When LP is the Cure for Your Matching Woes: Improved Bounds for Stochastic Matchings

Jian Li, University of Maryland, College Park

Joint work with Nikhil Bansal (IBM), Anupam Gupta (CMU),

Julian Mestre (MPI), Viswanath Nagarajan (IBM), Atri Rudra (SUNY-Buffalo)

Problem Definition

Stochastic Matching [Chen, Immorlica, Karlin, Mahdian, and Rudra. '09]

- Given:
  - A probabilistic graph $G(V,E)$.
  - Existential prob. p_e for each edge e.
  - Patience level t_v for each vertex v.
- **Probing** $e=(u,v)$: The only way to know the existence of e.
  - We can probe (u,v) only if $t_u>0, t_v>0$.
  - If e indeed exists, we should add it to our matching.
  - If not, $t_u = t_u-1$, $t_v = t_v-1$.

Problem Definition

- Output: A strategy to probe the edges
  - Edge-probing: an (adaptive or non-adaptive) ordering of edges.
  - Matching-probing: *k* rounds; In each round, probe a matching.

- Objectives:
  - Unweighted: Max. *E[cardinality of the matching]*.
  - Weighted: Max. *E[weight of the matching]*.

Motivations

- **Online dating**
  - Existential prob. p_e : estimation of the success prob. based on users' profiles.

MOTIVATIONS

- **Online dating**
  - Existential prob. p_e : estimation of the success prob. based on users' profiles.
  - Probing edge $e=(u,v)$: u and v are sent to a date.

MOTIVATIONS

- **Online dating**
  - Existential prob. p_e : estimation of the success prob. based on users' profiles.
  - Probing edge $e=(u,v)$: u and v are sent to a date.
  - Patience level: obvious.

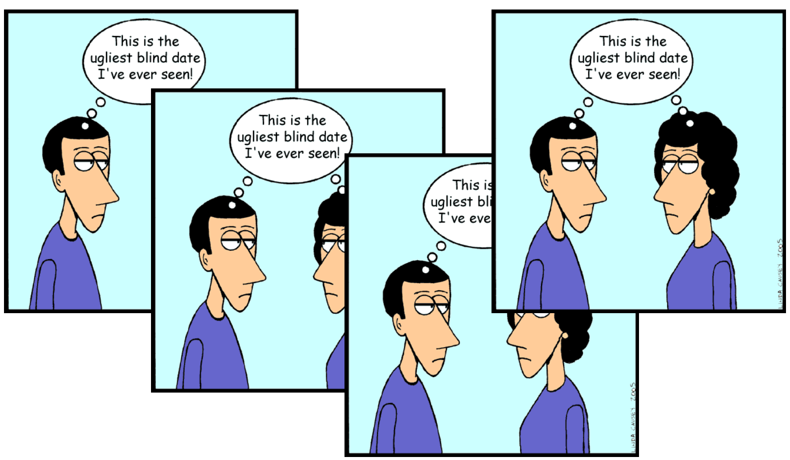

Motivations

- **Kidney exchange**
  - Existential prob. p_e : estimation of the success prob. based on blood type etc.
  - Probing edge $e=(u,v)$: the crossmatch test (which is more expensive and time-consuming).

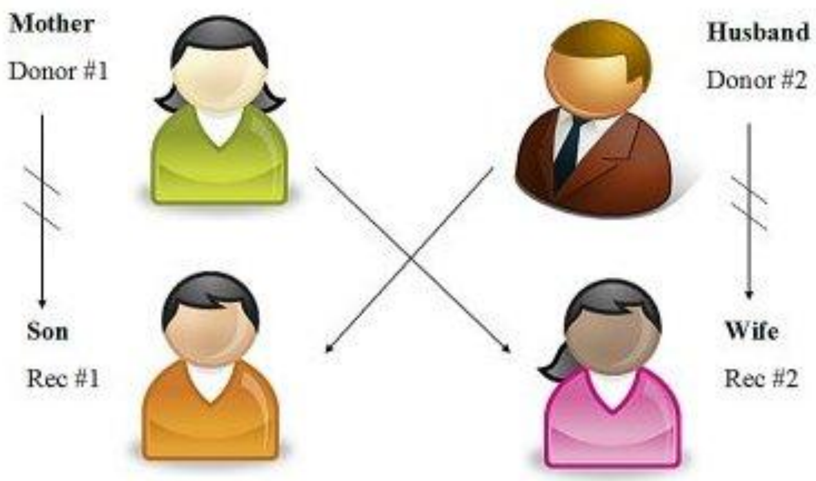

OUR RESULTS

- Previous results for unweighted version [Chen et al. '09]:
  - Edge-probing: Greedy is a 4-approx.
  - Matching-probing: O(log n)-approx.
- A simple 8-approx. for weighted stochastic matching.
  - For edge-probing model.
  - can be improved to 5.75 by a more careful analysis.
- An improved 3-approx. for bipartite graphs and 4-approx. for general graphs based on dependent rounding [Gandhi et al. '06].
  - For both edge-probing and matching-probing models.
  - This implies the gap between the best matching-probing strategy and the best edge-probing strategy is a small const.

OTHER RESULTS

- Stochastic online matching.
  - A set of items and a set of buyer types. A buyer of type *b* likes item *a* with probability p_{ab}. (G(items, buyer types): Expected graph)
  - The buyers arrive online (her type is an i.i.d. r.v.).
  - The algorithm shows the buyer (of type *b*) at most t_b items one by one.
  - The buyer buys the first item she likes or leaves without buying.
  - This generalizes the stochastic online matching problem of [Feldman et al. '09, Bahmani et al. '10, Saberi et al '10] where $p_e=\{0,1\}$.

  - We have a 7.92-approximation.

OTHER RESULTS

- Cardinality Constrained Matching in Rounds.
  - In each round, we can probe a matching of size ≤C.
  - An O(1)-approx.
  - Chen et al. obtained an O(min(k,C))-approx.
- A new proof for greedy.
  - An simple LP-based analysis: 5-approx.
  - The analysis by Chen et al. was based on decision trees.

Other Results

- Stochastic *k*-set packing.
  - Generalizing the stochastic matching problem.
  - *k=4.*

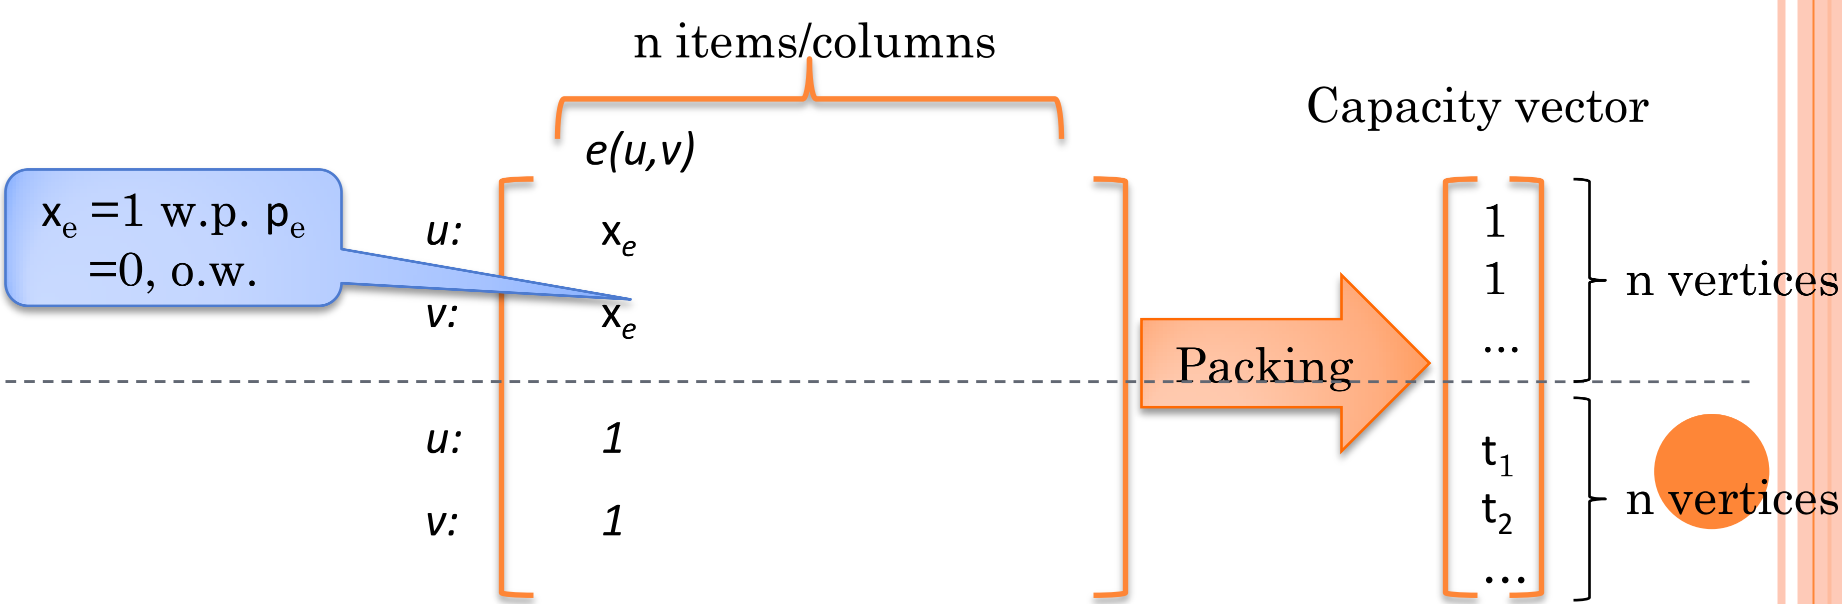

APPROXIMATION RATIO

- We compare our solution against the optimal (adaptive) strategy (not the offline optimal solution).

- An example:

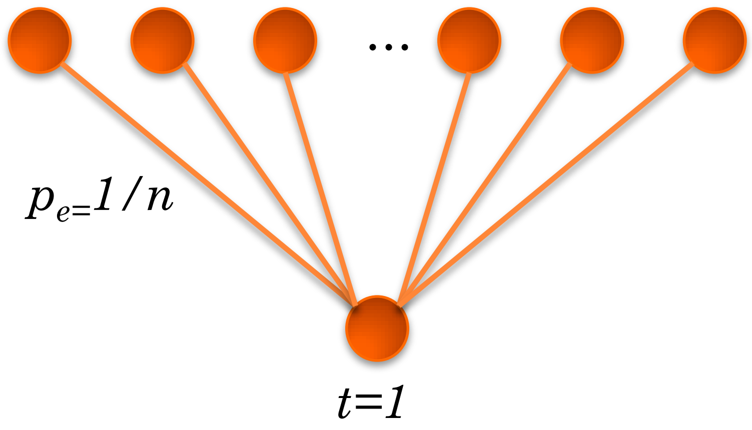

E[offline optimal] = $1-(1-1/n)^n \approx 1-1/e$

E[any algorithm] = $1/n$

A LP Upper Bound

- Variable y_e : Prob. that any algorithm probes *e.*

$$
\begin{aligned}
\text{maximize} \quad & \sum_{e \in E} w_e \cdot x_e \\
\text{subject to} \quad & \sum_{e \in \partial(v)} x_e \cdot 1 \quad \forall v \in V \\
& \sum_{e \in \partial(v)} y_e \cdot t_v \quad \forall v \in V \\
& x_e = p_e \cdot y_e \quad \forall e \in E \\
& 0 \cdot y_e \cdot 1 \quad \forall e \in E
\end{aligned}
$$

At most 1 edge in $\partial(v)$ is matched

At most t_v edges in $\partial(v)$ are probed

x_e: Prob. e is matched

A Simple 8-Approximation

An edge (u,v) is *safe* if $t_u>0, t_v>0$ and neither u nor v is matched

Algorithm:

- Pick a permutation π on edges uniformly at random
- For each edge e in the ordering π, do:
  - If e is not safe then do not probe it.
  - If e is safe then probe it w.p. y_e/α.

A SIMPLE 8-APPROXIMATION

Analysis:

Lemma: For any edge *(u,v)*, at the point when *(u,v)* is considered under *π*, *Pr(u loses its patience) ≤1/2α* .

Proof: Let *U* be #probes incident to u and before *e*.

$$\mathbb{E}[U] = \sum_{e \in \partial(u)} \Pr[\,\text{edge } e \text{ appears before } (u,v) \text{ in } \pi \text{ AND } e \text{ is probed}\,],$$

$$\cdot \sum_{e \in \partial(u)} \Pr[\,\text{edge } e \text{ appears before } (u,v) \text{ in } \pi\,] \cdot \frac{y_e}{\alpha},$$

$$= \sum_{e \in \partial(u)} \frac{y_e}{2\alpha} \cdot \frac{t_u}{2\alpha}.$$

By the Markov inequality: $\Pr[\,U \geq t_u\,] \cdot \frac{\mathbb{E}[U]}{t_u} \cdot \frac{1}{2\alpha}.$

A Simple 8-Approximation

Analysis:

Lemma: For any edge e=*(u,v)*, at the point when *(u,v)* is considered under π, *Pr(u is matched)* $\le 1/2\alpha$.

Theorem: The algorithm is a 8-approximation.

Proof: When e is considered,

Pr(e is not safe) $\le$ *Pr(u is matched)+ Pr(u loses its patience)+*

Pr(v is matched)+ Pr(v loses its patience) $\le 2/\alpha$

Therefore,

$E[\text{our solution}] = \Sigma_e w_e Pr(e \text{ is safe}) (y_e/ \alpha) p_e$

$\ge (1-2/\alpha) (1/ \alpha) \Sigma_e w_e y_e p_e \ge 1/8 \, OPT$ $(\alpha=4)$

Recall $\Sigma_e w_e y_e p_e$ is an upper bound of *OPT*

An Improved Approx. – Bipartite Graphs

Algorithm:

- *(x,y)* ← Optimal solution of the LP.
- *y'* ← Round *y* to an integral solution using *dependent rounding* [Gandhi et al. 06] *and Let* $E' = \{e \mid y'_e = 1\}$.
  - (Marginal distribution) $\Pr(y'_e=1)=y_e$;
  - (Degree preservation) $Deg_{E'}(v) \le t_v$; (Recall $\Sigma_{e \in \partial(\varpi)} y_e \le t_v$)
  - (Negative Correlation) $\Pr(\Lambda_{e\in \mathcal{S}} (y'_e=1)) \le \Pi_{e\in \mathcal{S}} y_e$.
- Probe the edges in *E'* in random order.

For matching-probe model:

- $M_1, \ldots, M_h$ ← Optimal edge coloring of *E'*.
- Probe $\{M_1, \ldots, M_h\}$ in random order.

Final Remarks and Open Questions

- Quite recently, Adamczyk has proved that the greedy algorithm is a 2-approximation for the unweighted version.
- Better approximations? (Unweighted: 2; Weighted bipartite: 3; Weighted: 4).
- o(k)-approximation for stochastic *k*-set packing? Or θ(k) is the best possible?
- Any lower bound?

Thanks

An Improved Approx. – Bipartite Graphs

Analysis (sketch):

Assume we have chosen E'.

Consider a particular edge *e=(u,v).*

Let *B(e,π)* be the set of incident edges that appear before *e in the random order π*.

$$\Pr[\, e \text{ is safe} \mid E' \,] \geq \mathbb{E}_\pi\Big[\prod_{f \in B(e,\pi)} (1-p_f) \mid E' \Big];$$

We claim that

$$\mathbb{E}_\sigma\Big[\prod_{f \in B(e,\sigma)} (1-p_f) \mid E' \Big] = \int_0^1 \prod_{f \in \partial_{E'}(e)} (1 - x p_f)\, \mathrm{d}x.$$

An Improved Approx. – Bipartite Graphs

Analysis cont: To see $\mathbb{E}_\sigma\Big[\prod_{f\in B(e,\sigma)}(1-p_f)\mid E'\Big] = \int_0^1 \prod_{f\in\partial_{E'}(e)}(1-xp_f)\,\mathrm{d}x.$

Consider this random experiment: For each edge in $\partial_{E'}(e)$, we pick a random real in [0,1]. This produces a uniformly random ordering. Let r.v. $A_f = (1-p_f)$ if f goes before e, and $A_f=1$ o.w.

Then, we consider $\mathbb{E}[\prod_{f\in\partial_{E'}(e)} A_f] = \int_0^1 \mathbb{E}\Big[\prod_{f\in\partial_{E'}(e)} A_f \mid a_e = x\Big]\,\mathrm{d}x$

An Improved Approx. – Bipartite Graphs

Analysis cont:

Define $\rho(r, p_{\max})$ to be the optimal value of

$$\begin{aligned} \text{maximize} \quad & = \int_0^1 \prod_{f \in \partial_{E'}(e)} (1 - x p_f) \, \mathrm{d}x \\ \text{subject to} \quad & \sum_{f \in \partial_{E'}(e)} p_f \cdot r \\ & 0 \cdot p_f \cdot p_{\max}. \end{aligned}$$

We can show $\rho(r, p_{\max})$ is convex and decreasing on *r*

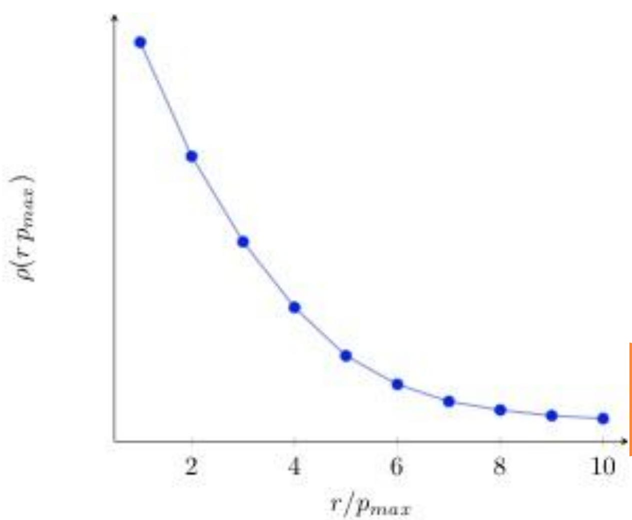

An Improved Approx. – Bipartite Graphs

Analysis cont:

$$\mathbb{E}\left[\mathsf{ALG}\right] = \sum w_e\, p_e\, \Pr[e \in E'] \cdot \Pr\left[e \text{ was safe } \mid e \in \widehat{E}\right]$$

Marginal Prob.

vexity.

ion.

An Improved Approx. – General Graphs

Algorithm:

- *(x,y)* ← Optimal solution of the LP.
- Randomly partition vertices into *A* and *B.*
- Run the previous algorithm on the bipartite graph G(A,B).

Thm: It is a $2/\rho(1, p_{max})$-approximation.

If $p_{max} \to 1$, the ratio tends to 4. If $p_{max} \to 0$, the ratio tends to $2/(1-1/e) \approx 3.15$