Computational Analysis of Perfect-Information Position Auctions

David Robert Martin Thompson and Kevin Leyton-Brown

June 16, 2009

Motivation

- Position auctions:
  - Billion dollar revenue stream for search engines
  - Auctions evolved in an ad hoc way
  - Auction theorists are catching up: starting to understand how the auctions perform under simplifying assumptions.
  - Performance: putting good ads in good spaces, and generating revenue
  - Which auction performs best?
- Our contribution: computational method for comparing auction performance quantitatively.

Outline

Outline

Use of position auctions

- **GFP**: Yahoo! and Overture 1997-2002
- **uGSP**: Yahoo! 2002-2007
- **wGSP**: Google, MSN Live, Yahoo! 2007-Present

Use of position auctions

- **GFP**: Yahoo! and Overture 1997-2002
- **uGSP**: Yahoo! 2002-2007
- **wGSP**: Google, MSN Live, Yahoo! 2007-Present

Question

Is wGSP better than GFP and uGSP?

Use of position auctions

- **GFP**: Yahoo! and Overture 1997-2002
- **uGSP**: Yahoo! 2002-2007
- **wGSP**: Google, MSN Live, Yahoo! 2007-Present

Question

Is wGSP better than GFP and uGSP?

- Better by what metric? Revenue, efficiency

How theorists study position auctions

- Terminology:
  - Nash equilibrium: every bidder is acting to maximize her own payoff.
  - Perfect-information game: every bidder knows every other's value / CTR.
  - VCG: a perfectly economically-efficient auction (a common theoretical benchmark)

How theorists study position auctions

- Terminology:
  - Nash equilibrium: every bidder is acting to maximize her own payoff.
  - Perfect-information game: every bidder knows every other's value / CTR.
  - VCG: a perfectly economically-efficient auction (a common theoretical benchmark)
- They also need a structural model of values / CTRs...

Edelman, Ostrovsky, Schwarz (2007)

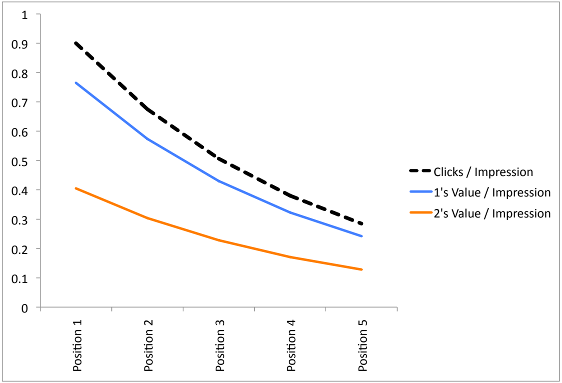

Varian (2007)

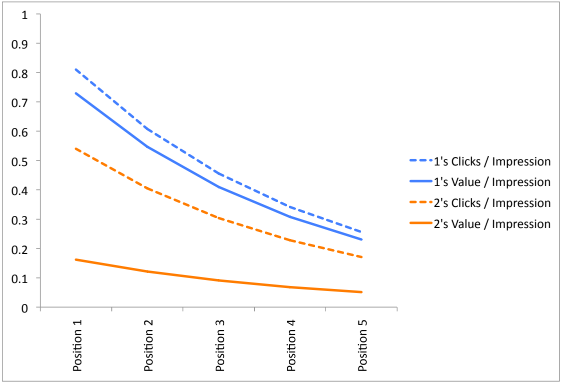
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Position 1
Position 2
Position 3
Position 4
Position 5
1's Clicks / Impression
1's Value / Impression
2's Clicks / Impression
2's Value / Impression

Blumrosen, Hartline, Nong (2008)

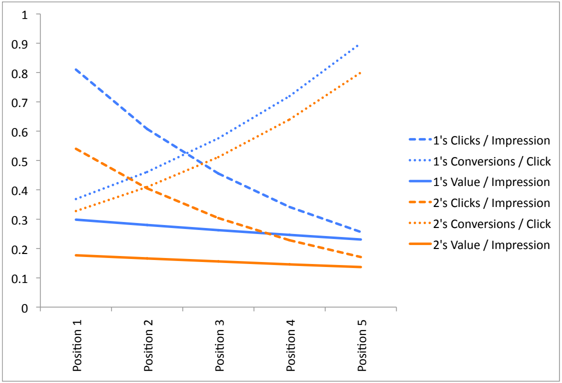

Benisch, Sadeh, Sandholm (2008)

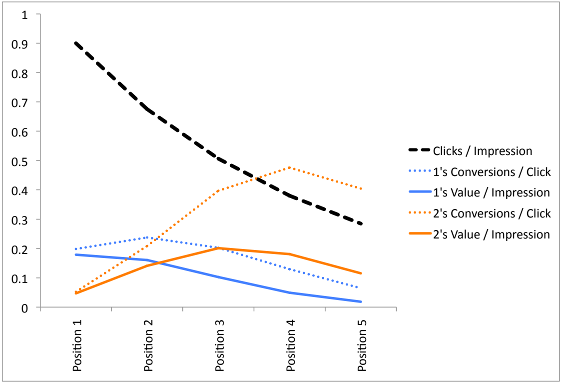

Outline

Action Graph Games [Bhat, Leyton-Brown, 2004]

- Graphical model like Bayes nets, GAI nets or graphical games
  - Nodes are variables, directed edges denote conditional independence
  - Representation is polynomial for graphs of bounded in-degree
- Nodes represent actions: variable = how many play that action?

Action Graph Games [Bhat, Leyton-Brown, 2004]

- Graphical model like Bayes nets, GAI nets or graphical games
  - Nodes are variables, directed edges denote conditional independence
  - Representation is polynomial for graphs of bounded in-degree
- Nodes represent actions: variable = how many play that action?
- Nodes can also be simple functions (e.g. sum, argmax)

Action Graph Games [Bhat, Leyton-Brown, 2004]

- Graphical model like Bayes nets, GAI nets or graphical games
  - Nodes are variables, directed edges denote conditional independence
  - Representation is polynomial for graphs of bounded in-degree
- Nodes represent actions: variable = how many play that action?
- Nodes can also be simple functions (e.g. sum, argmax)
- Expected utility is polynomial in input [Jiang, Leyton-Brown, 2006]
  - Exponential speedup for solvers that use expected utility in inner loop

Representing GFP as AGG

- n bidders, m bid increments ($O(nm)$ actions)

Representing GFP as AGG

- n bidders, m bid increments ($O(nm)$ actions)
- For each action, payoff only depends on position
- Sufficient statistic: How many bid the same? How many bid higher? $O(n^2)$
  - Easily computed with sum nodes

Representing GFP as AGG

- n bidders, m bid increments ($O(nm)$ actions)
- For each action, payoff only depends on position
- Sufficient statistic: How many bid the same? How many bid higher? $O(n^2)$
  - Easily computed with sum nodes
- AGG representation $O(n^3m)$ (vs. $O(nm^n)$ in normal form)

Representing GSP as AGG

- Additional sufficient statistic: What is the next highest bid? $O(nm)$
  - Easily computed with argmax nodes

Representing GSP as AGG

- Additional sufficient statistic: What is the next highest bid? $O(nm)$
  - Easily computed with argmax nodes
- AGG representation $O(n^4m^2)$ (vs. $O(nm^n)$ in normal form)

Outline

Problem instances

- 4 sizes (5-10 bidders, 5-40 increments)
- 4 preference distributions: EOS, V, BHN, BSS (assume uniform distributions where unspecified)
- 100 draws from each distribution, size $=$ 1600 "preference instances"

Problem instances

- 4 sizes (5-10 bidders, 5-40 increments)
- 4 preference distributions: EOS, V, BHN, BSS (assume uniform distributions where unspecified)
- 100 draws from each distribution, size
  $=$ 1600 "preference instances"
- 3 auctions: GFP, uGSP, wGSP
  $=$ 4800 games

Solving games

- Remove dominated strategies: bids above an agent's (maximum) value, strategically redundant bids

Solving games

- Remove dominated strategies: bids above an agent's (maximum) value, strategically redundant bids
- Two solvers: simpdiv [Scarf, 1967] and gnm [Govindan, Wilson, 2005]
  - implemented in Gambit [McKevley et al, 2006] with AGG dynamic programming optimizations [Jiang, Leyton-Brown, 2006]
  - Run solvers 10 times (with 5 minute cutoff).

Equilibrium selection

- Problem: These games have many equilibria, and equilibrium selection matters. (Enumerating equilibria is infeasible.)

Equilibrium selection

- Problem: These games have many equilibria, and equilibrium selection matters. (Enumerating equilibria is infeasible.)
- We use local search to find (locally) extreme equilibria: min/max revenue/efficiency (4 different objectives).
- SLS algorithm: start from existing equilibria, random improving moves, restart given local optimum.

Statistical methods

- Blocking, means-of-means, bootstrapping test (across a pair of auctions)
- Non-parametric confidence interval on mean difference
  - Significant if entire $1 - \alpha$ confidence interval ≥ 0

Statistical methods

- Blocking, means-of-means, bootstrapping test (across a pair of auctions)
- Non-parametric confidence interval on mean difference
  - Significant if entire $1 - \alpha$ confidence interval ≥ 0
- Used Bonferroni correction (divide by number of tests, $|T| = 80$)
- * denotes significant for $\alpha = 0.05/|T|$
- ** denotes significant for $\alpha = 0.01/|T|$

Outline

Efficiency: (simplified) theoretical predictions

- In EOS and V models, wGSP is efficient in every "envy-free" Nash equilibrium [Edelman, et al., 2007; Varian, 2007].
- There are cases in BHN and BSS models, wGSP is not efficient in any Nash equilibrium [Blumrosen, et al., 2008; Benisch, et al., 2008].

Worst-case efficiency

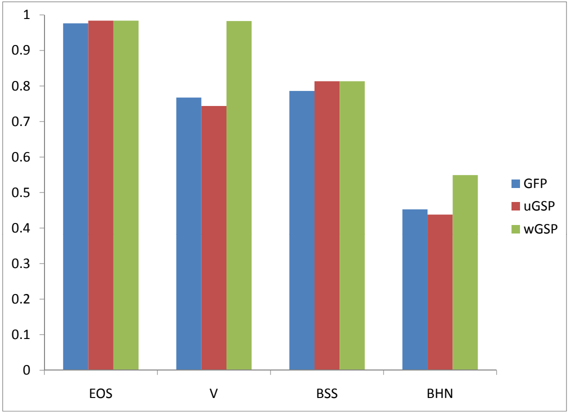

- (uGSP,GFP) $\leq$ wGSP $\leq$ discrete VCG $\leq$ VCG **

Revenue: (simplified) theoretical predictions

- In EOS and V models, wGSP beats VCG in every "envy-free" Nash equilibrium Edelman, et al., 2007; Varian, 2007].

V: revenue range

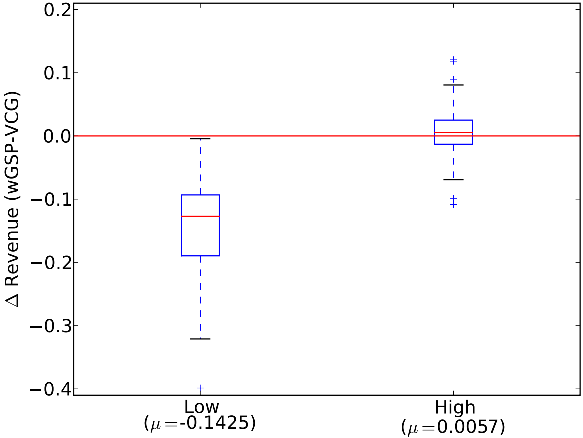

V: best-case revenue

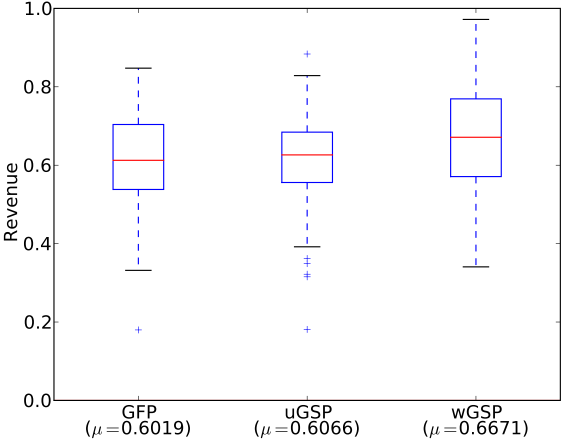

V: best-case revenue

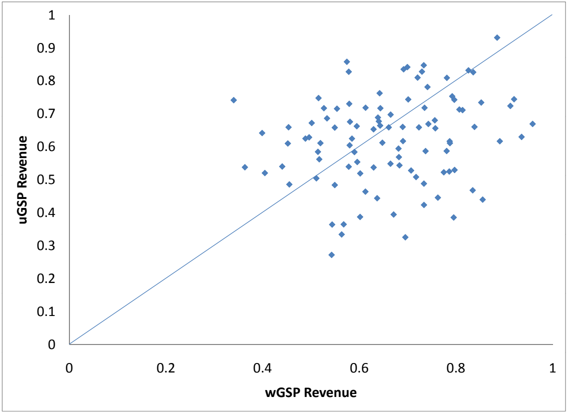

Outline

Conclusion

- This approach is possible and yields real economic insights!
- Efficiency: wGSP is more efficient (even in difficult models) and very robust to equilibrium selection.
- Revenue: Ranking is unclear. Equilibrium selection and instance details have large impact.
- Code and data are available at: `http://www.cs.ubc.ca/research/position_auctions/`

Future work

- Learning distributions from data
- Generalize representation to other models (e.g. with externalities)
- Better game solving techniques (e.g. provable bounds on revenue and welfare)
- Theoretical implications of results

The Quest for assets

- Our algorithm needs complete knowledge of advertisers' CTRs and values...

The Quest for assets

- Our algorithm needs complete knowledge of advertisers' CTRs and values...
- **The Good**: Lots of data on clicks and impressions
- **The Bad**: No data on bids or weights
- **The Wanted**: Data on conversions (or ideally, values)