Online Ad Serving: Theory and Practice

Aranyak Mehta
Vahab Mirrokni

June 7, 2011

Outline of this talk

- **Ad delivery for contract-based settings**
  - Ad Serving
  - Planning
- Ad serving in repeated auction settings
  - General architecture.
  - Allocation for budget constrained advertisers.
- Other interactions
  - Learning + allocation
  - Learning + auction
  - Auction + contracts

Contract-based Ad Delivery: Outline

- Basic Information
- Ad Serving.
  - Targeting.
  - Online Allocation
- Ad Planning: Reservation

Contract-based Online Advertising

- Pageviews (impressions) instead of queries.
- Display/Banner Ads, Video Ads, Mobile Ads.

Contract-based Online Advertising

- Pageviews (impressions) instead of queries.
- Display/Banner Ads, Video Ads, Mobile Ads.
- Cost-Per-Impression (CPM).
- Not Auction-based: offline negotiations + Online allocations.

Contract-based Online Advertising

- Pageviews (impressions) instead of queries.
- Display/Banner Ads, Video Ads, Mobile Ads.
- Cost-Per-Impression (CPM).
- Not Auction-based: offline negotiations + Online allocations.

Display/Banner Ads:

- Q1, 2010: One Trillion Display Ads in US. $2.7 billion.
- Top Advertiser: AT&T, Verizon, Scottrade.

Contract-based Online Advertising

- Pageviews (impressions) instead of queries.
- Display/Banner Ads, Video Ads, Mobile Ads.
- Cost-Per-Impression (CPM).
- Not Auction-based: offline negotiations + Online allocations.

Display/Banner Ads:

- Q1, 2010: One Trillion Display Ads in US. $2.7 billion.
- Top Advertiser: AT&T, Verizon, Scottrade.
- Ad Serving Systems e.g. Facebook, Google Doubleclick, AdMob.

Internet Advertising Revenues - 2010

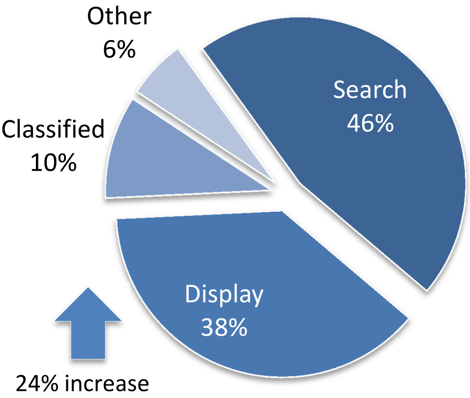

Total $26.0 billion

Display Ad Delivery: Overview

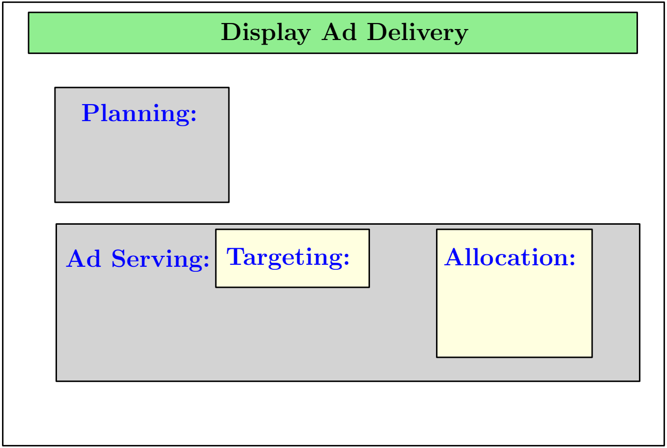

1. Planning: Contracts/Commitments with Advertisers.
2. Ad Serving:
   - Targeting: Predicting value of impressions.
   - Ad Allocation: Assigning Impressions to Ads Online.

Display Ad Delivery: Overview

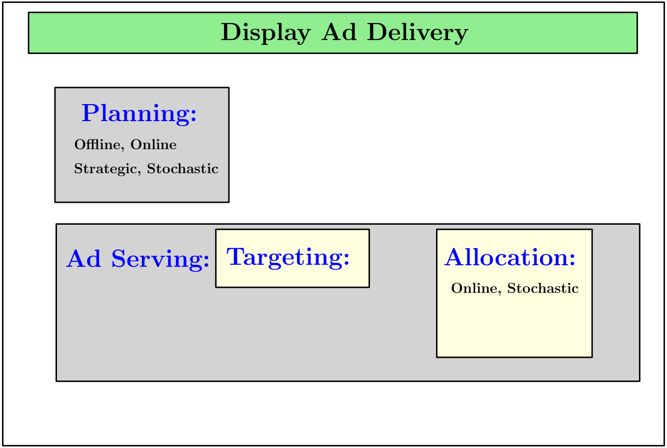

1. Planning: Contracts/Commitments with Advertisers.
2. Ad Serving:
   - Targeting: Predicting value of impressions.
   - Ad Allocation: Assigning Impressions to Ads Online.

Display Ad Delivery: Overview

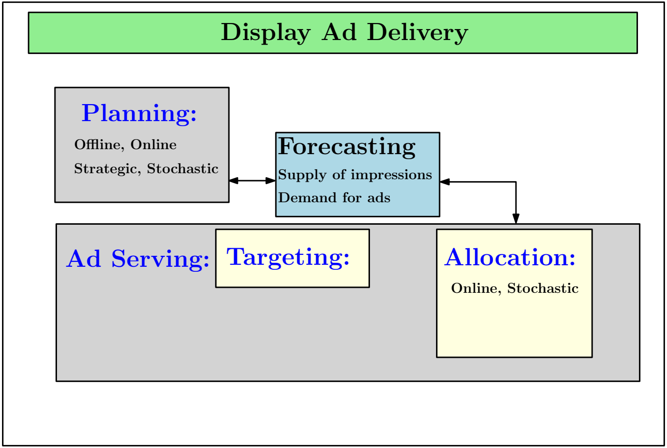

1. Planning: Contracts/Commitments with Advertisers.
2. Ad Serving:
   - Targeting: Predicting value of impressions.
   - Ad Allocation: Assigning Impressions to Ads Online.

Display Ad Delivery: Overview

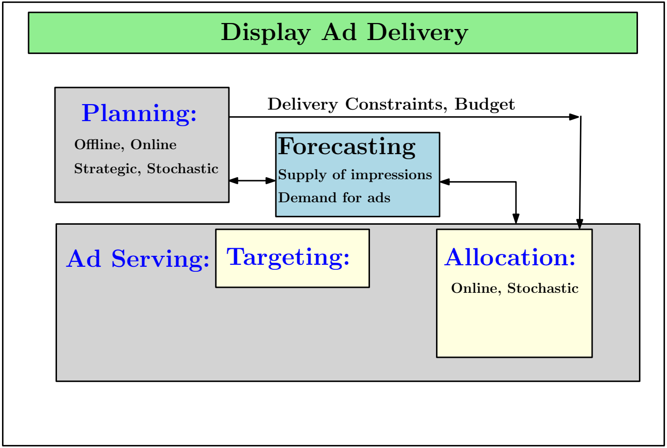

1. Planning: Contracts/Commitments with Advertisers.
2. Ad Serving:
   - Targeting: Predicting value of impressions.
   - Ad Allocation: Assigning Impressions to Ads Online.

Display Ad Delivery: Overview

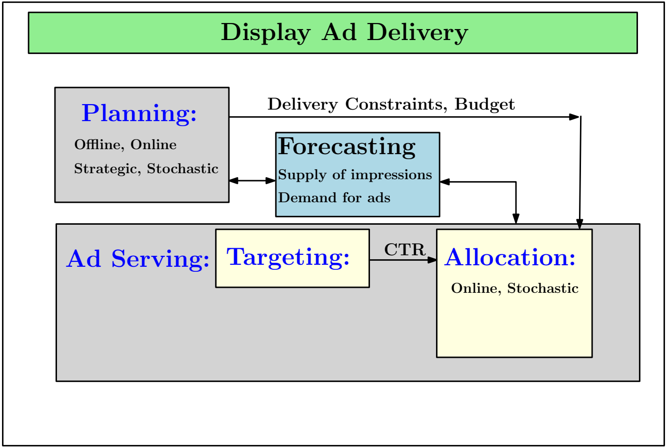

1. Planning: Contracts/Commitments with Advertisers.
2. Ad Serving:
   - Targeting: Predicting value of impressions.
   - Ad Allocation: Assigning Impressions to Ads Online.

Display Ad Delivery: Overview

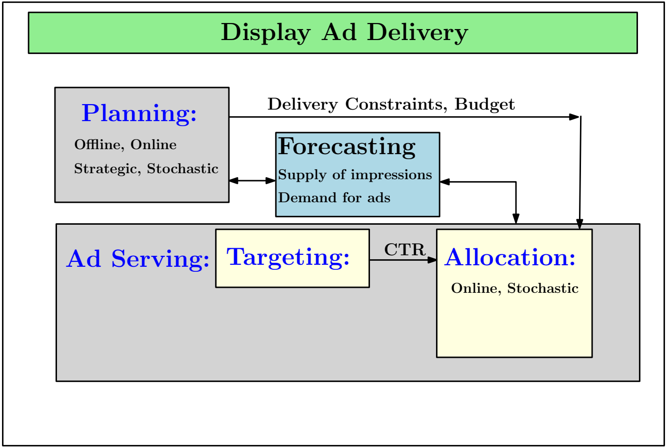

- **Objective Functions:**
  - Efficiency: Users and Advertisers. Revenue of the Publisher.
  - Smoothness, Fairness, Delivery Penalty.

Contract-based Ad Delivery: Outline

- Basic Information
- Ad Serving.
  - Targeting.
  - Online Ad Allocation
- Ad Planning: Reservation

Targeting

Estimating Value of an impression.

Targeting

Estimating Value of an impression.

- Behavioral Targeting
  - Interest-based Advertising.
  - Yan, Liu, Wang, Zhang, Jiang, Chen, 2009, How much can Behavioral Targeting Help Online Advertising?

Targeting

Estimating Value of an impression.

- Behavioral Targeting
  - Interest-based Advertising.
  - Yan, Liu, Wang, Zhang, Jiang, Chen, 2009, How much can Behavioral Targeting Help Online Advertising?
- Contextual Targeting
  - Information Retrieval (IR).
  - Broder, Fontoura, Josifovski, Riedel, A semantic approach to contextual advertising

Targeting

Estimating Value of an impression.

- Behavioral Targeting
  - Interest-based Advertising.
  - Yan, Liu, Wang, Zhang, Jiang, Chen, 2009, How much can Behavioral Targeting Help Online Advertising?
- Contextual Targeting
  - Information Retrieval (IR).
  - Broder, Fontoura, Josifovski, Riedel, A semantic approach to contextual advertising
- Creative Optimization
  - Experimentation

Predicting value of Impressions for Display Ads

- Estimating Click-Through-Rate (CTR).
  - Budgeted Multi-armed Bandit
- Probability of Conversion.

Predicting value of Impressions for Display Ads

- Estimating Click-Through-Rate (CTR).
  - Budgeted Multi-armed Bandit
- Probability of Conversion.
- Long-term vs. Short-term value of display ads?
  - Archak, Mirrokni, Muthukrishnan, 2010 Graph-based Models.
    - Computing Adfactors based on AdGraphs
    - Markov Models for Advertiser-specific User Behavior

Contract-based Ad Delivery: Outline

- Basic Information
- Ad Planning: Reservation
- Ad Serving.
  - Targeting.
  - Online Ad Allocation

Outline: Online Allocation

- Online Stochastic Assignment Problems
  - Online (Stochastic) Matching
  - Online Stochastic Packing
  - Online Generalized Assignment (with free disposal)
  - Experimental Results
- Online Learning and Allocation

Online Ad Allocation

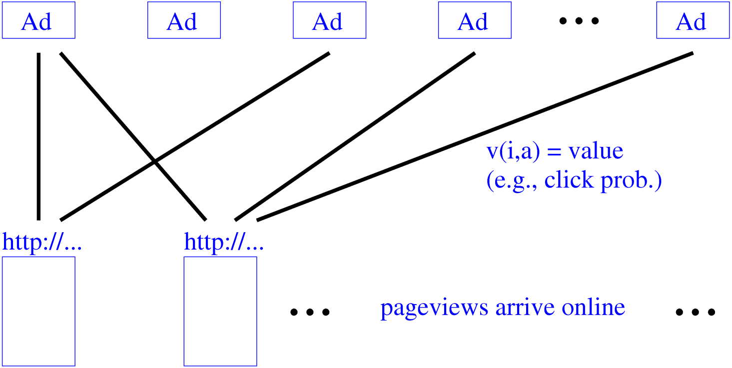

- When page arrives, assign an eligible ad.
  - value of assigning page i to ad a: v_{ia}

Online Ad Allocation

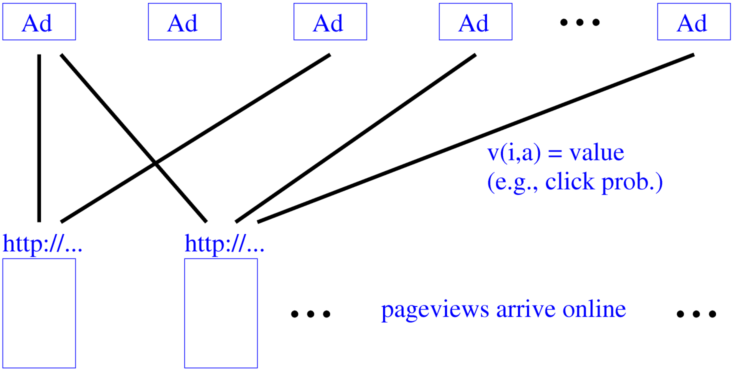

- When page arrives, assign an eligible ad.
  - value of assigning page i to ad a: v_{ia}
- Display Ads (DA) problem:
  - Maximize value of ads served: $\max \sum_{i,a} v_{ia} x_{ia}$
  - Capacity of ad a: $\sum_{i \in A(a)} x_{ia} \leq C_a$

Online Ad Allocation

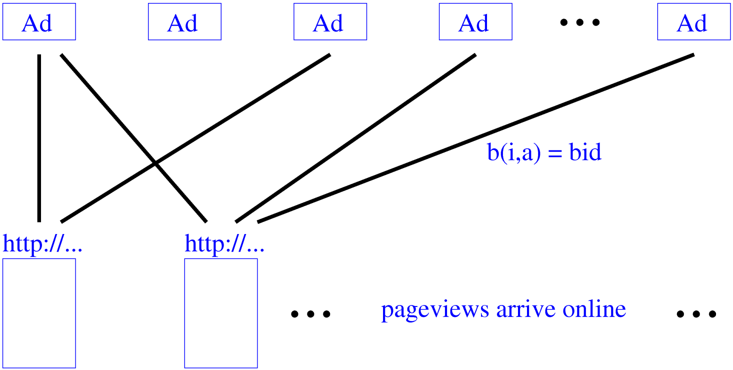

- When page arrives, assign an eligible ad.
  - revenue from assigning page i to ad a: b_{ia}
- "AdWords" (AW) problem:
  - Maximize revenue of ads served: $\max \sum_{i,a} b_{ia} x_{ia}$
  - Budget of ad a: $\sum_{i \in A(a)} b_{ia} x_{ia} \leq B_a$

General Form of LP

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

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

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$

General Form of LP

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

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

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst-Case	Greedy: $\frac{1}{2}$, [KVV]: $1 - \frac{1}{e}$-aprx		[MSVV,BJN]: $1 - \frac{1}{e}$-aprx

Ad Allocation: Problems and Models

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst Case	Greedy: $\frac{1}{2}$, [KVV]: $1 - \frac{1}{e}$-aprx	?	[MSVV,BJN]: $1 - \frac{1}{e}$-aprx

Ad Allocation: Problems and Models

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst Case	Greedy: $\frac{1}{2}$, [KVV]: $1 - \frac{1}{e}$-aprx	?	[MSVV,BJN]: $1 - \frac{1}{e}$-aprx
Stochastic (i.i.d.)		?	[DH09]: $1 - \epsilon$-aprx, if $\text{OPT} \gg \max v_{ia}$

Stochastic i.i.d model:

- i.i.d model with known distribution
- random order model (i.i.d model with unknown distribution)

Ad Allocation: Problems and Models

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst Case	Greedy: $\frac{1}{2}$, [KVV]: $1 - \frac{1}{e}$-aprx	?	[MSVV,BJN]: $1 - \frac{1}{e}$-aprx
Stochastic (i.i.d.)	[FMMM09,MOS11]: 0.702-aprx i.i.d with known distribution	?	[DH09]: $1 - \epsilon$-aprx, if $\text{OPT} \gg \max v_{ia}$

Stochastic i.i.d model:

- i.i.d model with known distribution
- random order model (i.i.d model with unknown distribution)

Ad Allocation: Problems and Models

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst Case	Greedy: $\frac{1}{2}$, [KVV]: $1-\frac{1}{e}$-aprx	?	[MSVV,BJN]: $1-\frac{1}{e}$-aprx
Stochastic (i.i.d.)	[FMMM09]: 0.67-aprx i.i.d with known distribution	?	[DH09]: $1-\epsilon$-aprx, if $\mathrm{OPT} \gg \max v_{ia}$

Ad Allocation: Problems and Models

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst Case	Greedy: $\frac{1}{2}$, [KVV]: $1 - \frac{1}{e}$-aprx	?	[MSVV,BJN]: $1 - \frac{1}{e}$-aprx
Stochastic (i.i.d.)	[FMMM09,MOS11]: 0.702-aprx i.i.d with known distribution	[FHKMS10,AWY]: $1 - \epsilon$-aprx, if $\text{OPT} \gg \max v_{ia}$ and $C_a \gg \max s_{ia}$	[DH09]: $1 - \epsilon$-aprx, if $\text{OPT} \gg \max v_{ia}$

random order = i.i.d. model with unknown distribution

Stochastic DA: Dual Algorithm

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

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

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

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

Stochastic DA: Dual Algorithm

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

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

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

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

Algorithm:

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

Stochastic DA: Dual Algorithm

Feldman, Henzinger, Korula, M., Stein 2010

Thm[FHKMS10,AWY]: W.h.p, this algorithm is a $(1-O(\epsilon))$-aprx, as long as each item has low value ($v_{ia} \leq \frac{\epsilon \mathrm{OPT}}{m \log n}$), and large capacity ($C_a \geq \frac{m \log n}{\epsilon^3}$)

Stochastic DA: Dual Algorithm

Feldman, Henzinger, Korula, M., Stein 2010

Thm[FHKMS10,AWY]: W.h.p, this algorithm is a $(1-O(\epsilon))$-aprx, as long as each item has low value ($v_{ia} \leq \frac{\epsilon \mathrm{OPT}}{m \log n}$), and large capacity ($C_a \geq \frac{m \log n}{\epsilon^3}$)

Fact: If optimum β^*_a are known, this alg. finds OPT

- Proof: Comp. slackness. Given β^*_a, compute x^* as follows: $x^*_{ia} = 1$ if $a = \text{argmax}(v_{ia} - \beta^*_a)$.

Stochastic DA: Dual Algorithm

Feldman, Henzinger, Korula, M., Stein 2010

Thm[FHKMS10,AWY]: W.h.p, this algorithm is a $(1-O(\epsilon))$-aprx, as long as each item has low value ($v_{ia} \leq \frac{\epsilon \text{OPT}}{m \log n}$), and large capacity ($C_a \geq \frac{m \log n}{\epsilon^3}$)

Fact: If optimum β_a^* are known, this alg. finds OPT

- Proof: Comp. slackness. Given β_a^*, compute x^* as follows: $x_{ia}^* = 1$ if $a = \text{argmax}(v_{ia} - \beta_a^*)$.

Lemma: In the random order model, W.h.p., the sample β_a' are close to β_a^*.

- Extending DH09.

General Stochastic Packing LPs

- m fixed *resources* with capacity C_a
- *Items* i arrive online with *options* O_i, *values* v_{io}, *rsrc. use* s_{ioa}.
  - Choose $o \in O_i$, using up capacity s_{ioa} in *all* a.

Thm[FHKMS10,AWY]: W.h.p, the PD algorithm is a $(1-O(\epsilon))$-aprx, as long as items have low value ($v_{io} \leq \frac{\epsilon \mathrm{OPT}}{\log n}$) and small size ($s_{ioa} \leq \frac{\epsilon^3 C_a}{\log n}$).

General Stochastic Packing LPs

- m fixed *resources* with capacity C_a
- *Items* i arrive online with *options* O_i, *values* v_{io}, *rsrc. use* s_{ioa}.
  - Choose $o \in O_i$, using up capacity s_{ioa} in *all* a.

Thm[FHKMS10,AWY]: W.h.p, the PD algorithm is a $(1-O(\epsilon))$-aprx, as long as items have low value ($v_{io} \leq \frac{\epsilon \mathrm{OPT}}{\log n}$) and small size ($s_{ioa} \leq \frac{\epsilon^3 C_a}{\log n}$).

Other Results and Extensions (random order model):

- Agrawal, Wang, Ye: Updating dual variables by periodic solution of the dual program: $C_a \geq \frac{m \log n}{\epsilon^2}$ or $s_{ioa} \leq \frac{\epsilon^2 C_a}{M}$

General Stochastic Packing LPs

- m fixed *resources* with capacity C_a
- *Items* i arrive online with *options* O_i, *values* v_{io}, *rsrc. use* s_{ioa}.
  - Choose $o \in O_i$, using up capacity s_{ioa} in *all* a.

Thm[FHKMS10,AWY]: W.h.p, the PD algorithm is a $(1-O(\epsilon))$-aprx, as long as items have low value ($v_{io} \leq \frac{\epsilon \mathrm{OPT}}{\log n}$) and small size ($s_{ioa} \leq \frac{\epsilon^3 C_a}{\log n}$).

Other Results and Extensions (random order model):

- Agrawal, Wang, Ye: Updating dual variables by periodic solution of the dual program: $C_a \geq \frac{m \log n}{\epsilon^2}$ or $s_{ioa} \leq \frac{\epsilon^2 C_a}{M}$
- Vee, Vassilvitskii , Shanmugasundaram 2010: extension to convex objective functions: Using KKT conditions.

Ad Allocation: Problems and Models

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst Case	Greedy: $\frac{1}{2}$, [KVV]: $1 - \frac{1}{e}$-aprx	?	[MSVV,BJN]: $1 - \frac{1}{e}$-aprx
Stochastic (i.i.d.)	[FMMM09,MOS11]: 0.702-aprx i.i.d with known distribution	[FHKMS10,AWY]: $1 - \epsilon$-aprx, if $\mathrm{OPT} \gg \max v_{ia}$ and $C_a \gg \max s_{ia}$	[DH09]: $1 - \epsilon$-aprx, if $\mathrm{OPT} \gg \max v_{ia}$

Ad Allocation: Problems and Models

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst Case	Greedy: $\frac{1}{2}$, [KVV]: $1 - \frac{1}{e}$-aprx	Free Disposal [FKMMP09]: $1 - \frac{1}{e}$-aprx: if $C_a \gg \max s_{ia}$	[MSVV,BJN]: $1 - \frac{1}{e}$-aprx
Stochastic (i.i.d.)	[FMMM09,MOS11]: 0.702-aprx i.i.d with known distribution	[FHKMS10,AWY]: $1 - \epsilon$-aprx, if $\text{OPT} \gg \max v_{ia}$ and $C_a \gg \max s_{ia}$	[DH09]: $1 - \epsilon$-aprx, if $\text{OPT} \gg \max v_{ia}$

DA: Free Disposal Model

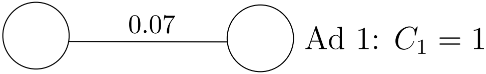

DA: Free Disposal Model

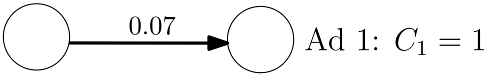

DA: Free Disposal Model

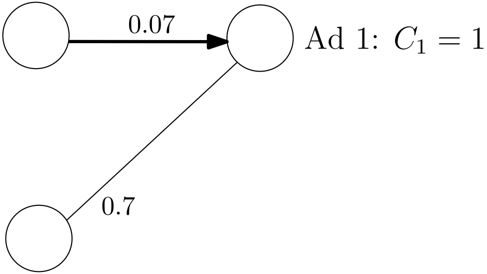

DA: Free Disposal Model

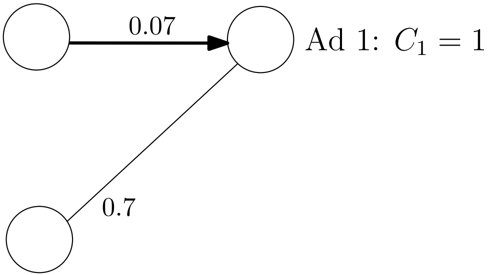

- Advertisers may not complain about extra impressions, but no bonus points for extra impressions, either.

DA: Free Disposal Model

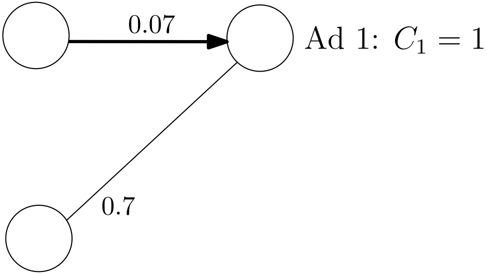

- Advertisers may not complain about extra impressions, but no bonus points for extra impressions, either.
- Value of advertiser = sum of values of top C_a items she gets.

Greedy Algorithm

Assign impression to an advertiser
maximizing Marginal Gain = (imp. value - min. impression value).

Greedy Algorithm

Assign impression to an advertiser
maximizing Marginal Gain = (imp. value - min. impression value).

- Competitive Ratio: 1/2. [NWF78]
  - Follows from submodularity of the value function.

Greedy Algorithm

Assign impression to an advertiser
maximizing Marginal Gain = (imp. value - min. impression value).

- Competitive Ratio: 1/2. [NWF78]
  - Follows from submodularity of the value function.

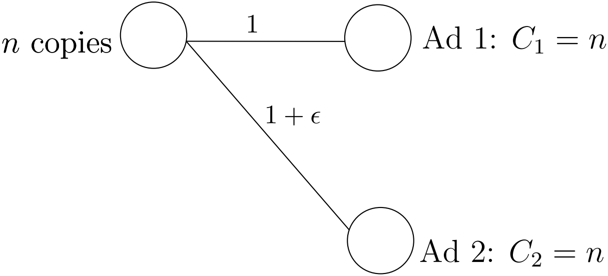

Greedy Algorithm

Assign impression to an advertiser
maximizing Marginal Gain = (imp. value - min. impression value).

- Competitive Ratio: 1/2. [NWF78]
  - Follows from submodularity of the value function.

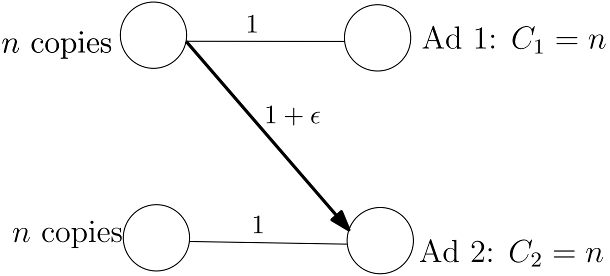

Greedy Algorithm

Assign impression to an advertiser
maximizing Marginal Gain = (imp. value - min. impression value).

- Competitive Ratio: 1/2. [NWF78]
  - Follows from submodularity of the value function.

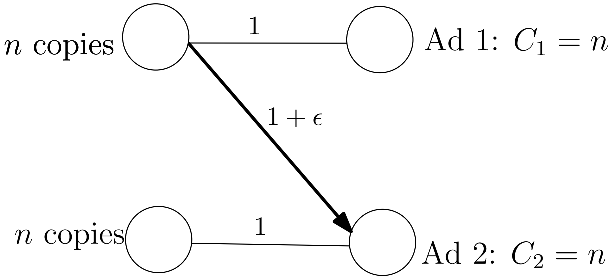

Evenly Split?

A better algorithm?

Assign impression to an advertiser a
maximizing (imp. value - β_a),
where $\beta_a =$ average value of top C_a impressions assigned to a.

A better algorithm?

Assign impression to an advertiser a
maximizing (imp. value - β_a),
where β_a = average value of top C_a impressions assigned to a.

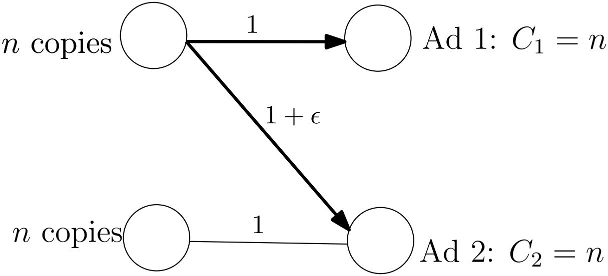

A better algorithm?

Assign impression to an advertiser a
maximizing (imp. value - β_a),
where $\beta_a =$ average value of top C_a impressions assigned to a.

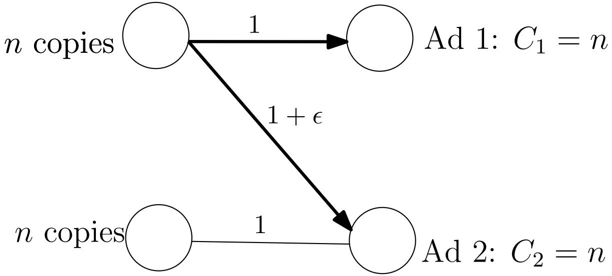

- Competitive Ratio: $\frac{1}{2}$ if $C_a >> 1$. [FKMMP09]
  - Primal-Dual Approach.

An Optimal Algorithm

Assign impression to an advertiser a:
maximizing (imp. value - β_a),

- Greedy: $\beta_a =$ min. impression assigned to a.
- Better (pd-avg): $\beta_a =$ average value of top C_a impressions assigned to a.

An Optimal Algorithm

Assign impression to an advertiser a:
maximizing (imp. value - β_a),

- Greedy: $\beta_a =$ min. impression assigned to a.
- Better (pd-avg): $\beta_a =$ average value of top C_a impressions assigned to a.
- Optimal (pd-exp): order value of edges assigned to a: $v(1) \geq v(2) \ldots \geq v(C_a)$:

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

An Optimal Algorithm

Assign impression to an advertiser a:
maximizing (imp. value - β_a),

- Greedy: $\beta_a =$ min. impression assigned to a.
- Better (pd-avg): $\beta_a =$ average value of top C_a impressions assigned to a.
- Optimal (pd-exp): order value of edges assigned to a: $v(1) \geq v(2) \ldots \geq v(C_a)$:

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

- Thm: pd-exp achieves optimal competitive Ratio: $1 - \frac{1}{e} - \epsilon$ if $C_a > O(\frac{1}{\epsilon})$. [Feldman, Korula, M., Muthukrishnan, Pal 2009]

Online Generalized Assignment (with free disposal)

- Multiple Knapsack: Item i may have different value (v_{ia}) and different size s_{ia} for different ads a.
- DA: $s_{ia} = 1$, AW: $v_{ia} = s_{ia}$.

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

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

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

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

Online Generalized Assignment (with free disposal)

- Multiple Knapsack: Item i may have different value (v_{ia}) and different size s_{ia} for different ads a.
- DA: $s_{ia} = 1$, AW: $v_{ia} = s_{ia}$.

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

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

- Offline Optimization: $1 - \frac{1}{e} - \delta$-aprx[FGMS07,FV08].
- Thm[FKMMP09]: There exists a $1 - \frac{1}{e} - \epsilon$-approximation algorithm if $\frac{C_a}{\max s_{ia}} \geq \frac{1}{\epsilon}$.

Proof Idea: Primal-Dual Analysis [BJN]

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

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

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

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

Proof Idea: Primal-Dual Analysis [BJN]

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

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

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

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

- 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 $\Delta(\text{Dual}) \leq (1 - \frac{1}{e})\Delta(\text{Primal})$.

Ad Allocation: Problems and Models

	Online Matching: $v_{ia} = s_{ia} = 1$	Disp. Ads (DA): $s_{ia} = 1$	AdWords (AW): $s_{ia} = v_{ia}$
Worst Case	Greedy: $\frac{1}{2}$, [KVV]: $1 - \frac{1}{e}$-aprx	Free Disposal [FKMMP09]: $1 - \frac{1}{e}$-aprx: if $C_a \gg \max s_{ia}$	[MSVV,BJN]: $1 - \frac{1}{e}$-aprx
Stochastic (random arrival order)	[FMMM09,MOS11]: 0.702-aprx	[FHKMS10,AWY]: $1 - \epsilon$-aprx, if $\text{OPT} \gg \max v_{ia}$ and $C_a \gg \max s_{ia}$	[DH09]: $1 - \epsilon$-aprx, if $\text{OPT} \gg \max v_{ia}$

Outline: Online Allocation

- Online Stochastic Assignment Problems
  - Online Stochastic Packing
  - Online Generalized Assignment (with free disposal)
  - Experimental Evaluation
  - Online Stochastic Weighted Matching

Dual-based Algorithms in Practice

- Algorithm:
  - Assign each item i to ad a that maximizes $v_{ia} - \beta_a$.

Dual-based Algorithms in Practice

- Algorithm:
  - Assign each item i to ad a that maximizes $v_{ia} - \beta_a$.
- More practical compared to Primal Algorithms:
  - Just keep one number β_a per advertiser.
  - Suitable for Distributed Ad Serving Schemes.

Dual-based Algorithms in Practice

- Algorithm:
  - Assign each item i to ad a that maximizes $v_{ia} - \beta_a$.
- More practical compared to Primal Algorithms:
  - Just keep one number β_a per advertiser.
  - Suitable for Distributed Ad Serving Schemes.
- Training-based Algorithms
  - Compute β_a based on historical/sample data.

Dual-based Algorithms in Practice

- Algorithm:
  - Assign each item i to ad a that maximizes $v_{ia} - \beta_a$.
- More practical compared to Primal Algorithms:
  - Just keep one number β_a per advertiser.
  - Suitable for Distributed Ad Serving Schemes.
- Training-based Algorithms
  - Compute β_a based on historical/sample data.
- Hybrid approach (see also [MNS07]):
  - Start with trained β_a (past history), blend in online algorithm.

Experiments: setup

- Real ad impression data from several large publishers
- 200k - 1.5M impressions in simulation period
- 100 - 2600 advertisers
- Edge weights = predicted click probability
- Efficiency: free disposal model
- Algorithms:
  - greedy: maximum marginal value
  - pd-avg, pd-exp: pure online primal-dual from [FKMMP09].
  - dualbase: training-based primal-dual [FHKMS10]
  - hybrid: convex combo of training based, pure online.
  - lp-weight: optimum efficiency

Experimental Evaluation: Summary

Algorithm	Avg Efficiency%
opt	100
greedy	69
pd-avg	77
pd-exp	82
dualbase	87
hybrid	89

- pd-exp & pd-avg outperform greedy by 9% and 14% (with more improvements in *tight* competition.)
- dualbase outperforms pure online algorithms by 6% to 12%.
- Hybrid has a mild improvement of 2% (up to 10%).
- pd-avg performs much better than the theoretical analysis.

Other Metrics: Fairness

- Qualititative definition: advertisers are “treated equally.”

Other Metrics: Fairness

- Qualititative definition: advertisers are "treated equally."
- One suggestion[FHKMS10]: Compute "fair" solution x^*, measure ℓ_1 distance to x^*.

Other Metrics: Fairness

- Qualititative definition: advertisers are "treated equally."
- One suggestion[FHKMS10]: Compute "fair" solution x^*, measure ℓ_1 distance to x^*.
- Fair solution:
  - Each a chooses best C_a impressions (highest v_{ia})
  - Repeat:
    - Impressions shared among those who chose them.
    - If some a not receiving C_a imps, a chooses an additional imp.

Experiments: highlights

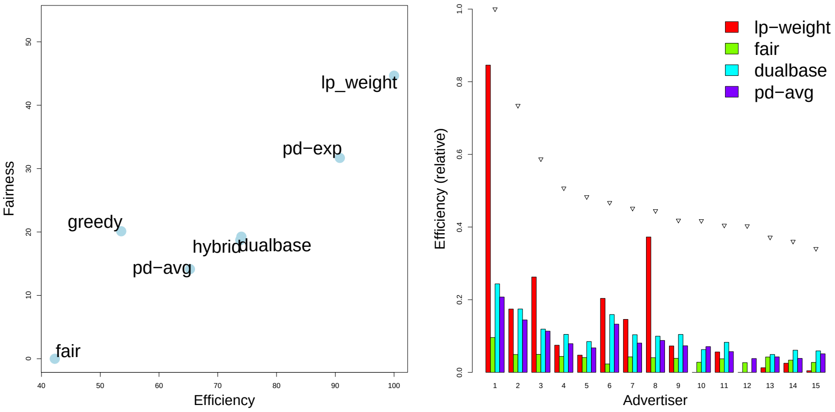

Experiments: highlights

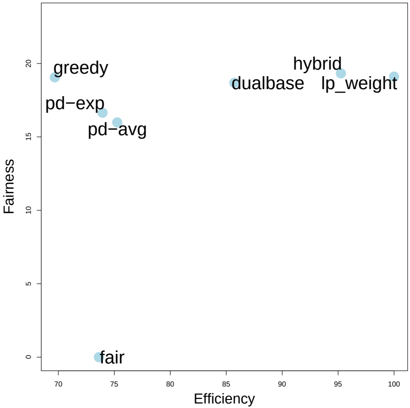

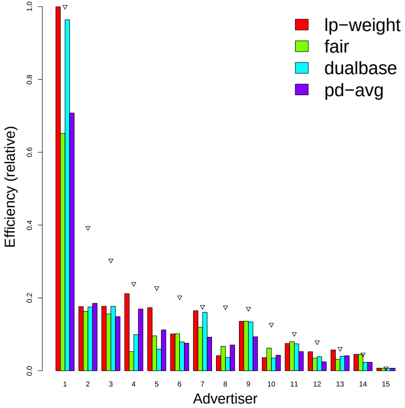

In Production

- Smooth Delivery of Display Ads
  - Delivery of impressions throughout time should follow the traffic smoothly.
  - Model this with multiple nested capacity constraints: $1 - 1/e$-competitive algorithm for this extension.
  - Bhalgat, Feldman, M., 2011

In Production

- Smooth Delivery of Display Ads
  - Delivery of impressions throughout time should follow the traffic smoothly.
  - Model this with multiple nested capacity constraints: $1 - 1/e$-competitive algorithm for this extension.
  - Bhalgat, Feldman, M., 2011
- Combined Allocation with Ad Exchange
  - Yield Optimization of Display Advertising with Ad Exchange
  - Belsairo, Feldman, M., Muthukrishnan, 2011

In Production

- Smooth Delivery of Display Ads
  - Delivery of impressions throughout time should follow the traffic smoothly.
  - Model this with multiple nested capacity constraints: $1 - 1/e$-competitive algorithm for this extension.
  - Bhalgat, Feldman, M., 2011
- Combined Allocation with Ad Exchange
  - Yield Optimization of Display Advertising with Ad Exchange
  - Belsairo, Feldman, M., Muthukrishnan, 2011
- Re-act adaptively and quickly to changes in traffic:
  - Use a control loop on the dual variable.
  - Assymptotically optimal policy: B. Tan and R. Srikant, 2011

Online Ad Allocation: Interesting Problems

- Online Stochastic DA
  - Simultaneous online worst-case & stochastic optimization.
  - Tradeoff between delivery penalty and efficiency: Covering Constraints?
  - More complex stochastic modeling (drift, seasonality, etc.)

Outline: Online Allocation

- Online Stochastic Ad Allocation
  - Online Stochastic Packing
  - Online Generalized Assignment (with free disposal)
  - Experimental Results
  - Online Stochastic Weighted Matching

Online Stochastic Weighted Matching

"ALG is α-approximation?" if $\frac{E[\mathrm{ALG}(H)]}{E[\mathrm{OPT}(H)]} \geq \alpha$

Power of Two Choices

- Offline:
  1. Find an optimal fractional solution x_e^* to a discounted matching LP, where $x_e \leq 1 - \frac{1}{e}$.
  2. Sample a matching M_s from x^*.
  3. Let $M' = M_1 \backslash M_s$ where M_1 is the maximum weighted matching.
- Online: try the edge in M_s first, and if it doesn't work, try M'.

Online Stochastic Weighted Matching

"ALG is α-approximation?" if $\frac{E[\mathrm{ALG}(H)]}{E[\mathrm{OPT}(H)]} \geq \alpha$

Power of Two Choices

- Offline:
  1. Find an optimal fractional solution x_e^* to a discounted matching LP, where $x_e \leq 1 - \frac{1}{e}$.
  2. Sample a matching M_s from x^*.
  3. Let $M' = M_1 \backslash M_s$ where M_1 is the maximum weighted matching.
- Online: try the edge in M_s first, and if it doesn't work, try M'.
- Thm: Approximation factor is better than 0.66.

(Haeupler, M., ZadiMoghaddam, 2011).

Open Problems

- Online Stochastic Display Ad Allocation
  - Simultaneous online worst-case & stochastic optimization.
  - Tradeoff between delivery penalty and efficiency: Covering Constraints?
  - More complex stochastic modeling (drift, seasonality, etc.)
- Online Stochastic Weighted Matching
  - Online Stochastic Matching: Close gap between 0.703 & 0.81.
  - Online Stochastic Weighted Matching: Power of many choices? Better lower bound?
  - Online Weighted Matching (with Free Disposal): Is $1 - 1/e$-approximation possible?

Contract-based Ad Delivery: Outline

- Basic Information
- Ad Serving.
  - Targeting.
  - Online Allocation
- Ad Planning: Reservation

Display Ad Delivery: Overview

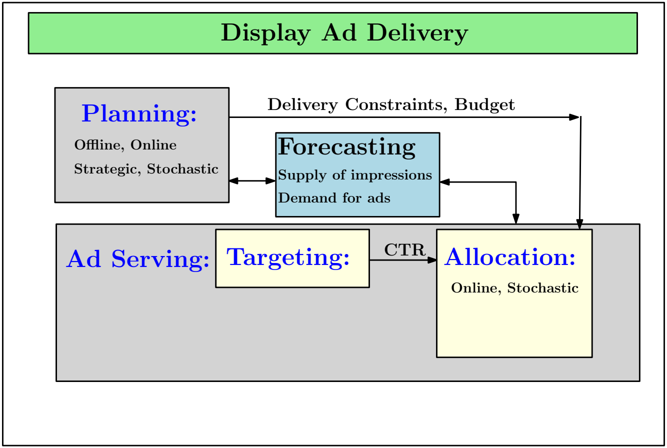

Ad Planning: Research Issues

Which set of contracts should we accept?

Ad Planning: Research Issues

Which set of contracts should we accept?

Related Research Issues.

- Pricing Uncertain Inventory.
  - Stochastic Supply and Demand
  - Bundling Opportunities
- Online Mechanisms for Signing Contracts.
- Contracts with Delivery Penalty.
- Offline Optimization of Contracts.

General Ad Planning: Weighted Matching

- n advertisers, and set Y of impressions (items).
- Each advertiser i
  - Interested in a set J_i of impressions, (e.g, young women in Seattle),
  - Needs d_i impressions (Demand),
  - Value v_{it} (or Bid b_i) for each impression t,

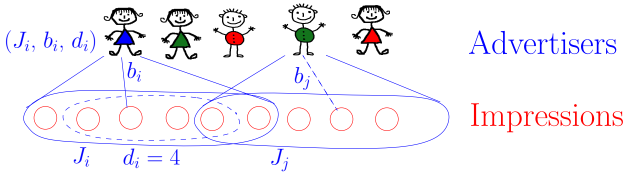

Efficiency (or Revenue) Maximization: Find an assignment with the maximum value.

Contracts with Delivery Penalty

If we don't meet the demand of an advertiser this month, we should give him/her free impressions in the next month.

Contracts with Delivery Penalty

If we don't meet the demand of an advertiser this month, we should give him/her free impressions in the next month.

- Each advertiser i
  - Needs d_i impressions,
  - Bids b_i for each impression,
  - Penalty λb_i for not satisfying each unit (Guaranteed Delivery).

Ad Planning With Penalties

- n advertisers, and set Y of impressions (items).
- Each advertiser i
  - Interested in set J_i of impressions, (e.g, young women in Seattle),
  - Bids b_i for each impression,
  - Needs d_i impressions,
  - Penalty λb_i for not satisfying each unit (Guaranteed Delivery).

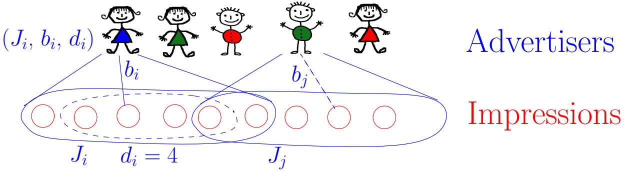

Ad Planning With Penalties

- n advertisers, and set Y of impressions (items).
- Each advertiser i
  - Interested in set J_i of impressions, (e.g, young women in Seattle),
  - Bids b_i for each impression,
  - Needs d_i impressions,
  - Penalty λb_i for not satisfying each unit (Guaranteed Delivery).

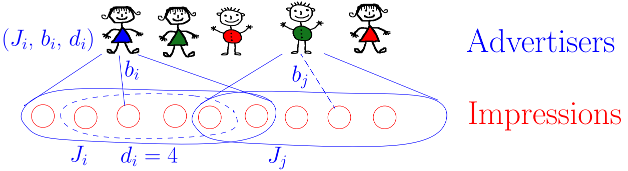

Goal: Choose a set T of advertisers to maximize revenue, $f(T)$.

Relation to Weighted Matching

If we give q items to advertiser i, we get

$$qb_i - \lambda b_i(d_i - q) = q(1+\lambda)b_i - d_i\lambda b_i = q(1+\lambda)b_i - c_i$$

Relation to Weighted Matching

If we give q items to advertiser i, we get

$$qb_i - \lambda b_i(d_i - q) = q(1+\lambda)b_i - d_i\lambda b_i = q(1+\lambda)b_i - c_i$$

If we commit to a set T of advertisers:

$$f(T) = P(T) - \sum_{i\in T} c_i = P(T) - C(T).$$

- $C(T) = \sum_{i\in T} c_i$ where $c_i = d_i\lambda b_i$.
- $P(T)$ be the maximum weighted matching in the following bipartite graph with $(1+\lambda)b_i$ as weights of edges.
  - Advertiser $i \in T$ needs at most d_i ads.
  - Each ad can go to at most one advertiser.

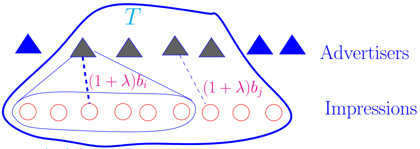

$P(T) =$ Maximum weighted matching for this graph.

Discussion Summary

Given a set T of advertisers, maximizing $f(T)$ is easy as it is a maximum weighted matching.

Discussion Summary

Given a set T of advertisers, maximizing $f(T)$ is easy as it is a maximum weighted matching.

Discussion Summary

Given a set T of advertisers, maximizing $f(T)$ is easy as it is a maximum weighted matching.

Challenge: Which set of advertisers T should we accept to maximize $f(T)$?

Note:

- Offline Optimization.
- Can be used in a negotiation process.
- Can have different penalty factors λ_i for each advertiser i.

Ad Planning with Penalties

Feige, Immorlica, M., Nazerzadeh, 2008.

- Hardness: No Constant-factor Approximation.
- Heuristic Greedy Algorithms:
  - Simple Greedy Algorithm: Bicriteria Approximation.
    - At each step, add the advertiser with the maximum profit-per-impression fixing the existing assignment.

Ad Planning with Penalties

Feige, Immorlica, M., Nazerzadeh, 2008.

- Hardness: No Constant-factor Approximation.
- Heuristic Greedy Algorithms:
  - Simple Greedy Algorithm: Bicriteria Approximation.
    - At each step, add the advertiser with the maximum profit-per-impression fixing the existing assignment.

  Theorem: This is a good approximation compared to the optimum with larger penalty factor.

Ad Planning with Penalties

Feige, Immorlica, M., Nazerzadeh, 2008.

- Hardness: No Constant-factor Approximation.
- Heuristic Greedy Algorithms:
  - Simple Greedy Algorithm: Bicriteria Approximation.
    - At each step, add the advertiser with the maximum profit-per-impression fixing the existing assignment.

    Theorem: This is a good approximation compared to the optimum with larger penalty factor.
  - Greedy-Rate Algorithm: Structural Approximation.

Greedy Algorithms

Greedy-Rate Algorithm:

- At each step, add the advertiser with the maximum marginal-profit per marginal-cost ratio.

Greedy Algorithms

Greedy-Rate Algorithm:

- At each step, add the advertiser with the maximum marginal-profit per marginal-cost ratio.

T: Set of advertisers we committed to (initialized to $\emptyset$)

- At each step, add an advertiser $i \in X \setminus S$ to T that maximizes $\frac{P(S \cup \{i\}) - P(S)}{c_i}$, if $P(S \cup \{i\}) - P(S) - c_i > 0$.

Greedy Algorithms

Greedy-Rate Algorithm:

- At each step, add the advertiser with the maximum marginal-profit per marginal-cost ratio.

T: Set of advertisers we committed to (initialized to $\emptyset$)

- At each step, add an advertiser $i \in X \backslash S$ to T that maximizes $\frac{P(S \cup \{i\}) - P(S)}{c_i}$, if $P(S \cup \{i\}) - P(S) - c_i > 0$.

Theorem: The greedy-rate algorithm achieves the best structural approximation.

Greedy Algorithms

Greedy-Rate Algorithm:

- At each step, add the advertiser with the maximum marginal-profit per marginal-cost ratio.

T: Set of advertisers we committed to (initialized to $\emptyset$)

- At each step, add an advertiser $i \in X \backslash S$ to T that maximizes $\frac{P(S \cup \{i\}) - P(S)}{c_i}$, if $P(S \cup \{i\}) - P(S) - c_i > 0$.

Theorem: The greedy-rate algorithm achieves the best structural approximation.
Proof: Uses submodularity of $P(T)$, and $f(T)$.

- The approximation factor of the algorithm is a function of the structure of the solution, i.e., a Signature: $\alpha = \frac{C(\text{OPT})}{P(\text{OPT})}$.
- Greedy-rate algorithm achieves at least factor $\frac{1-\alpha-\alpha \ln \frac{1}{\alpha}}{1-\alpha}$ (improving factor $1 + \alpha - 2\sqrt{\alpha}$).

Ad Planning with Delivery Penalty

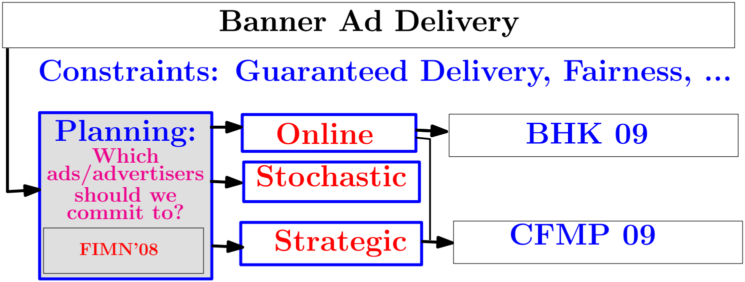

- Babaiof, Hartline, Kleinberg: Online Algorithms with Buyback
- Special cases: Single item, Matroid, Knapsack

- Constantin, Feldman, Muthkrishnan, Pal: Ad slotting with Cancellations.
- Special case: Demand=1, Truthful Mechanism: Constant-factor

Ad Planning with Delivery Penalty

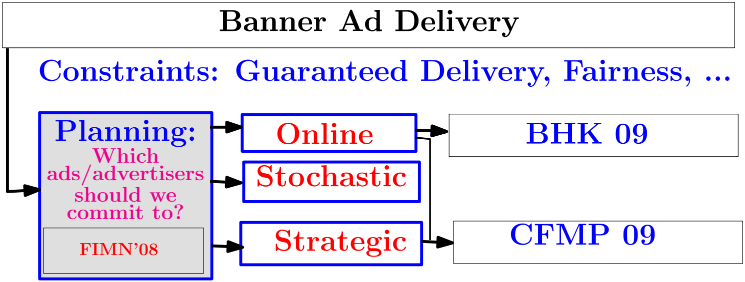

- Babaiof, Hartline, Kleinberg: Online Algorithms with Buyback
- Special cases: Single item, Matroid, Knapsack

- Constantin, Feldman, Muthkrishnan, Pal: Ad slotting with Cancellations.
- Special case: Demand=1, Truthful Mechanism: Constant-factor

Interesting Algorithmic Problem:

- **Online mechanism for general ad planning with delivery penalty: bicriteria approximation?**

Ad Planning: Offline Optimization of Contracts

- Fast Algorithms to Verify Feasibility of Contracts:
  - Lopsided Bipartite Graphs.
  - Improved Algorithms for Bipartite Network Flow: Ahuja, Orlin, Stein, Tarjan 94
    - Faster Algorithm for max-flow: running time depends on the size of smaller part.

Ad Planning: Offline Optimization of Contracts

- Fast Algorithms to Verify Feasibility of Contracts:
  - Lopsided Bipartite Graphs.
  - Improved Algorithms for Bipartite Network Flow: Ahuja, Orlin, Stein, Tarjan 94
    - Faster Algorithm for max-flow: running time depends on the size of smaller part.
  - Sampling for Max-Cardinality Matching: Charles, Chickering, Devanur, Jain, Sanghi 2010,
    - Sampling and Concise Allocation.
    - Algorithm: Iteratively assign nodes, and minimize the future failure probability.
    - WHP verifies if there exists a feasible matching.

Ad Planning: Offline Optimization of Contracts

- Fast Algorithms to Verify Feasibility of Contracts:
  - Lopsided Bipartite Graphs.
  - Improved Algorithms for Bipartite Network Flow: Ahuja, Orlin, Stein, Tarjan 94
    - Faster Algorithm for max-flow: running time depends on the size of smaller part.
  - Sampling for Max-Cardinality Matching: Charles, Chickering, Devanur, Jain, Sanghi 2010,
    - Sampling and Concise Allocation.
    - Algorithm: Iteratively assign nodes, and minimize the future failure probability.
    - WHP verifies if there exists a feasible matching.
- Online algorithms for accepting contracts: Alaei, Arcuate, Khuller, Ma, Malekian, Tomlin 2009
  - Utility model to combine contract-based advertisers & sales-based advertisers.
  - Online algorithm for accepting contracts (under assumptions)

Outline of this talk

- Ad delivery for contract based settings
  - Planning
  - Ad Serving
- **Ad serving in repeated auction settings**
  - General architecture.
  - Allocation for budget constrained advertisers.
- Other interactions
  - Learning + allocation
  - Learning + auction
  - Auction + contracts

Outline of this talk

- Ad delivery for contract based settings
  - Planning
  - Ad Serving
- Ad serving in repeated auction settings
  - General architecture.
  - Allocation for budget constrained advertisers.
- **Other interactions**
  - Learning + allocation
  - Learning + auction
  - Auction + contracts

Three main theory/practice problems

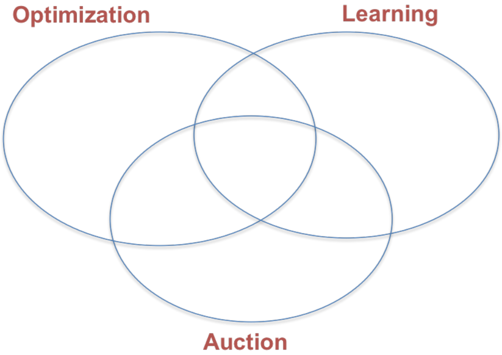

Combined Allocation with AdX: Objective

Short-term: boost revenue from AdX

VS

Long-term: prioritize quality of guaranteed contracts.

Find allocation policy to maximize

$$\text{yield} = \text{revenue(AdX)} + \gamma \cdot \text{quality(advertisers)},$$

where $\gamma \geq 0$ is a tradeoff paramater (or Lagrange multiplier).

Publisher's Decisions

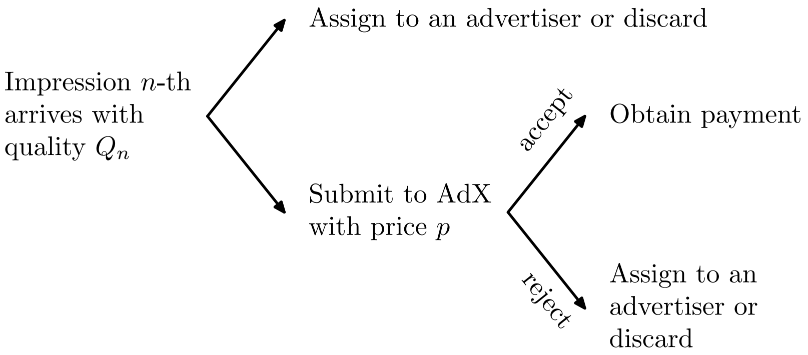

Publisher's Decisions

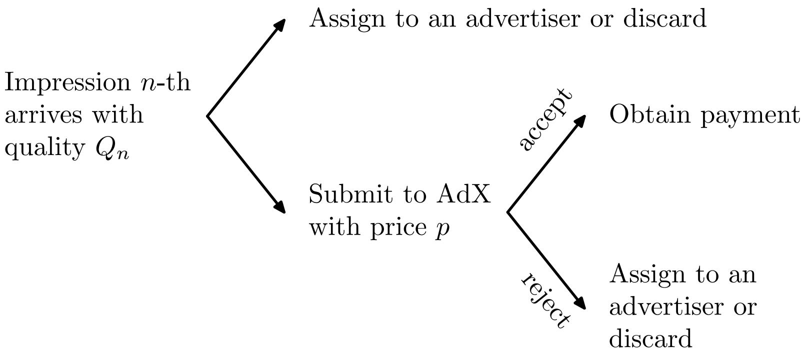

Optimal to always test the exchange!

Impact of AdX

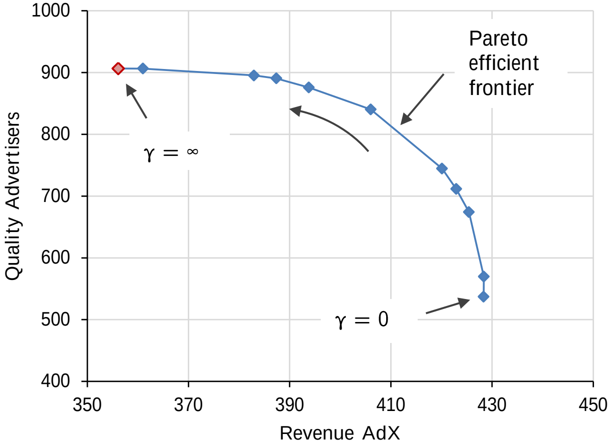

- $\gamma = 0$: maximum revenue from AdX.
- $\gamma = \infty$: maximum placement quality for contracts.

Display Ad Delivery

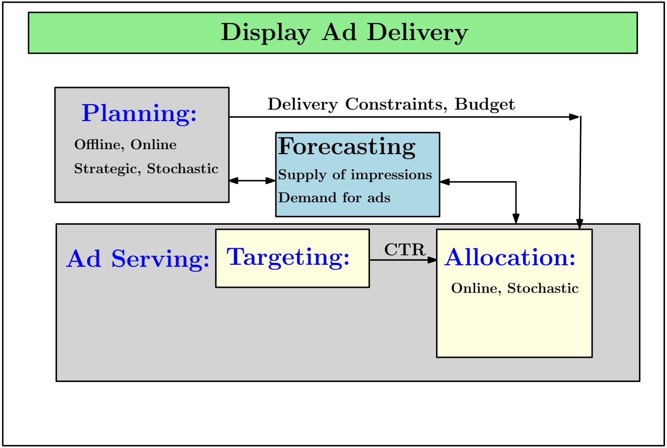

Display Ad Delivery

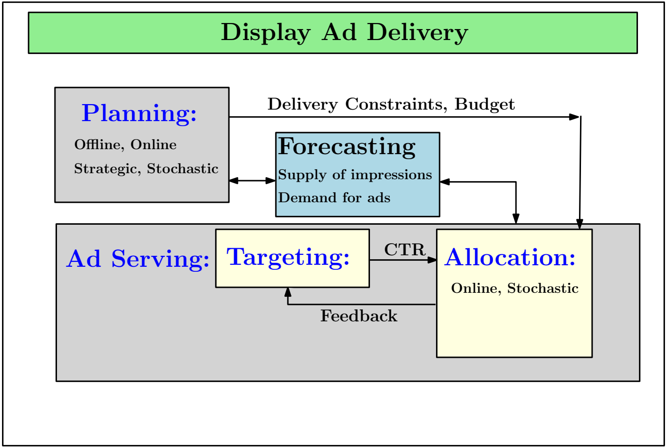

Online Learning & Allocation

- Value: Estimated Click-Through-Rate (CTR).

Online Learning & Allocation

- Value: Estimated Click-Through-Rate (CTR).
- Combined online capacity planning & learning?
  - Budgeted Active Learning
    - Madani, Lizotte, Greiner 2004, Active Model Selection.

Online Learning & Allocation

- Value: Estimated Click-Through-Rate (CTR).
- Combined online capacity planning & learning?
  - Budgeted Active Learning
    - Madani, Lizotte, Greiner 2004, Active Model Selection.
  - Bayesian Budgeted Multi-armed Bandits:
    - Guha, Munagala, Multi-armed Bandits with Metric Switching Costs.
    - Goel, Khanna, Null, The Ratio Index for Budgeted Learning, with Applications.
    - Guha, Munagala, Pal, Multi-armed Bandit with Delayed Feedback.

Online Learning & Allocation

- Value: Estimated Click-Through-Rate (CTR).
- Combined online capacity planning & learning?
  - Budgeted Active Learning
    - Madani, Lizotte, Greiner 2004, Active Model Selection.
  - Bayesian Budgeted Multi-armed Bandits:
    - Guha, Munagala, Multi-armed Bandits with Metric Switching Costs.
    - Goel, Khanna, Null, The Ratio Index for Budgeted Learning, with Applications.
    - Guha, Munagala, Pal, Multi-armed Bandit with Delayed Feedback.
  - Budgeted Unknown-CTR Multi-armed Bandit
    - Pandey, Olston 2007, Handling Advertisement of Unknown Quality.

Online CTR Learning: Mixed Explore/Exploit

- Pandey, Olston 2007, Handling Advertisement of Unknown Quality.

Online CTR Learning: Mixed Explore/Exploit

- Pandey, Olston 2007, Handling Advertisement of Unknown Quality.
- Algorithm: Revised Greedy
  - Upon arrival of query of type i, assign it to an ad a maximizing $P_{ia} = (\hat{c}_{ia} + \sqrt{\frac{2 \ln n_i}{n_{ia}}})b_{ia}$ where $\hat{c}_{ia}$ is the current estimate of CTR, n_{ia} is the number of times i has been assigned to a, n_i is the number of queries of type i so far.

Online CTR Learning: Mixed Explore/Exploit

- Pandey, Olston 2007, Handling Advertisement of Unknown Quality.
- Algorithm: Revised Greedy
  - Upon arrival of query of type i, assign it to an ad a maximizing $P_{ia} = (\hat{c}_{ia} + \sqrt{\frac{2 \ln n_i}{n_{ia}}}) b_{ia}$ where $\hat{c}_{ia}$ is the current estimate of CTR, n_{ia} is the number of times i has been assigned to a, n_i is the number of queries of type i so far.
- Thm[PO07]: ALG $\geq \frac{\text{OPT}}{2} - O(\ln n)$ where n is the number of arrivals.

Hybrid ad serving: Contracts + Spot Auctions

Given a page view, and two types of advertisers:

- Contract-based.
- Auction-based.

Hybrid ad serving: Contracts + Spot Auctions

Given a page view, and two types of advertisers:

- Contract-based.
- Auction-based.

- Decide who wins and how much do they pay.
- Requirements:
  - For each contract-advertiser, meet its demand.
  - Implement the scheme using proxy-bidding for contract-advertisers in the spot auction.

Hybrid ad serving: Contracts + Spot Auctions

- Naive solution: If a contract-adv is eligible and has not finished demand, then let it win the spot. Bid infinity for all auctions.

Hybrid ad serving: Contracts + Spot Auctions

- Naive solution: If a contract-adv is eligible and has not finished demand, then let it win the spot. Bid infinity for all auctions.

- Optimize for revenue: If the auction pressure (price) is low then let the contract-adv win. Bid a low bid for all auctions.

Hybrid ad serving: Contracts + Spot Auctions

- Naive solution: If a contract-adv is eligible and has not finished demand, then let it win the spot. Bid infinity for all auctions.

- Optimize for revenue: If the auction pressure (price) is low then let the contract-adv win. Bid a low bid for all auctions.
  - Unfair to contract-adv, since low auction-price $\Rightarrow$ it is a lower value impression.

Hybrid ad serving: Contracts + Spot Auctions

- Naive solution: If a contract-adv is eligible and has not finished demand, then let it win the spot. Bid infinity for all auctions.

- Optimize for revenue: If the auction pressure (price) is low then let the contract-adv win. Bid a low bid for all auctions.
  - Unfair to contract-adv, since low auction-price $\Rightarrow$ it is a lower value impression.

- Ideally:
  - Provide contract-adv with a representative allocation, an equal slice of impressions from each price-point.
  - A price-oblivious scheme, i.e., bid without seeing the auction bids.
  - Revenue per auction: average auction-price of impressions given away to contract-advertisers is at most some target t.

Obtaining representative allocations

Two main ideas:

1. Can implement any decreasing function $a(p)$ for fraction of impressions of auction-price p.

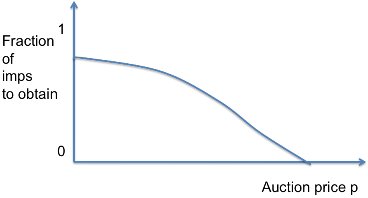

Obtaining representative allocations

Two main ideas:

1. Can implement any decreasing function $a(p)$ for fraction of impressions of auction-price p.

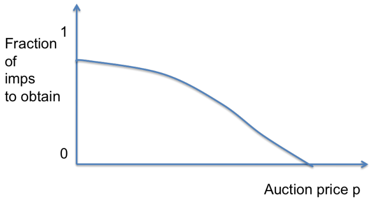

2. Solve the system for well chosen distance functions:

$$\text{Minimize dist(U, a)}$$

$$\text{s.t.: } \int_p a(p)f(p)dp = d$$

$$\int_p pa(p)f(p)dp \leq td$$

Online Learning & Auction Incentives

[Devanur,Kakade'09, Babaioff,Sharma,Slivkins'09]

- Multi-Armed Bandit algorithms achieve an "implicit" exploration-exploitation tradeoff to get a regret of $O(\sqrt{T})$ (e.g., UCB).
- Can these be run in tandem with truthful auctions? (e.g., 2nd price for a single slot).

Online Learning & Auction Incentives

[Devanur,Kakade'09, Babaioff,Sharma,Slivkins'09]

- Multi-Armed Bandit algorithms achieve an "implicit" exploration-exploitation tradeoff to get a regret of $O(\sqrt{T})$ (e.g., UCB).
- Can these be run in tandem with truthful auctions? (e.g., 2nd price for a single slot).
- A naive explore-exploit method gets $O(T^{2/3})$ regret:
  - Explore ads for the first *phase*, giving them out for free.
  - Fix the CTRs thus learned in the first phase.
  - Run 2nd price auction for the 2nd phase.

Online Learning & Auction Incentives

[Devanur,Kakade'09, Babaioff,Sharma,Slivkins'09]

- Multi-Armed Bandit algorithms achieve an "implicit" exploration-exploitation tradeoff to get a regret of $O(\sqrt{T})$ (e.g., UCB).
- Can these be run in tandem with truthful auctions? (e.g., 2nd price for a single slot).
- A naive explore-exploit method gets $O(T^{2/3})$ regret:
  - Explore ads for the first *phase*, giving them out for free.
  - Fix the CTRs thus learned in the first phase.
  - Run 2nd price auction for the 2nd phase.
- Can you do better that this simpe decoupling?
- No!

Theorem

[DK09,BSS09] For every truthful auction (under certain assumptions), there exist bids, ctrs, s.t. regret $= \Omega(T^{2/3})$.

Display Ad Delivery

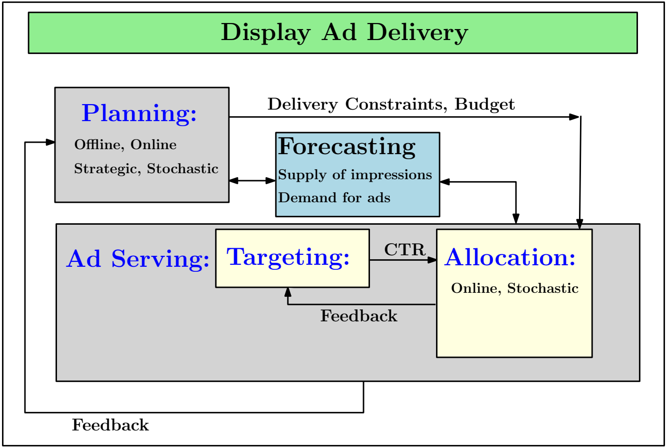

Open Problems:

- **Optimal combined online allocation & learning.**
- **Feature selection and correlation in learning CTR.**
- **Optimal combined stochastic planning and serving?**

Thank You