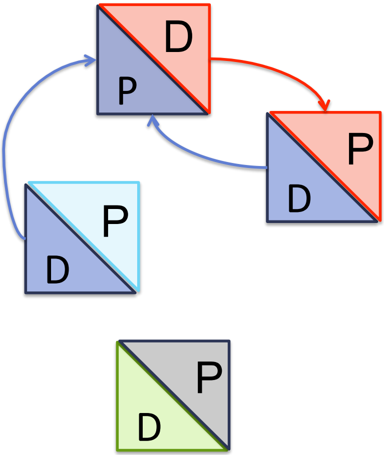

kidney exchange

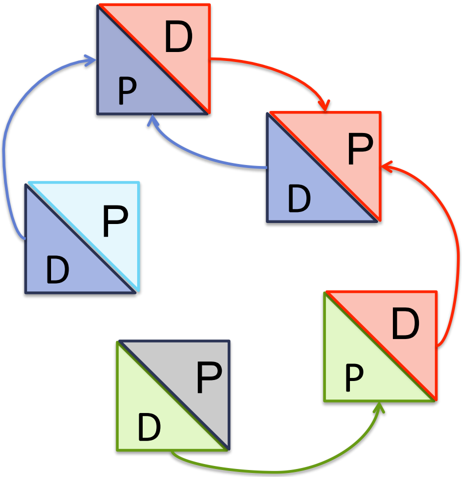

data

- blood-type compatibility:

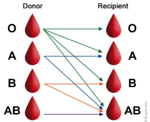

- tissue-type compatibility: the likelihood of being compatible with a random donor (from a compatible blood type) is measured by the PRA (panel reactive antibodies): low PRA = highly compatible.

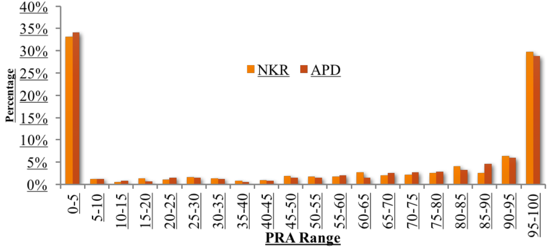

data and model - thickness and heterogeneity

- thickness: how easy is it for a typical patient to get a "match"?
- tissue-type compatibility:
  - high PRA = "hard" to match
  - low PRA = "easy" to match

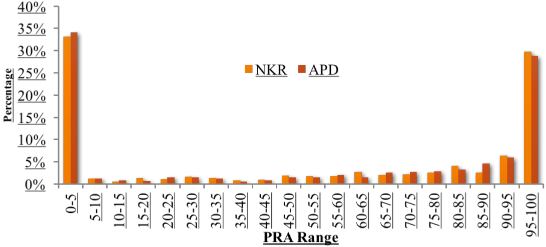

- bimodal distribution ⇒ two-type model:
  - hard-to-match (H): typical probability $p_H = 2.5\%$
  - easy-to-match (E): typical probability $p_E = 90\%$

main questions

- how to match in an heterogeneous environment?
- how do myopic (greedy in time) algorithms perform?
- how does the composition of the pool impact the types of exchanges needed to achieve the best performance?

main questions

- how to match in an heterogeneous environment?
- how do myopic (greedy in time) algorithms perform?
- how does the composition of the pool impact the types of exchanges needed to achieve the best performance?

main questions

- how to match in an heterogeneous environment?
- how do myopic (greedy in time) algorithms perform?
- how does the composition of the pool impact the types of exchanges needed to achieve the best performance?

selected work

- RSU07: Roth, Sönmez, Ünver, 2007, [compatibility-based static matching]
- U10: Ünver, 2010, [dynamic kidney exchange, blood-type compatibility, homogeneous case]
- AR11: Ashlagi, Roth, 2011, compatibility-based static matching]
- AGRR12: Ashlagi, Gamarnik, Rees, Roth, 2012 [compatibility-based static matching]
- AAKG13: Anderson, Ashlagi, Kanoria, Gamarnik, 2013, [compatibility-based dynamic matching, homogeneous case]
- ALO14: Akbarbour, Li, Oveis Gharan, 2014,[compatibility-based dynamic matching, with departure]
- ABJM16: Ashlagi, Burq, J., Manshadi, 2016, [compatibility-based dynamic matching, heterogeneous case]

modeling assumptions

- an incompatible patient-donor pair is represented as a node.
- nodes arrive over time: one node arrival at each time step.
- arriving node has type H with probability θ (and type E w.p. $1-\theta$).

modeling assumptions

- an incompatible patient-donor pair is represented as a node.
- nodes arrive over time: one node arrival at each time step.
- arriving node has type H with probability θ (and type E w.p. $1-\theta$).
- compatibility between nodes is modeled as a random graph, where the probability of an arc depends on the type of the receiving node:
  - $\mathbb{P}_{H\to H} = \mathbb{P}_{E\to H} = p_H = o(1)$.
  - $\mathbb{P}_{H\to E} = \mathbb{P}_{E\to E} = p_E = \Theta(1)$.

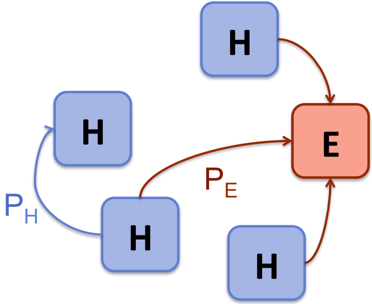

modeling assumptions

- an incompatible patient-donor pair is represented as a node.
- nodes arrive over time: one node arrival at each time step.
- arriving node has type H with probability θ (and type E w.p. $1-\theta$).
- compatibility between nodes is modeled as a random graph, where the probability of an arc depends on the type of the receiving node:
  - $\mathbb{P}_{H\to H} = \mathbb{P}_{E\to H} = p_H = o(1)$.
  - $\mathbb{P}_{H\to E} = \mathbb{P}_{E\to E} = p_E = \Theta(1)$.

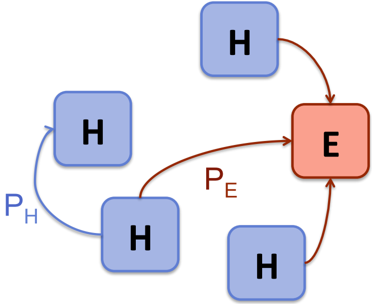

- no endogenous departures.

matching technology options

- in real kidney exchanges, matches are conducted over time:
  - bilaterally,
  - in multi-way cycles,
  - through chains.

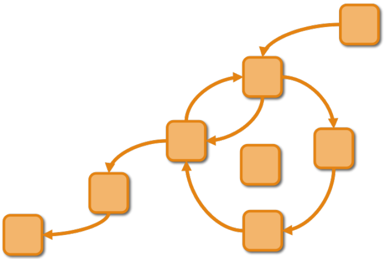

- in our model, we only consider:
  - bilateral matchings
  - chain matchings

our measure of efficiency

average waiting times

- infinite horizon, dynamic system, steady-state behavior.
- everyone eventually gets matched
- the E nodes typically wait very little
- $\Rightarrow$ objective : minimize w_H, the average waiting time for H nodes in steady-state.

our measure of efficiency

average waiting times

- infinite horizon, dynamic system, steady-state behavior.
- everyone eventually gets matched
- the E nodes typically wait very little
- $\Rightarrow$ objective : minimize w_H, the average waiting time for H nodes in steady-state.

note: by Little's law, if n_H is the average number of H nodes waiting in steady state, then: $n_H = \theta w_H$.

lower bound on best-possible algorithms

Proposition 1 (ABJM16)

- *under any matching policy that reaches a steady state, average waiting times are such that* $w_H + w_E = \Omega(\frac{1}{p_H})$.

intuition: by contradiction: if the pool size is too small, an arriving node has a small probability of being matched immediately, and must wait a "long" time to get at least an incoming arc; by Little's law, this in turn would imply a "large" pool size.

lower bound on best-possible algorithms

Proposition 1 (ABJM16)

- *under any matching policy that reaches a steady state, average waiting times are such that* $w_H + w_E = \Omega(\frac{1}{p_H})$.

intuition: by contradiction: if the pool size is too small, an arriving node has a small probability of being matched immediately, and must wait a "long" time to get at least an incoming arc; by Little's law, this in turn would imply a "large" pool size.

Proposition 2 (ABJM16)

- *under any bilateral matching policy that reaches a steady state, the average waiting time of node* H *is such that* $w_H = \Omega(\frac{1}{p_H^2})$

intuition: main idea is to show that a significant fraction of H nodes have to match to each other as a necessary condition for steady-state.

bilateral matching

BilateralMatch-H algorithm:

- only 2-cycles are considered.
- matches are conducted as soon as possible (myopic policy).
- in case of ties, priority is given to H agents.

properties:

- arcs present at t are never used for a matching at a time $> t$.
- system is characterized by only $N_t = (N_H(t), N_E(t))$.
- N_t is a positive recurrent Markov Chain.

BilateralMatch-H performance

Theorem 1 (ANJM16)

- *if* $\theta < 1/2$, *BilateralMatch-H achieves a waiting time* $w_H = \Theta\left(\frac{1}{p_H}\right)$
- *if* $\theta \geq 1/2$, *BilateralMatch-H leads to much larger waiting time* $w_H = \Theta\left(\frac{1}{p_H^2}\right)$

BilateralMatch-H performance

Theorem 1 (ANJM16)

- *if $\theta < 1/2$, BilateralMatch-H achieves a waiting time $w_H = \Theta\left(\frac{1}{p_H}\right)$*
- *if $\theta \geq 1/2$, BilateralMatch-H leads to much larger waiting time $w_H = \Theta\left(\frac{1}{p_H^2}\right)$*

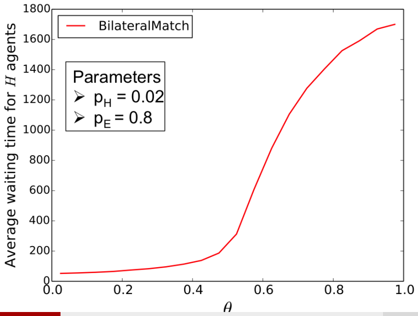

intuition of the proof

case $\theta < 1/2$: more E than H nodes:

- when there are a lot of E nodes, H nodes can easily match to them bilaterally; the probability of that happening is p_H for each H node waiting in the system.
- $\Rightarrow$ there needs to be only $1/p_H$ nodes waiting for one to be matched with high probability at every time step ...

case $\theta > 1/2$: more H than E nodes:

- when there are few E nodes however, some H nodes have to match to each other; the probability of this happening drops to p_H^2
- $\Rightarrow$ there needs to be $1/p_H^2$ nodes waiting for one to be matched with high probability at every time step ...

priorities - what about *BilateralMatch-E* performance?

Theorem 1 also holds for a policy that prioritizes for E nodes instead of H nodes (proved when $p_E = 1$, conjectured when $p_E < 1$)

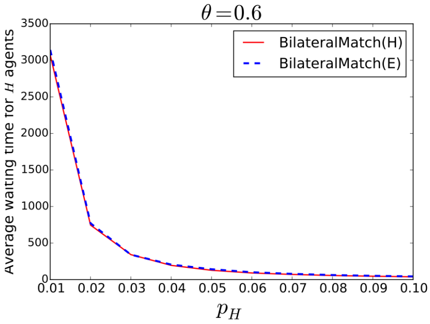

chain matching

motivations

- "non-directed" donors allow for chain matching.
- very good performance in static systems.
- very widely used in kidney exchange platforms.

known results

- chains can be used in dynamic matching [Anderson et al.].
- often at high computation cost (longest chains in the graph).

our results

- extension to heterogeneous systems.
- good performance with myopic (greedy in time and in chain exploration) policy.
- heterogeneity helps.

chain matching

ChainMatch algorithm:

- arrival of an altruistic agent.

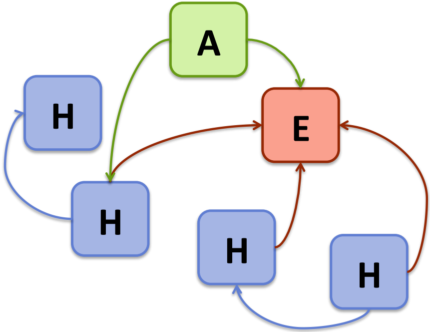

chain matching

ChainMatch algorithm:

- arrival of an altruistic agent.
- priority for H agents in case of a tie.

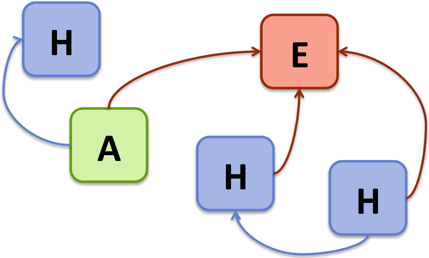

chain matching

ChainMatch algorithm:

- arrival of an altruistic agent.
- priority for H agents in case of a tie.
- we match agents as soon as possible.

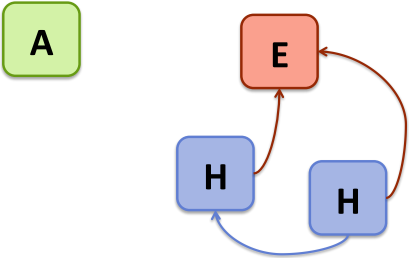

chain matching

ChainMatch algorithm:

- arrival of an altruistic agent.
- priority for H agents in case of a tie.
- we match agents as soon as possible.
- the next agent is chosen randomly

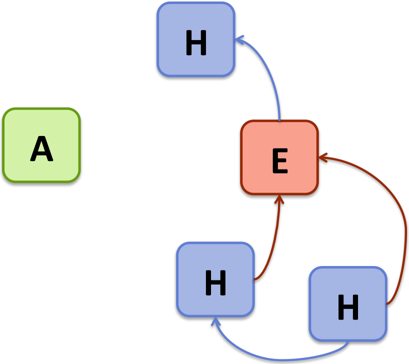

chain matching

ChainMatch algorithm:

- arrival of an altruistic agent.
- priority for H agents in case of a tie.
- we match agents as soon as possible.
- the next agent is chosen randomly

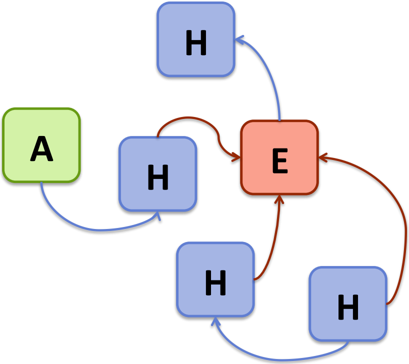

chain matching

ChainMatch algorithm:

- arrival of an altruistic agent.
- priority for H agents in case of a tie.
- we match agents as soon as possible.
- the next agent is chosen randomly

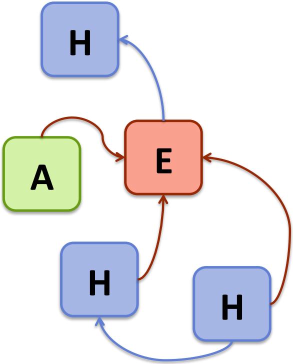

chain matching

ChainMatch algorithm:

- arrival of an altruistic agent.
- priority for H agents in case of a tie.
- we match agents as soon as possible.
- the next agent is chosen randomly

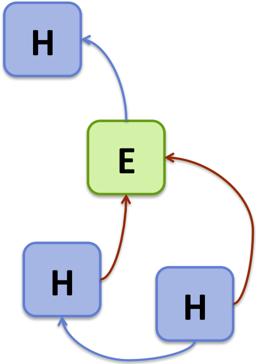

chain matching

ChainMatch **algorithm:**

- arrival of an altruistic agent.
- priority for H agents in case of a tie.
- we match agents as soon as possible.
- the next agent is chosen randomly

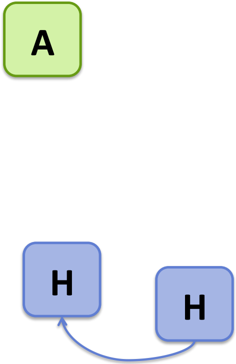

properties:

- system is defined by $(N_H(t), N_E(t))$.
- Markov property.

ChainMatch performance

Theorem 2 (ANJM16)

in the special case $p_E = 1$,

- *if $\theta < 1$, ChainMatch achieves a waiting time $w_H = \Theta\left(\frac{1}{p_H}\right)$*
- *if $\theta = 1$, ChainMatch achieves a waiting time $w_H = O\left(\frac{\ln(1/p_H)}{p_H}\right)$*

summary results: chain vs bilateral

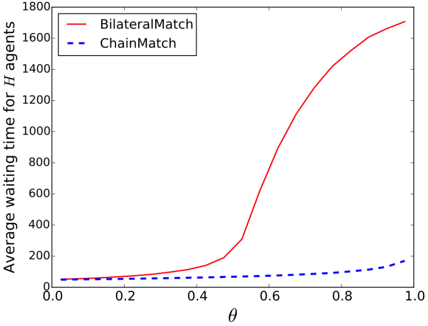

Parameters of the simulation

- 100,000 iterations
- $p_H = 0.02$
- $p_E = 0.8$

discussion

performance of myopic algorithms

- batching does not help reduce waiting times.
- greedy chains perform as well as longest chains.

value of heterogeneity.

- for $\theta < 1$, waiting times decrease as $\frac{1}{1-\theta}$.
- having at least 50% of E agents allows for near-optimal performance using bilateral matching.

takeaways

benefits of heterogeneity

- tradeoff between complexity of the matching mechanism and performance.
- simpler chain-matching algorithm while keeping optimal performances.

Thank You !

Online Optimization for Dynamic Matching Markets

Patrick Jaillet

Department of Electrical Engineering and Computer Science
Laboratory for Information and Decision Systems
Operations Research Center
Massachusetts Institute of Technology

November 8, 2016

Research funded in part by the U.S. National Science Foundation (NSF) and Office of Naval Research (ONR)

Collaborators: Itai Ashlagi (Stanford), Maximilien Burq (MIT), Xin Lu (Amazon), Vahideh Manshadi (Yale)