Online Stochastic Ad Allocation

Vahab Mirrokni

Google Research, New York

Outline

- Overview: Related Work and Results.
- Online Stochastic Matching: iid with known distribution
- Online Stochastic Packing: Random order
- Online Generalized Assignment (with free disposal)

Online Ad Allocation

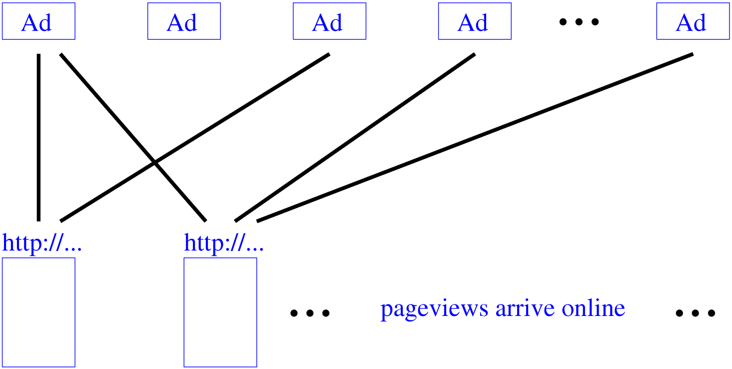

- When page arrives, assign an eligible ad.
  - value of assigning page i to ad a: v_{ia}
- Capacity of ad a: C_a

Online Ad Allocation

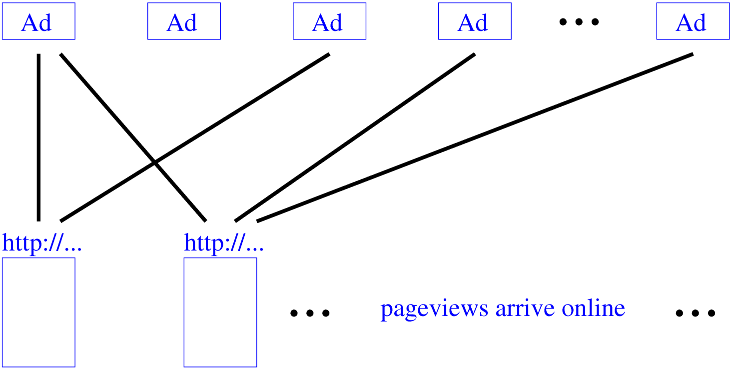

- When page arrives, assign an eligible ad.
  - value of assigning page i to ad a: v_{ia}
- Capacity of ad a: C_a
- Online Matching ($v_{ia} = 1$): Maximize number of ads served.
- Online Weighted Matching: Maximize value of ads served.

Online (Stochastic) Allocation: Known Results

- Arbitrary order:
  - Greedy is 0.5-approx.
  - [KVV90] Online $1 - 1/e \approx 0.632$ alg. This is tight.
  - [MSVV05,BJN07] Adwords (Weights, budgets): $1 - 1/e$-approx.

Online (Stochastic) Allocation: Known Results

- Arbitrary order:
  - Greedy is 0.5-approx.
  - [KVV90] Online $1 - 1/e \approx 0.632$ alg. This is tight.
  - [MSVV05,BJN07] Adwords (Weights, budgets): $1 - 1/e$-approx.
- Random order, or i.i.d with unknown distribution:
  - [GM08] Greedy is $1 - 1/e$ opt. 3/4-hard (5/6 rand).
  - Weights, budgets: greedy is $1 - 1/e$ opt.

Online (Stochastic) Allocation: Known Results

- Arbitrary order:
  - Greedy is 0.5-approx.
  - [KVV90] Online $1 - 1/e \approx 0.632$ alg. This is tight.
  - [MSVV05,BJN07] Adwords (Weights, budgets): $1 - 1/e$-approx.
- Random order, or i.i.d with unknown distribution:
  - [GM08] Greedy is $1 - 1/e$ opt. 3/4-hard (5/6 rand).
  - Weights, budgets: greedy is $1 - 1/e$ opt.
- i.i.d. model with known distribution:
  - [GM08] Greedy is $1 - 1/e$ opt.

Results: Three Recent Papers

- Online Stochastic Matching: Beating $1 - \frac{1}{e}$, FOCS 2009.
  - online stochastic matching in iid model with known dist.
  - 0.67-approximation (idea: power of two choices)
  - Feldman, Mehta, M., Muthukrishnan

Results: Three Recent Papers

Google

- Online Stochastic Matching: Beating $1 - \frac{1}{e}$, FOCS 2009.
  - online stochastic matching in iid model with known dist.
  - 0.67-approximation (idea: power of two choices)
  - Feldman, Mehta, M., Muthukrishnan

- Online Stochastic Packing applied to Display Ad Allocation, Arxive 2010.
  - Online stoch. packing in random order model: online routing.
  - $1 - \epsilon$-approximation under assumptions (idea: learn dual variables.)
  - Feldman, Henzinger, Korula, M., Stein

Results: Three Recent Papers

- Online Stochastic Matching: Beating $1 - \frac{1}{e}$, FOCS 2009.
  - online stochastic matching in iid model with known dist.
  - 0.67-approximation (idea: power of two choices)
  - Feldman, Mehta, M., Muthukrishnan

- Online Stochastic Packing applied to Display Ad Allocation, Arxive 2010.
  - Online stoch. packing in random order model: online routing.
  - $1 - \epsilon$-approximation under assumptions (idea: learn dual variables.)
  - Feldman, Henzinger, Korula, M., Stein

- Online Ad Assignment with Free Disposal, WINE 2009.
  - online generalized assignment problems with free disposal.
  - $1 - \frac{1}{e}$-competitive algorithm (idea: primal-dual analysis.)
  - Feldman, Korula, M., Muthukrishnan, Pal

Outline

- Overview: Related Work and Results.
- Online Stochastic Matching: iid with known distribution
- Online Stochastic Packing: Random order
- Online Generalized Assignment (with free disposal)

Online Stochastic Matching: iid (known dist.)

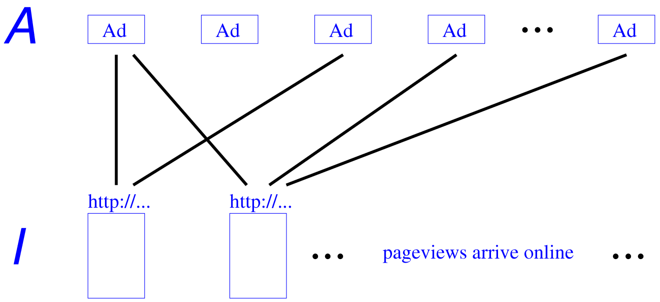

Given (offline):

- Bipartite graph $G = (A, I, E)$,
- Distribution D over I.

Online:

- n indep. draws from D.
- Must assign nodes upon arrival.

Online Stochastic Matching

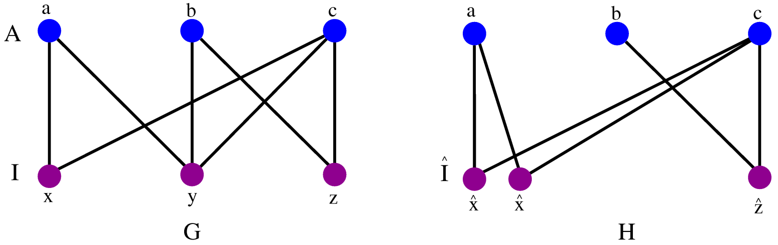

- Let $H = (A, \hat{I}, \hat{E})$ be the "realization" of G
  - (i.e., $\hat{I}$ are the nodes that actually arrive).

Online Stochastic Matching

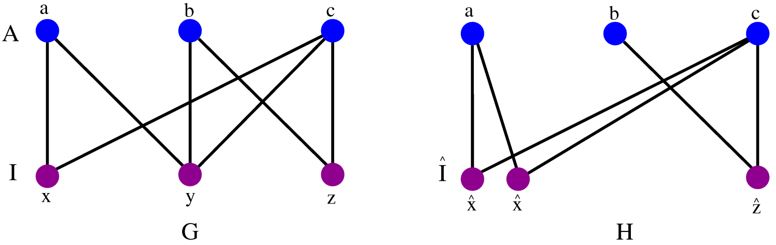

- Let $H = (A, \hat{I}, \hat{E})$ be the "realization" of G
  - (i.e., $\hat{I}$ are the nodes that actually arrive).
- Approximation ratio $= \mathrm{ALG}(H)/\mathrm{OPT}(H)$.

Online Stochastic Matching

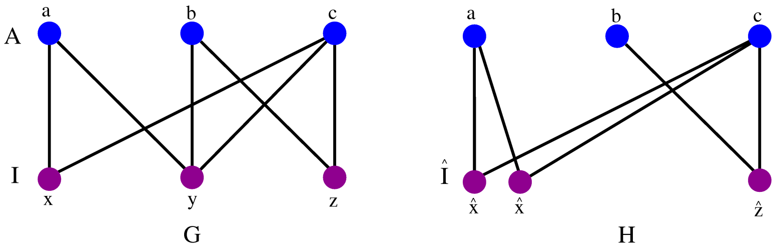

- Let $H = (A, \hat{I}, \hat{E})$ be the "realization" of G
  - (i.e., $\hat{I}$ are the nodes that actually arrive).
- Approximation ratio $= \mathrm{ALG}(H)/\mathrm{OPT}(H)$.
- "ALG is α-opt...?"

$$E\left[\frac{\mathrm{ALG}(H)}{\mathrm{OPT}(H)}\right] \geq \alpha \qquad \frac{E[\mathrm{ALG}(H)]}{E[\mathrm{OPT}(H)]} \geq \alpha \qquad \text{w.h.p.,} \frac{\mathrm{ALG}(H)}{\mathrm{OPT}(H)} \geq \alpha$$

Online Stochastic Matching

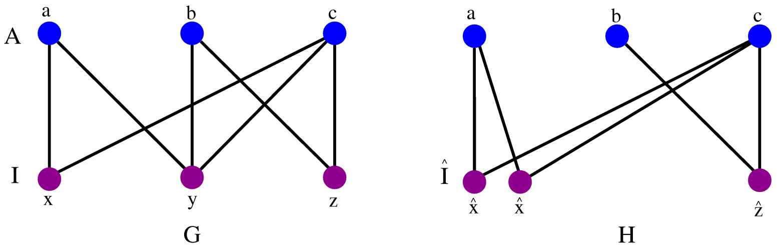

- Let $H = (A, \hat{I}, \hat{E})$ be the "realization" of G
  - (i.e., $\hat{I}$ are the nodes that actually arrive).
- Approximation ratio $= \mathrm{ALG}(H)/\mathrm{OPT}(H)$.
- "ALG is α-opt...?"

  $E\left[\frac{\mathrm{ALG}(H)}{\mathrm{OPT}(H)}\right] \geq \alpha \qquad \frac{E[\mathrm{ALG}(H)]}{E[\mathrm{OPT}(H)]} \geq \alpha \qquad$ w.h.p., $\frac{\mathrm{ALG}(H)}{\mathrm{OPT}(H)} \geq \alpha$
- This talk: D is uniform, $n = |A| = |I|$.

Results[FMMM09]: i.i.d with known distribution

- New algorithm for "online stochastic matching:"
  - Offline optimization on "expected graph"
  - Online: some ideas from "power of two choices" [Azar, Broder, Karlin, '99], [Mitzenmacher, '01]

Results[FMMM09]: i.i.d with known distribution

- New algorithm for "online stochastic matching:"
  - Offline optimization on "expected graph"
  - Online: some ideas from "power of two choices" [Azar, Broder, Karlin, '99], [Mitzenmacher, '01]
- With high probability,

$$\frac{\mathrm{ALG}}{\mathrm{OPT}} \geq \frac{1 - \frac{2}{e^2}}{\frac{4}{3} - \frac{2}{3e}} \simeq 0.67$$

Results[FMMM09]: i.i.d with known distribution

- New algorithm for "online stochastic matching:"
  - Offline optimization on "expected graph"
  - Online: some ideas from "power of two choices" [Azar, Broder, Karlin, '99], [Mitzenmacher, '01]
- With high probability,

$$\frac{\text{ALG}}{\text{OPT}} \geq \frac{1 - \frac{2}{e^2}}{\frac{4}{3} - \frac{2}{3e}} \simeq 0.67$$

- Analysis of algorithm is tight
  - $\exists$ example matching this bound.

Results[FMMM09]: i.i.d with known distribution

- New algorithm for "online stochastic matching:"
  - Offline optimization on "expected graph"
  - Online: some ideas from "power of two choices" [Azar, Broder, Karlin, '99], [Mitzenmacher, '01]
- With high probability,

$$\frac{\text{ALG}}{\text{OPT}} \geq \frac{1 - \frac{2}{e^2}}{\frac{4}{3} - \frac{2}{3e}} \simeq 0.67$$

- Analysis of algorithm is tight
  - $\exists$ example matching this bound.
- No algorithm can get $\frac{\text{ALG}}{\text{OPT}} \geq \frac{6e^3-23}{6e^3-22} \simeq .9898$.

Background: Balls in bins

- Suppose n balls thrown into n bins, i.i.d. uniform.

Background: Balls in bins

- Suppose n balls thrown into n bins, i.i.d. uniform.
- # non-empty bins concentrates:

Background: Balls in bins

- Suppose n balls thrown into n bins, i.i.d. uniform.
- # non-empty bins concentrates:
  - $B =$ particular subset of bins.

Background: Balls in bins

- Suppose n balls thrown into n bins, i.i.d. uniform.
- # non-empty bins concentrates:
  - $B =$ particular subset of bins.
  - $s = \#$ bins in B with ≥ 1 ball.

Background: Balls in bins

- Suppose n balls thrown into n bins, i.i.d. uniform.
- # non-empty bins concentrates:
  - $B =$ particular subset of bins.
  - $s = \#$ bins in B with ≥ 1 ball.
  - Then w.h.p., $s \approx |B|(1 - \frac{1}{e})$.

First Attempt: “Suggested matching”

1. Find a maximum matching in G.
2. Use that matching as nodes arrive online.

First Attempt: “Suggested matching”

1. Find a maximum matching in G.
2. Use that matching as nodes arrive online.

- Does no better than $1 - 1/e$.

First Attempt: “Suggested matching”

1. Find a maximum matching in G.
2. Use that matching as nodes arrive online.

- Does no better than $1 - 1/e$.
- Proof:
  - Suppose $G =$ complete graph.

First Attempt: "Suggested matching"

1. Find a maximum matching in G.
2. Use that matching as nodes arrive online.

- Does no better than $1 - 1/e$.
- Proof:
  - Suppose G = complete graph.
  - Then OPT(H) $= n$.

First Attempt: "Suggested matching"

1. Find a maximum matching in G.
2. Use that matching as nodes arrive online.

- Does no better than $1 - 1/e$.
- Proof:
  - Suppose $G =$ complete graph.
  - Then OPT(H) $= n$.
  - But w.h.p. only $1 - 1/e$ fraction of I will ever arrive.
    $\implies$ ALG $\approx (1 - 1/e)n$.

First Attempt: "Suggested matching"

1. Find a maximum matching in G.
2. Use that matching as nodes arrive online.

- Does no better than $1 - 1/e$.
- Proof:
  - Suppose $G =$ complete graph.
  - Then OPT(H) $= n$.
  - But w.h.p. only $1 - 1/e$ fraction of I will ever arrive.
    $\implies$ ALG $\approx (1 - 1/e)n$.
- In fact, this algorithm does achieve $1 - 1/e$ (in paper).

New ALG: “Two suggested matchings”

1. Offline: Find two disjoint matchings
2. Online: try the first one, then if that doesn’t work, try the second one.

New ALG: "Two suggested matchings"

Warmup: complete graph

- Two disjoint perfect matchings: blue (1-ary), red (2-ary).

New ALG: "Two suggested matchings"

Warmup: complete graph

- Two disjoint perfect matchings: blue (1-ary), red (2-ary).
- Union of matchings = cycles with alt. blue and red edges

New ALG: "Two suggested matchings"

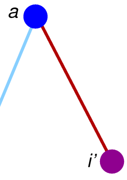

For particular node $a \in A$:

$$
\begin{aligned}
\Pr[a \text{ is chosen }] &\geq \Pr[i \text{ arrives once, or } i' \text{ arrives twice}] \\
&= 1 - \Pr[i \text{ never arrives \& } i' \text{ arrives } \leq \text{ once}] \\
&= 1 - \left((1-2/n)^n + n(1/n)(1-2/n)^{n-1}\right) \\
&\approx 1 - 2/e^2
\end{aligned}
$$

Thus, E[# nodes in A chosen] $\approx (1-2/e^2)n \approx .729n$
(This also concentrates...)

Algorithm (Offline)

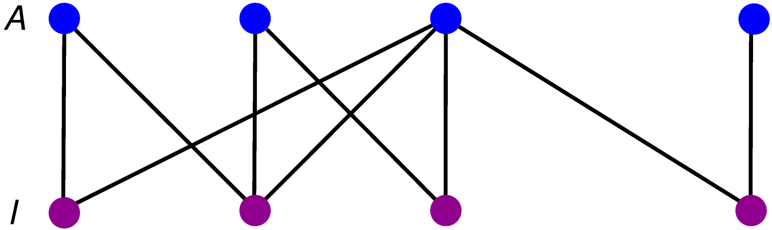

- How to find a matching with flow.

Algorithm (Offline)

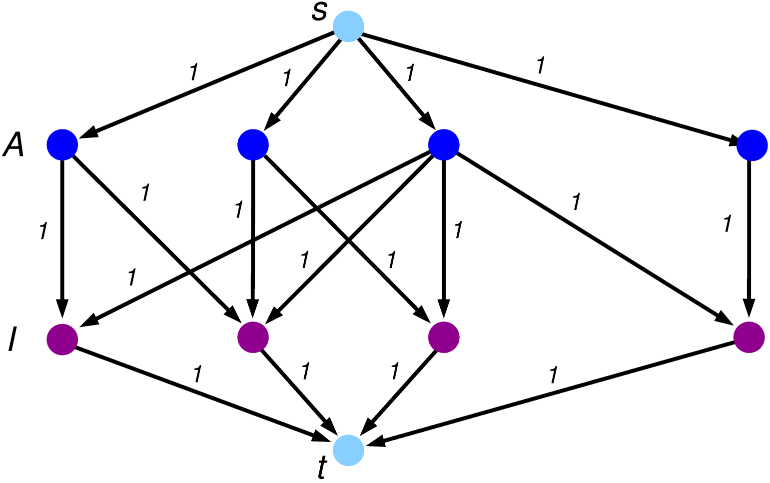

- How to find a matching with flow.

Algorithm (Offline)

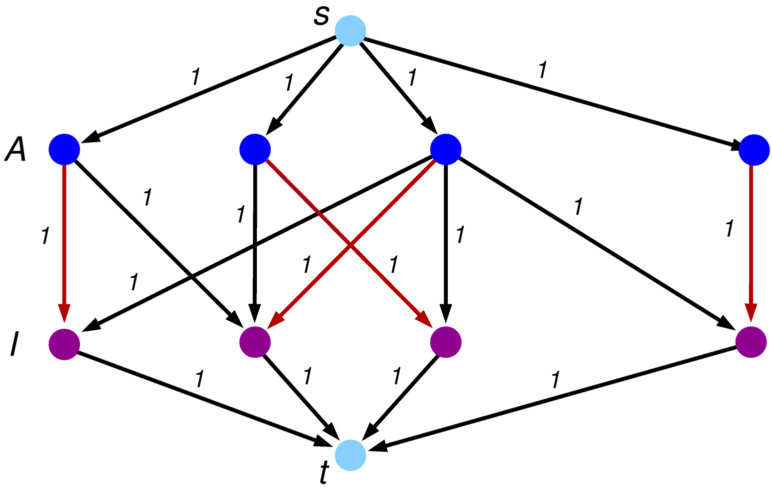

- How to find a matching with flow.

Algorithm (Offline)

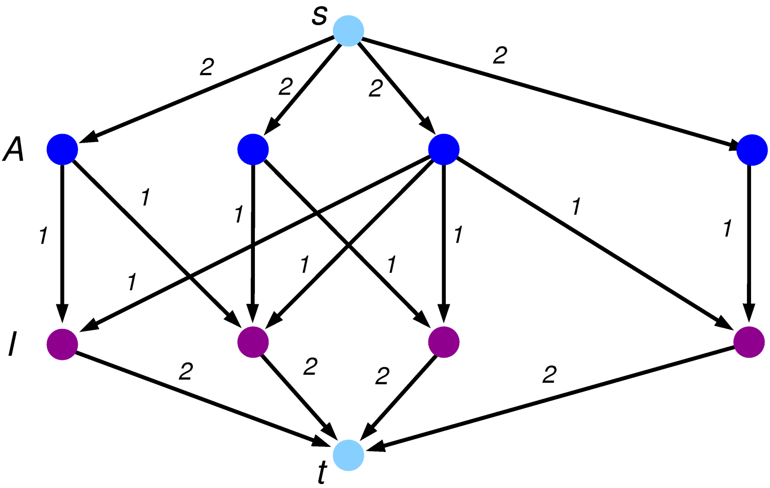

- Solve an "augmented flow" problem instead.

Algorithm (Offline)

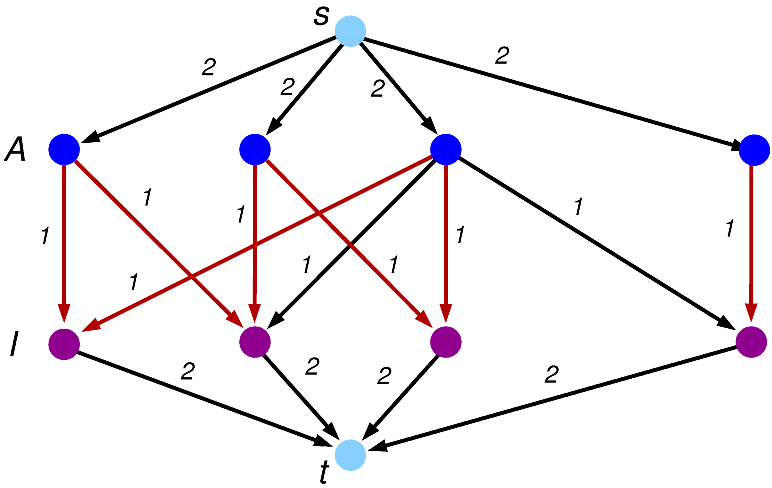

- Solve an "augmented flow" problem instead.

Algorithm (Offline)

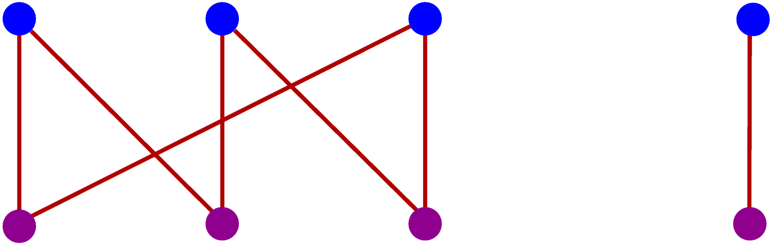

- Examine edges in flow.

Algorithm (Offline)

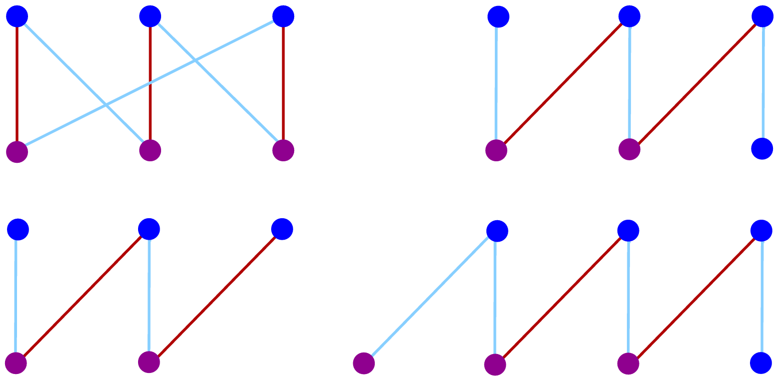

- Color the edges as shown

Algorithm (Online)

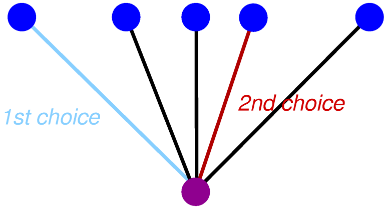

- When node $i \in I$ arrives:
  - Try the blue edge first, then the red edge.

Algorithm (Online)

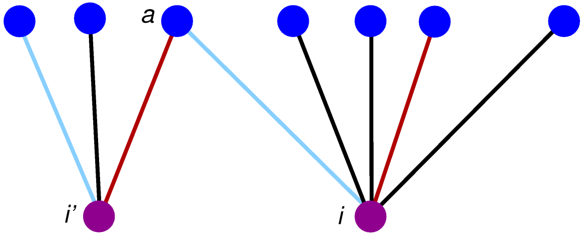

- Consider a node $a \in A$:
  - $\Pr[a \text{ is chosen }] \geq \Pr[i \text{ arrives once, or } i' \text{ arrives twice}]$

Performance of the Algorithm

- Classify $a \in A$ based on its neighbors in the flow.

$$|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$$

Performance of the Algorithm

- Classify $a \in A$ based on its neighbors in the flow.

$$|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$$

- Using Balls-in-bins concentration results (Azuma's inequality):

Performance of the Algorithm

- Classify $a \in A$ based on its neighbors in the flow.

$$|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$$

- Using Balls-in-bins concentration results (Azuma's inequality):
  - $a \in A_B$. We get at least $|A_B|(1 - 1/e)$.

Performance of the Algorithm

- Classify $a \in A$ based on its neighbors in the flow.

$$|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$$

- Using Balls-in-bins concentration results (Azuma's inequality):
  - $a \in A_B$. We get at least $|A_B|(1 - 1/e)$.
  - $a \in A_{BR}$. We get at least $|A_{BR}|(1 - 2/e^2)$.

Performance of the Algorithm

- Classify $a \in A$ based on its neighbors in the flow.

$$|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$$

- Using Balls-in-bins concentration results (Azuma's inequality):
  - $a \in A_B$. We get at least $|A_B|(1 - 1/e)$.
  - $a \in A_{BR}$. We get at least $|A_{BR}|(1 - 2/e^2)$.
  - $a \in A_{BB}$. We get at least $|A_{BB}|(1 - 1/e^2)$.

Performance of the Algorithm

- Classify $a \in A$ based on its neighbors in the flow.

$$|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$$

- Using Balls-in-bins concentration results (Azuma's inequality):
  - $a \in A_B$. We get at least $|A_B|(1 - 1/e)$.
  - $a \in A_{BR}$. We get at least $|A_{BR}|(1 - 2/e^2)$.
  - $a \in A_{BB}$. We get at least $|A_{BB}|(1 - 1/e^2)$.
  - $a \in A_R$. We get at least $|A_R|(1 - 2/e)$.

Performance of the Algorithm

- Classify $a \in A$ based on its neighbors in the flow.

$$|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$$

- Using Balls-in-bins concentration results (Azuma's inequality):
  - $a \in A_B$. We get at least $|A_B|(1 - 1/e)$.
  - $a \in A_{BR}$. We get at least $|A_{BR}|(1 - 2/e^2)$.
  - $a \in A_{BB}$. We get at least $|A_{BB}|(1 - 1/e^2)$.
  - $a \in A_R$. We get at least $|A_R|(1 - 2/e)$.
- Bound on ALG in terms of flow (using $|B| \geq |R|$):

$$ALG \geq \left(1 - \frac{1}{e^2}\right)|A_{BB}| + \left(1 - \frac{2}{e^2}\right)|A_{BR}| + \left(1 - \frac{3}{2e}\right)(|A_B| + |A_R|)$$

Bounding OPT

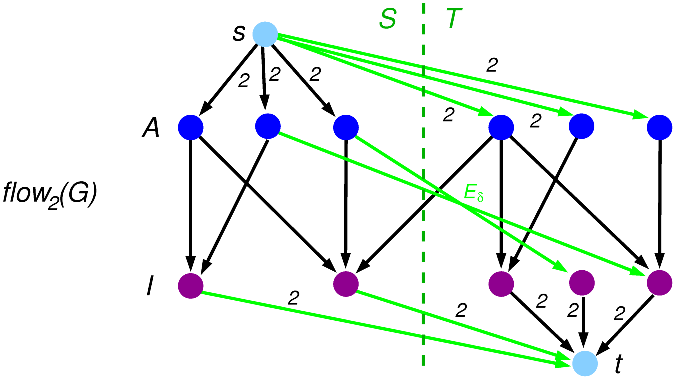

- Find min-cut in augmented flow graph (from G).
- E_δ is a matching.
- By max-flow min-cut,

$$|\text{flow}| = 2(|A_T| + |I_S|) + |E_\delta|.$$

Bounding OPT

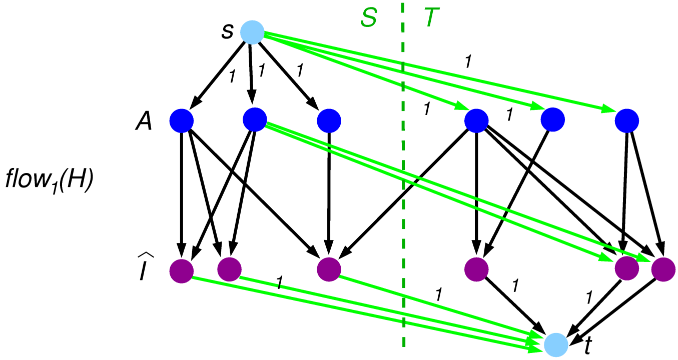

- OPT $\leq$ cut(H). (Remember $H = (A, \hat{I}, \hat{E})$.)
- Use min-cut in G as "guide" for cut in H.
- W.h.p., $|I_S| \approx |\hat{I}_S|$. E_δ?
- For any node $a \in S$ with an edge in the cut in $\hat{E}(H)$, move it to $T \Rightarrow$ # nonempty nodes in $E_\delta \Rightarrow (1 - \frac{1}{e})E_\delta$.

Putting things together

- Eventually (after moving a few nodes around) you get
  - $OPT \lesssim |I_S| + |A_T| + (1 - 1/e)|E_\delta|$.

Putting things together

- Eventually (after moving a few nodes around) you get
  - $OPT \lesssim |I_S| + |A_T| + (1 - 1/e)|E_\delta|$.
- A lemma relating the decomposition to the cut gives
  - $|E_\delta| \leq \frac{2}{3}|A_{BR}| + \frac{4}{3}|A_{BB}| + |A_B| + \frac{1}{3}|A_R|$,

Putting things together

- Eventually (after moving a few nodes around) you get
  - $OPT \lesssim |I_S| + |A_T| + (1 - 1/e)|E_\delta|$.
- A lemma relating the decomposition to the cut gives
  - $|E_\delta| \leq \frac{2}{3}|A_{BR}| + \frac{4}{3}|A_{BB}| + |A_B| + \frac{1}{3}|A_R|$,
- which, when combined with
  - $|\text{flow}| = 2(|A_T| + |I_S|) + |E_\delta|$
  - $|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$,
  - $\text{ALG} \geq (1 - \frac{1}{e^2})|A_{BB}| + (1 - \frac{2}{e^2})|A_{BR}| + (1 - \frac{3}{2e})(|A_B| + |A_R|)$,
- gives
  - $\frac{ALG}{OPT} \geq min\{\frac{1-1/e^2}{5/3-4/3e}, \frac{1-2/e^2}{4/3-2/3e}, \frac{1-3/2e}{1-1/e}\}$
  - $\frac{ALG}{OPT} \geq .67$

Putting things together

- Eventually (after moving a few nodes around) you get
  - $OPT \lesssim |I_S| + |A_T| + (1 - 1/e)|E_\delta|$.
- A lemma relating the decomposition to the cut gives
  - $|E_\delta| \leq \frac{2}{3}|A_{BR}| + \frac{4}{3}|A_{BB}| + |A_B| + \frac{1}{3}|A_R|$,
- which, when combined with
  - $|\text{flow}| = 2(|A_T| + |I_S|) + |E_\delta|$
  - $|\text{flow}| = 2|A_{BR}| + 2|A_{BB}| + |A_B| + |A_R|$,
  - $\text{ALG} \geq (1 - \frac{1}{e^2})|A_{BB}| + (1 - \frac{2}{e^2})|A_{BR}| + (1 - \frac{3}{2e})(|A_B| + |A_R|)$,
- gives
  - $\frac{ALG}{OPT} \geq min\{\frac{1-1/e^2}{5/3-4/3e}, \frac{1-2/e^2}{4/3-2/3e}, \frac{1-3/2e}{1-1/e}\}$
  - $\frac{ALG}{OPT} \geq .67$
- The analysis is tight.

Lower bound

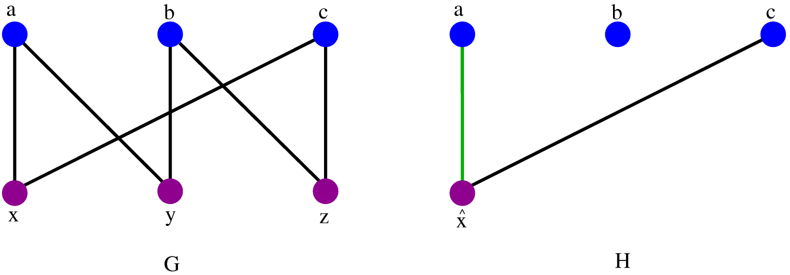

- Suppose wlog x arrives first, ALG assigns to a (if anywhere).

Lower bound

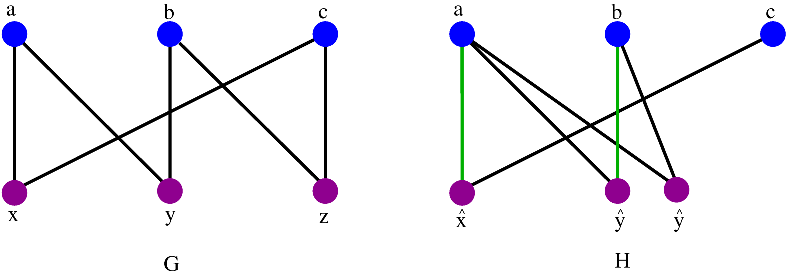

- Suppose wlog x arrives first, ALG assigns to a (if anywhere).
- If the next two arrivals are both y, then ALG ≤ 2, OPT $= 3$.

Lower bound

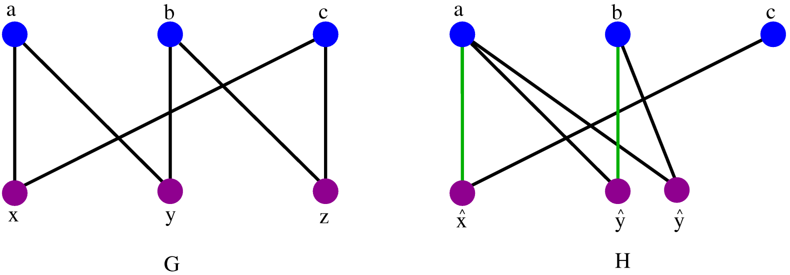

- Suppose wlog x arrives first, ALG assigns to a (if anywhere).
- If the next two arrivals are both y, then ALG ≤ 2, OPT $= 3$.
- Therefore E[ALG/OPT] $\leq (1/9)(2/3) + (8/9) = 26/27$

Lower bound

- But what if one demands an instance that grows with n ?

Lower bound

- But what if one demands an instance that grows with n ?
- Use collection of 6-cycles, same “bad event,” Azuma, Chernoff...

Lower bound

- But what if one demands an instance that grows with n ?
- Use collection of 6-cycles, same “bad event,” Azuma, Chernoff...
- Theorem: Even for $n \geq n_0$, no algorithm can do better than $\frac{6e^3-23}{6e^3-22} \simeq .9898$, even in expectation.

Outline

- Overview: Related Work and Results.
- Online Stochastic Matching: iid with known distribution
- Online Stochastic Packing: Random order
- Online Generalized Assignment (with free disposal)

Results[FHKMS10]: random order (iid w. unknown dist.)

- Generalize the primal-dual technique used by Devanour & Hayes to General Online Stochastic Packing LPs:
  - Online Stochastic Routing Problems
  - Online Stochastic Combinatorial Auctions

Results[FHKMS10]: random order (iid w. unknown dist.)

- Generalize the primal-dual technique used by Devanour & Hayes to General Online Stochastic Packing LPs:
  - Online Stochastic Routing Problems
  - Online Stochastic Combinatorial Auctions
- Thm: Under the following conditions, w.h.p there exists a $1-\epsilon$-approximation:
  - OPT is much larger than the maximum value of any item for any resource, i.e, $\frac{\mathrm{OPT}}{\max v_{ia}} \geq \frac{m \log n}{\epsilon}$.
  - Each item takes a small fraction of the total capacity of any resource, i.e., $\frac{C_a}{\max s_{ia}} \geq \frac{m \log n}{\epsilon^3}$.

$$\max \sum_{i,a} v_{ia} x_{ia}$$

$$\begin{aligned} \sum_a x_{ia} &\leq 1 && (\forall\ i) \\ \sum_i s_{ia} x_{ia} &\leq C_a && (\forall\ a) \end{aligned}$$

$$\min \sum_a C_a \beta_a + \sum_i z_i$$

$$\begin{aligned} s_{ia}\beta_a + z_i &\geq v_{ia} && (\forall i, a) \\ \beta_a, z_i &\geq 0 && (\forall i, a) \end{aligned}$$

Results[FHKMS10]: random order (iid w. unknown dist.)

- Thm: Under the following conditions, w.h.p there exists a $1-\epsilon$-approximation:
  - OPT is much larger than the maximum value of any item for any resource, i.e, $\frac{\text{OPT}}{\max v_{ia}} \geq \frac{m \log n}{\epsilon}$.
  - Each item takes a small fraction of the total capacity of any resource, i.e., $\frac{C_a}{\max s_{ia}} \geq \frac{m \log n}{\epsilon^3}$.

Results[FHKMS10]: random order (iid w. unknown dist.)

- Thm: Under the following conditions, w.h.p there exists a $1-\epsilon$-approximation:
  - OPT is much larger than the maximum value of any item for any resource, i.e, $\frac{\text{OPT}}{\max v_{ia}} \geq \frac{m \log n}{\epsilon}$.
  - Each item takes a small fraction of the total capacity of any resource, i.e., $\frac{C_a}{\max s_{ia}} \geq \frac{m \log n}{\epsilon^3}$.

- Algorithm:
  - Observe the first ϵ fraction sample of items.
  - Learn a dual variable for each ad β_a, by solving the dual program on the sample.
  - Assign each item i to ad a that maximizes $v_{ia} - \beta_a$.

Results[FHKMS10]: random order (iid w. unknown dist.)

- Thm: Under the following conditions, w.h.p there exists a $1 - \epsilon$-approximation:
  - $\textsc{OPT}$ is much larger than the maximum value of any item for any resource, i.e, $\frac{\textsc{OPT}}{\max v_{ia}} \geq \frac{m \log n}{\epsilon}$.
  - Each item takes a small fraction of the total capacity of any resource, i.e., $\frac{C_a}{\max s_{ia}} \geq \frac{m \log n}{\epsilon^3}$.
- Algorithm:
  - Observe the first ϵ fraction sample of items.
  - Learn a dual variable for each ad β_a, by solving the dual program on the sample.
  - Assign each item i to ad a that maximizes $v_{ia} - \beta_a$.
- More general: multiple resources in one option o. Maximize $v_{io} - \sum_{a \in o_i} \beta_a$

Outline

- Overview: Related Work and Results.
- Online Stochastic Matching: iid with known distribution
- Online Stochastic Packing: Random order
- Online Generalized Assignment (with free disposal)

Results[FKMMP09]: arbitrary order (with free disposal)

- Online Weighted Matching (with free disposal): [NWF78]: 0.5-approx.
- Online AdWord (Budgeted): [MSVV,BJN]: $1 - \frac{1}{e}$-approx

Results[FKMMP09]: arbitrary order (with free disposal)

- Online Weighted Matching (with free disposal): [NWF78]: 0.5-approx.
- Online AdWord (Budgeted): [MSVV,BJN]: $1 - \frac{1}{e}$-approx
- Online Generalized Assignment (with free disposal):
  - Items i may have different value (v_{ia}) and different size s_{ia} for different ads a.
  - Online Weighted Matching: $s_{ia} = 1$.
  - Online AdWord Assignment[MSVV]: $v_{ia} = s_{ia}$.

$$\max \sum_{i,a} v_{ia} x_{ia}$$

$$\begin{aligned} \sum_{a} x_{ia} &\leq 1 && (\forall\ i) \\ \sum_{i} s_{ia} x_{ia} &\leq C_a && (\forall\ a) \\ x_{ia} &\geq 0 && (\forall\ i, a) \end{aligned}$$

$$\min \sum_{a} C_a \beta_a + \sum_{i} z_i$$

$$\begin{aligned} s_{ia}\beta_a + z_i &\geq v_{ia} && (\forall i, a) \\ \beta_a, z_i &\geq 0 && (\forall i, a) \end{aligned}$$

Results[FKMMP09]: arbitrary order (with free disposal)

- Thm: There exists a $1 - \frac{1}{e} - \epsilon$-approximation algorithm if:
  - For online weighted matching, $C_a \geq O(\frac{1}{\epsilon})$.
  - For online GAP, each item takes a small fraction of the total capacity of any bin, i.e., $\frac{C_a}{\max s_{ij}} \geq \frac{1}{\epsilon}$.

Results[FKMMP09]: arbitrary order (with free disposal)

- Thm: There exists a $1 - \frac{1}{e} - \epsilon$-approximation algorithm if:
  - For online weighted matching, $C_a \geq O(\frac{1}{\epsilon})$.
  - For online GAP, each item takes a small fraction of the total capacity of any bin, i.e., $\frac{C_a}{\max s_{ij}} \geq \frac{1}{\epsilon}$.

- Algorithm:
  - Initialize $\beta = 0$.
  - Online: Assign item i to ad a that maximizes $v_{ia} - \beta_a$, and update β_a.
  - If the top C_a items assigned to a have values: $v(1) \geq v(2) \ldots \geq v(C_a)$:
    Greedy: $\beta_a = v(C_a)$, or Average: $\beta_a = \frac{\sum_{j=1}^{C_a} v(j)}{C_a}$, or

$$\beta_a = \frac{1}{C_a(e-1)} \sum_{j=1}^{C_a} v(j)(1 + \frac{1}{C_a})^{j-1}.$$

Proof Idea: arbitrary order (with free disposal)

$$\max \sum_{i,a} v_{ia} x_{ia}$$

$$\begin{aligned} \sum_a x_{ia} &\leq 1 & (\forall\ i) \\ \sum_i s_{ia} x_{ia} &\leq C_a & (\forall\ a) \\ x_{ia} &\geq 0 & (\forall\ i, a) \end{aligned}$$

$$\min \sum_a C_a \beta_a + \sum_i z_i$$

$$\begin{aligned} s_{ia}\beta_a + z_i &\geq v_{ia} & (\forall i, a) \\ \beta_a, z_i &\geq 0 & (\forall i, a) \end{aligned}$$

- Proof:
  1. Start from feasible primal and dual ($x_{ia} = 0$, $\beta_a = 0$, and $z_i = 0$, i.e., Primal=Dual=0).
  2. After each assignment, update x, β, z variables and keep primal and dual solutions.
  3. Show $\Delta(\text{Dual}) \leq (1 - \frac{1}{e})\Delta(\text{Primal})$.

Proof Idea: arbitrary order (with free disposal)

$$\max \sum_{i,a} v_{ia} x_{ia}$$

$$\begin{aligned} \sum_a x_{ia} &\leq 1 && (\forall\ i) \\ \sum_i s_{ia} x_{ia} &\leq C_a && (\forall\ a) \\ x_{ia} &\geq 0 && (\forall\ i, a) \end{aligned}$$

$$\min \sum_a C_a \beta_a + \sum_i z_i$$

$$\begin{aligned} s_{ia}\beta_a + z_i &\geq v_{ia} && (\forall i, a) \\ \beta_a, z_i &\geq 0 && (\forall i, a) \end{aligned}$$

- Proof:
  1. Start from feasible primal and dual ($x_{ia} = 0$, $\beta_a = 0$, and $z_i = 0$, i.e., Primal=Dual=0).
  2. After each assignment, update x, β, z variables and keep primal and dual solutions.
  3. Show $\Delta(\text{Dual}) \leq (1 - \frac{1}{e})\Delta(\text{Primal})$.

- Thm: There exists a $1 - \frac{1}{e} - \epsilon$-approximation algorithm if:
  For online weighted matching: $C_a \geq O(\frac{1}{\epsilon})$

Conclusions

- 0.67-approximation for online stochastic matching (iid, known distribution): power of two choices.
- $1-\epsilon$-approximation for online stochastic packing (random order, under assumptions): learning dual variables.
- $1-\frac{1}{e}$-approximation for online GAP with free disposal (with small elements, adversarial): primal-dual approach.

Conclusions

- 0.67-approximation for online stochastic matching (iid, known distribution): power of two choices.
- $1 - \epsilon$-approximation for online stochastic packing (random order, under assumptions): learning dual variables.
- $1 - \frac{1}{e}$-approximation for online GAP with free disposal (with small elements, adversarial): primal-dual approach.
- Open Problems:
  - Gap: $0.98 > 0.67$.
  - Apply power of two choices to other problems?
  - Power of "many choices"?
  - An algorithm that achieves good performance both in stochastic setting and worst case?

Thank you.