

Algorithmic Game Theory

Auction theory in practice

Vangelis Markakis
markakis@gmail.com

Allocation rules and truthful mechanisms


- We recall first some definitions we saw in previous lectures
- Consider a mechanism with allocation rule \mathbf{x}
- Definition: An allocation rule is **monotone** if for every i , and every profile \mathbf{b}_{-i} , the allocation $x_i(z, \mathbf{b}_{-i})$ to i is non-decreasing in z
 - i.e., bidding higher can only get you more stuff

[Myerson '81]

- Theorem: For every single-parameter environment,
 - An allocation rule \mathbf{x} can be turned into a truthful mechanism if and only if it is monotone
 - If \mathbf{x} is monotone, then there is a unique payment rule \mathbf{p} , so that (\mathbf{x}, \mathbf{p}) is a truthful mechanism

Myerson's lemma and payment formula

- For the payment rule, we need to look for each bidder at the allocation function $x_i(z, \mathbf{b}_{-i})$
- For the single-item truthful auction:
 - Fix \mathbf{b}_{-i} and let $b^* = \max_{j \neq i} b_j$




Facts:

- For any fixed \mathbf{b}_{-i} , the allocation function is piecewise linear with 1 jump
- The Vickrey payment is precisely the value at which the jump happens
- The jump changes the allocation from 0 to 1 unit

Myerson's lemma and payment formula

For most scenarios of interest

- The allocation is piecewise linear with multiple jumps
- The jump determines how many extra units the bidder wins




- Suppose bidder i bids b_i
- Look at the jumps of $x_i(z, b_i)$ in the interval $[0, b_i]$
- Suppose we have k jumps
- Jump at $z_1 = w_1$
- Jump at $z_2 = w_2 - w_1$
- Jump at $z_3 = w_3 - w_2$
- ...
- Jump at $z_k = w_k - w_{k-1}$

Myerson's lemma and payment formula

For most scenarios of interest

- The allocation is piecewise linear with multiple jumps
- The jump determines how many extra units the bidder wins



Payment formula

- For each bidder i at a profile b , find all the jump points within $[0, b_i]$
- $$p_i(b) = \sum_j z_j \cdot [\text{jump at } z_j]$$
$$= \sum_j z_j \cdot [w_j - w_{j-1}]$$
- The formula can also be generalized for monotone but not piecewise linear functions

Sponsored Search Auctions

What is sponsored search?

Google™ Αναζήτηση [Σύνθετη Αναζήτηση](#) [Προτιμήσεις](#)

Αναζήτηση: παγκόσμιος ιστός σελίδες στα Ελληνικά σελίδες από Ελλάδα

Παγκόσμιος ιστός Αποτελέσματα 1 - 10 από περίπου 4.140.000 για **crm software**. (0,17 δευτερόλεπτα)

Interworks Web CRM Σύνδεσμοι διαφημιζομένων
www.interworks.gr Το πρώτο Web CRM στην Ελλάδα Δοκιμάστε το δωρεάν για 30 ημέρες

Goldmine CRM
www.alexandermooore.com Αυξήστε τις Πωλήσεις με το Νο1 CRM στις ΗΠΑ & 10 χρόνια στην Ελλάδα

Crm Software
www.CRMdesk.com Web-based Help Desk, Customer Service and Online Support Software

Σύνδεσμοι διαφημιζομένων

AuraPortal: BPMS with CRM
5 in 1: Process, CRM/E-Business, Intranet, Documents, ECM Portals
www.AuraPortal.com

SalesManager Hellas CRM
Διεθνώς καταξιωμένη λύση CRM πλήρως προσαρμοσμένη στην Ελληνική Αγορά
www.salesmanager.gr

Sales Plus CRM
Το CRM με εκατοντάδες εγκαταστάσεις σε Ελλάδα και εξωτερικό
www.orbit.gr/sales.html

EasyConsole eCRM
Σύστημα Διαχείρισης Πελατών (CRM) για Μικρές και Μεγάλες επιχειρήσεις
www.dynamicworks.eu

Εξοπλισμός κομμωτηρίων
Αναβαθμιστείτε σήμερα! 210.6396.937
Πελατολόγιο, Βαφές, Προϊόντα
www.easytouch.gr


Συμβουλή: [Αναζήτηση αποτελεσμάτων μόνο σε Ελληνικά](#). Μπορείτε να επιλέξετε τη γλώσσα αναζήτησης στη σελίδα [Προτιμήσεις](#)

[Διαχείριση πελατολογίου, Συναλλαγών, Πελατών, Πελατολόγιο ...](#)
Greek CRM software, database software, ΕΣΟΔΑ, ΕΞΟΔΑ, ... crm network fax software, καταχώρηση τιμολογίων, πρόγραμμα πελατων, εσοδα εξοδα, κεφαλαιο, ...
www.starmessage.gr/crm_software.html - 66k - [Προσωρινά αποθηκευμένη](#) - [Παρόμοιες σελίδες](#)

[CRM Software, Customer Relationship Management, CRM Solutions from ...](#) - [[Μετάφραση αυτής της σελίδας](#)]
CRM from Oncontact. Your source for customer relationship management or CRM software, CRM solutions and customer relationship management software.
www.oncontact.com/ - 12k - [Προσωρινά αποθηκευμένη](#) - [Παρόμοιες σελίδες](#)

CRM SOFTWARE - SalesManager Hellas - Customer Relationship .
Μία από τις πλέον σύγχρονες τάσεις της επιχειρηματικότητας αφορά στην « Διαχείριση των Σχέσεων με τους Πελάτες / Customer Relationship Management» ή « CRM». ...

What is sponsored search?



Advertising slots

Also try: [graduation gifts](#), [spencer gifts](#), [anniversary gifts](#), [More...](#)

Purse Hooks - Unique Gifts for Women
CWDdesignLLC.com - great collection of gold, silver and crystal purse h...
www.cwdesignllc.com

Gift Ideas: Christmas Gifts, Anniversary And Birthday
Gift ideas from popular to unique. Gifts.com's gift experts find great gifts for every occasion, including gift ideas for wedding presents, birthdays and anniversary.
www.gifts.com

| | |
|-------------------------------|-----------------------------------|
| Christmas | Gifts For Women |
| Anniversary | Teens |
| Birthday | Online Gift Cards |
| Gifts For Men | Wedding Registry |

[more results from gifts.com »](#)

RedEnvelope Official Site - Gifts for All Occasions
Find unique gifts of affordable luxury including jewelry and men's accessories as well as gourmet foods, flowers, gift baskets, toys, and personalized gifts.
Quick Links: [Jewelry](#) - [Sale](#) - [What's New](#)
www.redenvelope.com

Gift Ideas from Surprise.com
A community site where consumers share their best gift ideas with each other, organized into categories like unusual ... recommends creative gifts and shows ...
www.surprise.com

Purchase Gourmet Gifts Dean Deluca Online from &
Indulge your taste buds with Gourmet Gifts Dean Deluca from Dean & DeLuca, featuring a premium selection of fine gourmet food, gift baskets and more.

Try Our Auction Misspelling Tool & Save \$
Use the Free Misspelling Tool. Find Misspelled Auctions & Save Big.
www.AwesomeAuctionDeals.com

Find Christian Gifts
Join our community of Christians for Bible study, fellowship and more.
www.worldslastchance.com

American Metal Arts
Custom made Metal Art Hand Crafted and Made in the USA.
www.AmericanMetalArtsStudios.com

Wedding Gifts
Create unique gifts with your photos. T-shirts, mugs, postage & more.
www.Zazzle.com

Collectible Stock Certificates
We buy and sell collectible stock certificates. Old stock certificates...
www.giftofhistory.com

How does it work?

- For a fixed search term (e.g. *ipod*)
 - n advertisers
 - k slots (typically $k \ll n$)
 - An auction is run for every single search
 - Each advertiser (bidder) is interested in getting himself displayed in one of the slots
 - And usually they prefer a slot as high up as possible
 - Same auction is also run for related keywords (e.g. “buy ipod”, “cheap ipod”, “ipod purchase”, ...)
 - The advertiser can determine for which phrases to participate

How does it work?

- Bidders submit an initial budget which they can refresh weekly or monthly
- Bidders also submit an initial bid which they can adjust as often as they wish
- The auction selects the winners to be displayed
- Different charging models exist: **Pay Per Click, Pay Per Impression, Pay Per Transaction**
- Currently, most popular is Pay Per Click
- A bidder is charged only if someone clicks on the bidder's ad

The Actors

- The Search engine:
 - Wants to make as much revenue as possible
 - At the same time, wants to make sure users receive meaningful ads and bidders do not feel that they were overcharged
 - Big percentage of Google's revenue has been due to these auctions!
- The Bidders:
 - Want to occupy a high slot and pay as little as possible
- The Searchers:
 - Want to find the most relevant ads with respect to what they are looking for

Analyzing sponsored search auctions

- We will focus on the bidders' side
- Model parameters for each bidder i
 - Private information: v_i = maximum amount willing to pay per click = value/happiness derived from a click (private information)
 - Each bidder i submits a bid b_i for willingness to pay per click (b_i may differ from v_i)
 - We will ignore the budget parameter
 - In many cases, it is large enough and cannot affect the game
 - Hence, we have a single-parameter problem

Analyzing sponsored search auctions

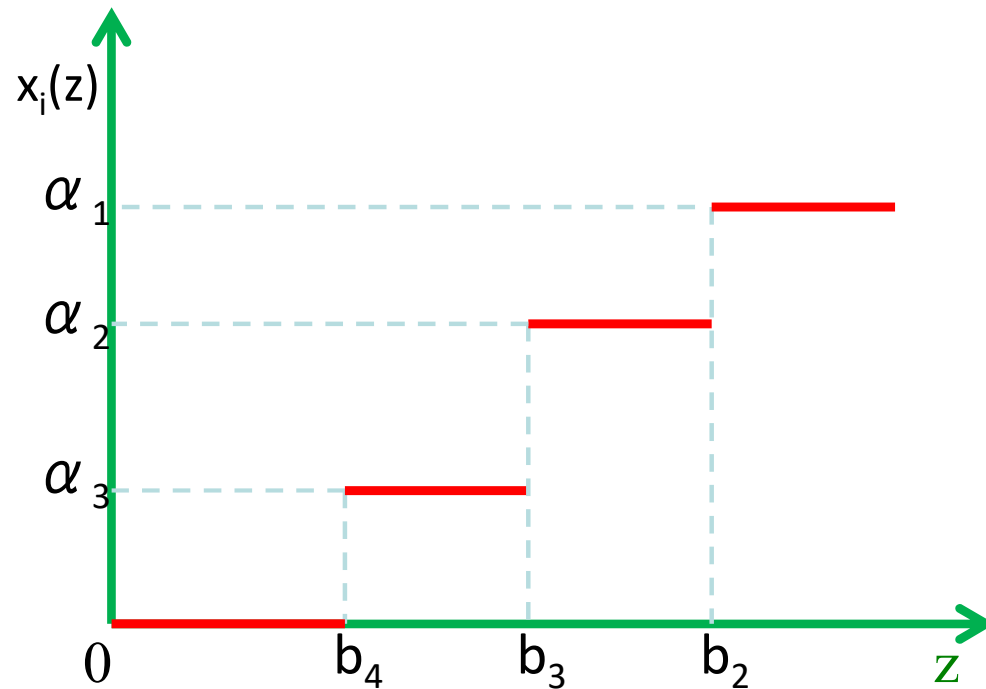
- We will focus on the bidders' side
- Model parameters for each slot j
 - α_j = Click-through rate (CTR) of slot j = probability that a user will click on slot j
 - Assume it is independent of who occupies slot j
 - We can generalize to the case where the rates are weighted by a quality score of the advertiser who takes each slot
 - The search engines update regularly the CTRs and statistics show that
$$\alpha_1 \geq \alpha_2 \geq \alpha_3 \geq \dots \geq \alpha_k$$
 - Users tend to click on higher slots
 - Validation also by eye-tracking experiments

Analyzing sponsored search auctions

- How shall we allocate the k slots to the n bidders?
- **Most natural allocation rule:** for $i=1$ to k , give to the i -th highest bidder the i -th best slot in terms of CTR
 - Remaining $n-k$ bidders do not win anything
- For convenience, assume that $b_1 \geq b_2 \geq b_3 \geq \dots \geq b_n$
- Expected value of a winning bidder i : $\alpha_i v_i$
- Is this rule monotone?
- Yes, bidding higher can only get you a better slot
- Hence we can apply Myerson's formula to find the payment rule
- For each bidder i , let $x_i(b_i, b_{-i}) \in \{0, \alpha_k, \alpha_{k-1}, \dots, \alpha_1\}$

Myerson's lemma for sponsored search auctions

- Let's analyze the highest bidder with bid b_1
- Suppose we have 3 slots and $n > 3$ bidders



- Look at the jumps of x_i in the interval $[0, b_1]$
- Jump at $b_4 = \alpha_3$
- Jump at $b_3 = \alpha_2 - \alpha_3$
- Jump at $b_2 = \alpha_1 - \alpha_2$

Total payment:

$$b_4 \alpha_3 + b_3 (\alpha_2 - \alpha_3) + b_2 (\alpha_1 - \alpha_2)$$

Myerson's lemma for sponsored search auctions

- More generally, for the i -th highest bidder, there will be $k-i+1$ jumps

$$p_i(\mathbf{b}) = \sum_{j=i}^k b_{j+1} [\alpha_j - \alpha_{j+1}]$$

- Under pay-per-click, no actual payment takes place at the end of every auction, unless there is a click by a user
- Need to scale so that expected per-click payment is $p_i(\mathbf{b})$
- Proposed per-click payment to bidder in i -th slot: $p_i(\mathbf{b}) / \alpha_i$
- By Myerson, no other payment can achieve truthfulness with the same allocation rule

Sponsored search auctions in practice

- In practice most engines do not use the payment of Myerson's lemma
- But they use the same allocation rule
- **The Generalized Second Price Mechanism (GSP)** - initial version:
 - The search engine ranks the bids in decreasing order:
 $b_1 \geq b_2 \geq \dots \geq b_n$
 - The i -th highest bidder takes the i -th best slot
 - Every time there is a click on slot i , bidder i pays b_{i+1}
 - There is also a reserve price (opening bid), initially the same for every keyword (\$0.1), later became keyword-dependent

The Generalized Second Price Mechanism (GSP)

- A better version:
 - The search engine keeps a quality score q_i for each bidder i
 - Yahoo, Bing (till a few years ago): q_i is the click-through rate of i (probability of a user clicking on an ad of bidder i)
 - Google: q_i depends on click-through rate, relevance of text and other factors
 - The search engine ranking is in decreasing order of $q_i \times b_i$
 $q_1 \times b_1 \geq q_2 \times b_2 \geq \dots \geq q_n \times b_n$
 - The first k bidders of the ranking are displayed in the k slots
 - Every time there is a click on slot i , bidder i pays the minimum bid required to keep his position, i.e. $(q_{i+1} \times b_{i+1}) / q_i$

The Generalized Second Price Mechanism (GSP)

- Myerson's lemma implies GSP cannot be truthful
 - Otherwise, its payment rule would coincide with the Myerson formula
- The deployment of GSP was probably just an educated guess
 - As an attempt to generalize the Vickrey auction and use something simple that looked close to truthful!
- Nevertheless...
 - For a long period, revenue from GSP was 95% of Google's revenue
 - Still nowadays an important percentage of search engines' revenue

Multi-unit auctions

Multi-unit Auctions

Auctions for selling multiple identical units of a **single good**

In practice:

- US Treasury notes, bonds
- UK electricity auctions (output of generators)
- Carbon allowance markets (pollution rights)
- Spectrum licences
- Various online sales

Multi-unit Auctions

Online sites offering multi-unit auctions

- US
 - www.onlineauction.com
- UK
 - uk.ebid.net
- Greece
 - www.ricardo.gr
 - Actually not any more...
- ...

Some Notation

- n bidders
- k available units of an indivisible good
- Bidder i has valuation function $v_i : \{0, 1, \dots, k\} \rightarrow \mathbb{R}$
 - $v_i(j)$ = value of bidder i for obtaining j units
- Representation with marginal valuations:
 - $m_i(j) = v_i(j) - v_i(j-1)$ = additional value for obtaining the j -th unit, if already given $j-1$ units
 - $(m_i(1), m_i(2), \dots, m_i(k))$: vector of marginal values

Submodular Valuations

- In the multi-unit setting, a valuation v_i is *submodular* iff

$$\forall x \leq y, \quad v_i(x + 1) - v_i(x) \geq v_i(y + 1) - v_i(y)$$

- Hence: $m_i(1) \geq m_i(2) \geq \dots \geq m_i(k)$ (decreasing marginal values)
- In many multi-unit auctions, bidders are asked to submit a submodular valuation
 - Makes sense due to the saturation of getting more and more units
- **Valuation compression:** Even if bidders are not submodular, they would still have to express their preferences by a submodular function

A Bidding Format for Multi-unit Auctions

- Used in various multi-unit auctions

[Krishna '02, Ch. 12-13, Milgrom '04, Ch. 7]

1. The auctioneer asks each bidder to submit a vector of decreasing marginal bids
 - $\mathbf{b}_i = (b_i(1), b_i(2), \dots, b_i(k))$
 - $b_i(1) \geq b_i(2) \geq \dots \geq b_i(k)$
2. The bids are ranked in decreasing order and the k highest win the units

Simplified format in some cases: Uniform bidding, i.e., ask for a bid per unit + number of units demanded

Example



$$\mathbf{b}_1 = (45, 42, 31, 22, 15)$$




$$\mathbf{b}_2 = (35, 27, 20, 12, 7)$$



$$\mathbf{b}_3 = (40, 33, 24, 14, 9)$$

Example



How should we charge the winners?


Pricing Rules

1. Multi-unit Vickrey auction (VCG) [Vickrey '61]
 - Each bidder pays the externality he causes to the others
 - Generalization of single-item 2nd price auction
 - Good theoretical properties, truthful, but barely used in practice
2. Discriminatory Price Auction (DPA)
 - Bidders pay their bids for the units won
 - Generalization of 1st price auction
 - Not truthful, but widely used in practice

Pricing Rules (cont'd)

3. Uniform Price Auction (UPA) [Friedman 1960]
 - Same price for every unit
 - Interval of prices to pick from:
 [highest losing bid, lowest winning bid]
 - This lecture: price = highest losing bid
 - For 1 unit, same as Vickrey auction
 - For ≥ 2 units, not truthful, but widely used in practice

Example Revisited



Interval of candidate prices for UPA = [31, 33]
 Uniform price = 31

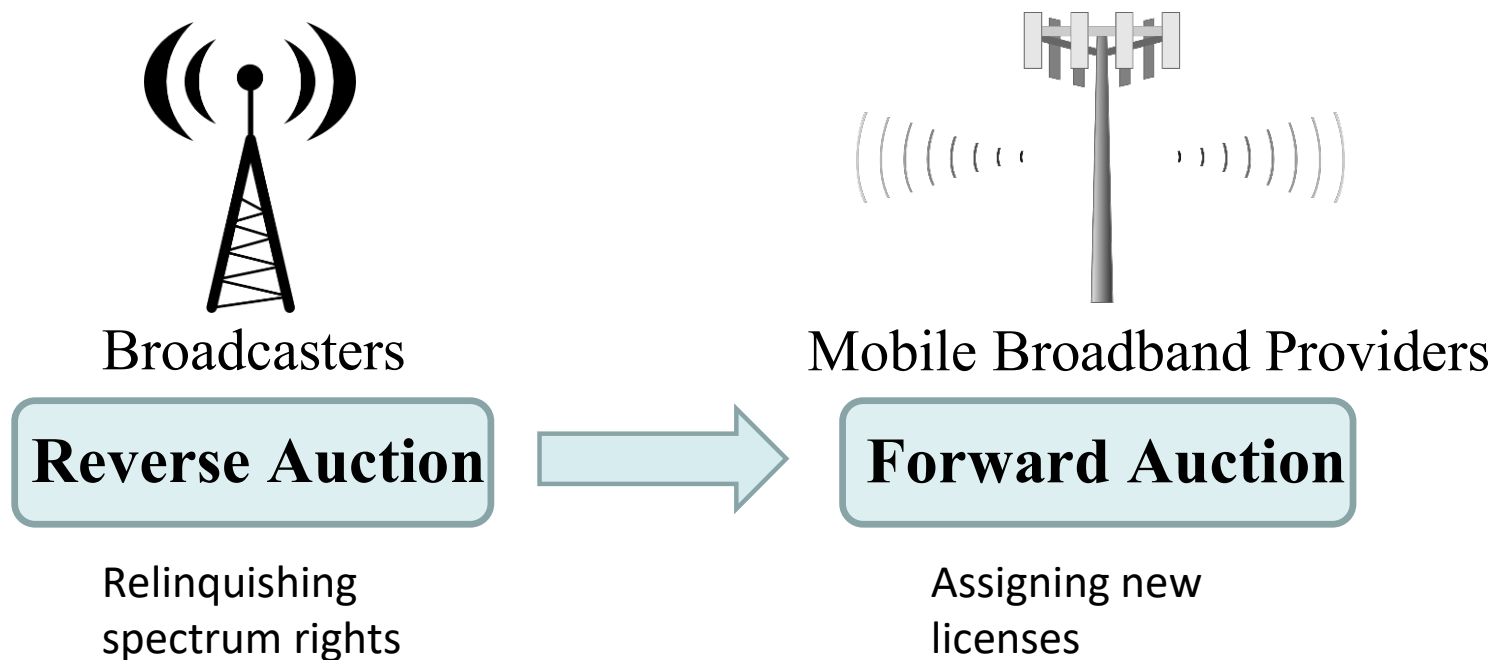
Uniform Price vs Discriminatory?

- Debate still going on for treasury auctions
- DPA is thought to raise more revenue (no formal justification though)
- UPA eliminates complaints arising from price discrimination (identical goods should cost the same!)

Examples of truthful auctions in practice

Spectrum Auctions

- Deferred Acceptance Auctions initiated by [Milgrom, Segal '14]
- Motivated by the design of the FCC “**Broadcast Incentive Auction**”



- Commenced on March 2016, closed on April 2017 for repurposing spectrum to align with consumer demand for broadband services

Basic Mechanism Design Setting

Main features:

- A **provider** of some service or resources
- A set of **single-parameter buyers** $N = \{1, 2, \dots, n\}$ interested in (some of) the resources
- Each buyer has a **valuation** v_i
- For each buyer: need to make an **accept/reject** decision
- **Feasible solutions:** Only specific subsets of buyers may be served simultaneously, due to problem constraints (e.g. interference constraints in spectrum auctions)

The framework of Deferred-Acceptance Auctions

- Backward greedy allocation algorithms
- They work in rounds, finalizing the decision for a single bidder in each round
- A_t = set of active bidders at round t
- Score of bidder i at round t : $\sigma_i^{A_t}(b_i, b_{N \setminus A_t})$
 - non-decreasing in b_i
 - Possible dependence on the set A_t (but not on the bids of active bidders)

1. Initially all bidders are **active** ($A_1=N$)
2. While accepting all active bidders in A_t is **infeasible**
 - **Reject** the bidder i with the lowest score
 - $A_{t+1} = A_t \setminus \{i\}$
3. Remaining bidders are **accepted** and pay **threshold prices**

Properties of Deferred-Acceptance Auctions

Incentive guarantees:

- Not hard to show that DA auctions are truthful
- In fact we can have much stronger incentive guarantees

Definition: A mechanism is **weakly group-strategyproof** if: for any coalition $S \subseteq N$, and any profile b_{-S} , there is no deviation by S , such that all members are strictly better off, i.e., such that:

$$u_i(b_S, b_{-S}) > u_i(v_S, b_{-S}), \text{ for every } i \in S$$

Lemma: DA auctions are weakly group-strategyproof

Properties of Deferred-Acceptance Auctions

Further advantages of DA auctions:

1. Practical and simple to implement **as long as**
 - Scoring function is simple
 - Checking feasibility of a solution is easy
2. They admit an implementation as an ascending clock auction
3. Using the ascending auction implementation:
 - Very easy to argue that truth-telling is a dominant strategy (**obvious strategyproofness [Li '15]**)
 - Privacy preservation: winners do not reveal their true value

Possible limitations:

1. They do not always guarantee a good approximation to the social welfare
2. Same for other objectives (e.g. revenue)
3. Solution returned may not be a maximal set w.r.t. problem constraints (drawback of backward greedy algorithms)



An illustration

Recall single-minded bidders from previous lectures



- The auctioneer has a set M of items for sale
- Each bidder i is interested in acquiring a specific subset of items, $S_i \subseteq M$ (known to the mechanism)
 - If the bidder does not obtain S_i (or a superset of it), his value is 0
- Each bidder submits a bid b_i for his value if he obtains the set
- Motivated by certain spectrum auctions
- Feasible allocations: the auctioneer needs to select winners who do not have overlapping sets

Single-minded bidders

Examples



- In the example above, the auctioneer can accept only 1 bidder as a winner
- In the example below, the auctioneer can accept up to 2 bidders as winners



A forward greedy algorithm for single-minded bidders

[Lehmann, O' Callaghan, Shoham '01]:


$$s_i = |S_i|$$

- Order the bidders in decreasing order of $b_i/\sqrt{s_i}$
- Accept each bidder in this order unless overlapping with previously accepted bidders



This algorithm achieves

- Monotonicity of the allocation (hence can be made truthful)
- $1/\sqrt{m}$ -approximation, where $m = |M|$
- $1/d$ -approximation, where $d = \max_i s_i$

Final conclusion: truthful polynomial time mechanism with the best possible approximation to the social welfare

Coalitions under the forward greedy mechanism

- The forward greedy mechanism is truthful but suppose players could also collude:




- What would forward greedy do?
 - Accept bid $\{c, d\}$
 - Reject bids $\{a, c\}$ and $\{b, d\}$
 - Accept bid $\{a, b\}$
 - Threshold price = 0
- The coalition $\{3, 4\}$ can change the outcome
- Threshold price still 0
- Both members better off!
- Forward greedy is not group-strategyproof

Scoring Functions for DA auctions

- Can we achieve similar welfare guarantees with backward greedy algorithms?

- How about a DA auction with scoring

 ?

- Backward greedy can do much worse than forward greedy

- Use conflict number ?

- $c_{i,t}$ = number of conflicts with other bidders at stage t

[Dutting, Gkatzelis, Roughgarden '14]:

- This does not work either
- Having s_i or $c_{i,t}$ in the denominator, raised to any power cannot achieve an $O(1/d)$ or $\tilde{O}(1/\sqrt{m})$ approximation

Positive results for DA auctions

[Dutting, Gkatzelis, Roughgarden '14]:

Theorem 1: There exists a DA auction that achieves an approximation ratio of $O(d)$

Theorem 2: There exists a DA auction that achieves an approximation ratio of $O(\sqrt{m} \log m)$

Main message:

We can have comparable approximations as in forward greedy, but with stronger incentive guarantees!

- And with a more complicated scoring function

Equilibrium analysis of non-truthful mechanisms

Non-truthful mechanisms

- As already seen, there are plenty of settings where the mechanism employed is not truthful
 - Sponsored search
 - Auctions for government bonds
 - Some types of auctions for telecom/spectrum licences (e.g., core-selecting auctions)
- Why?
 - Low revenue often achieved by truthful auctions, e.g., by VCG
 - Complexity: Social welfare maximization may turn out too difficult to solve (which is a required step in VCG-based mechanisms)
- **[Ausubel, Milgrom '06]:** The lovely but lonely Vickrey auction
 - Chapter 1 in the book “Combinatorial Auctions”

Non-truthful mechanisms

- How do we evaluate non-truthful mechanisms?
 - If the bidders are non-truthful, can we argue about the social welfare generated?
- We can think of the equilibria as the most likely outcomes to occur
 - If these games are played frequently, players may end up at an equilibrium by adjusting gradually their strategies
 - Thus, we can take the social welfare or revenue achieved at an equilibrium as an evaluation metric

PoA in auctions

- Consider an auction where v_i = actual valuation function of bidder i
 - It can be either single or multi-parameter
- Let \mathbf{b} be a pure Nash equilibrium with resulting allocation:
 $(x_1, \dots, x_n) = (x_1(\mathbf{b}), \dots, x_n(\mathbf{b}))$
- Social Welfare at \mathbf{b} : $SW(\mathbf{b}) = \sum v_i(x_i)$
- OPT = Optimal welfare (as determined by the valuations)

$$PoA = \sup_{\mathbf{b}} OPT/SW(\mathbf{b})$$

Where the supremum can be either over all pure or over all mixed equilibria

PoA in sponsored search auctions

- PoA can become unbounded in worst case
- [Lahaie '06]: $\text{PoA} \leq (\min_{1 \leq i \leq k-1} \min\{\alpha_{i+1}/\alpha_i, 1 - (\alpha_{i+2}/\alpha_{i+1})\})^{-1}$
 - For pure equilibria, when we have $k \geq 2$ slots
 - Where recall α_i is the CTR of slot i , and assume $\alpha_{k+1} = 0$
- For arbitrary auctions, the ratios of the CTRs can become arbitrarily high
- In some cases, the click data fit well with an exponential decay model (geometric CTRs): $\alpha_i \propto 1/\delta^i$ for a constant δ
 - [Feng, Bhargava, Pennock '07]: $\delta = 1.428$ using various empirical datasets
 - In these cases, $\text{PoA} \leq (\min\{1/\delta, 1-1/\delta\})^{-1}$
 - Hence, low inefficiency under geometric CTRs

PoA in sponsored search auctions

- One can also study PoA under restrictions on the set of equilibria under consideration
- E.g., some “bad” equilibria arise when some players overbid and at the same time some high-valued players underbid
- **The no-overbidding assumption:** Focus on equilibria where $b_i \leq v_i$
 - Such bidders are also referred to as conservative bidders
 - Initiated in [Christodoulou, Kovacs, Schapira '08], and assumed in several follow up works
- Can PoA be better under no-overbidding?

PoA in sponsored search auctions

- [Paes Leme, Tardos '10]: Under no-overbidding
 - $\text{PoA} \leq 1.618 (= \phi)$ for pure equilibria
 - $\text{PoA} \leq 4$ for mixed equilibria
- [Lucier, Paes Leme '11, Caragiannis et al. '11, '15]:
Currently best known:
 - $\text{PoA} \leq 1.28$ for pure equilibria
 - $\text{PoA} \leq 2.31$ for mixed equilibria
- For lower bounds, it is known that $\text{PoA} \geq 1.259$
- **Main conclusion:** For conservative bidders, selfish behavior does not lead to socially bad outcomes

PoA in multi-unit auctions

- A PoA analysis can be carried out for any other non-truthful auction
- For multi-unit auctions, PoA can be affected by the phenomenon of “demand reduction”
 - [Ausubel, Cramton '96]: Bidders may have incentives to hide their demand for items in order to achieve a better price

Example of Demand Reduction in UPA

Real profile



(1, 1, 1)



(2/3, 0, 0)



(1/2, 0, 0)

Equilibrium profile



(1, 0, 0)



(2/3, 0, 0)



(1/2, 0, 0)

$OPT = 3, SW(\mathbf{b}) = 13/6 \Rightarrow PoA \geq 18/13$ for UPA

- Revealing the true profile for bidder 1 results in a relatively high price
- Demand reduction discussed further in [Ausubel, Cramton '96]

PoA for pure equilibria

Can demand reduction create a huge loss of efficiency?

Theorem:

For the Discriminatory Price Auction (DPA), and **arbitrary** monotone valuations for the bidders, **PoA = 1**

- No need to assume no-overbidding
- All pure Nash equilibria (when they exist) are efficient
- Generalizes what holds for the single-item 1st price auction
- Existence of pure equilibria guaranteed under appropriate tie-breaking rules

PoA for pure equilibria

- The same is not true for UPA
- **Example:** Consider k units and the profiles:

Real profile



$(k, 0, 0, \dots, 0)$



$(1, 1, 1, \dots, 1)$

Equilibrium profile **b**



$(1, 1, 1, \dots, 1)$



$(0, 0, 0, \dots, 0)$

- $OPT = 2k-1$
- $SW(\mathbf{b}) = k$
- $PoA \geq (2k-1)/k = 2 - 1/k$ for UPA
- Can it get worse?

PoA for pure equilibria

- For non-conservative bidders, it can get unbounded
- The no-overbidding assumption in UPA:

$$\square \sum_{j=1}^s b_j(j) \leq v_i(s) \quad \forall i, \forall s \leq k$$

[Birmpas, Markakis, Telelis, Tsikiris '17]:

For the Uniform Price Auction (UPA), and for

- Submodular bidders
- No-overbidding pure equilibria,

$$\text{PoA} \leq 2.18$$

- Tight example even for 2 bidders

PoA for mixed equilibria

[de Keijzer, Markakis, Schaefer, Telelis '13]:

For submodular valuations, the PoA for mixed equilibria is

- $\leq e/e-1$ for DPA
- $\leq 3.146 < 2e/e-1$ for UPA

Remarks:

- $3.146\dots = |W_{-1}(-1/e^2)|$ (Lambert W function)
- Bounds hold both for standard bidding and for the simplified uniform bidding format
- The same bounds also hold for Bayesian games (PoA for Bayes-Nash equilibria)

PoA for mixed equilibria

- Currently known lower bounds: ≈ 1.1 for DPA, **2.18** for UPA
 - Far from tight in the case of mixed equilibria
- Our proof can be cast into the smoothness framework of **[Syrngkanis, Tardos '13]**



- Upper bounds carry over to simultaneous and sequential compositions of multi-unit auctions (e.g. combinatorial multi-unit auctions)
- Similar approaches and techniques used in other types of auctions as well (e.g. item-bidding auctions)
[Christodoulou, Kovacs, Schapira '08, Bhawalkar, Roughgarden '11, Feldman, Fu, Gravin, Lucier '13]

Conclusions on PoA

- **Take-home story:** simple auction formats used in practice perform quite well w.r.t. social welfare
- Upper bounds:
 - For pure equilibria, almost tight for sponsored search, completely tight for multi-unit auctions
 - Open if we can improve the bounds for mixed equilibria
 - PoA can also become even better if we focus on Nash equilibria in undominated strategies
- Lower bounds:
 - Much harder to get

Final conclusions

- A wide range of applications
- The full spectrum of incentive guarantees can be seen in practice
 - **Non-truthful and bad equilibria** (uniform price auction or sponsored search with overbidding)
 - **Non-truthful and efficient equilibria** (single-item first price auction)
 - **Non-truthful and relatively efficient equilibria** (sponsored search, uniform price auction, under no-overbidding)
 - **Truthful** (single-item Vickrey)
 - **Weakly group-strategyproof** (DA auctions)
- The choice of mechanism deployed may depend on:
 - Traditions and practices used in a specific application domain (not always easy to switch to a new format)
 - Complexity considerations (simplicity is often a must)
 - Legal issues (there exist governmental auctions where social welfare w.r.t. reported bids needs to be maximized)