Algorithmic Game Theory

Introduction to Mechanism Design for Single Parameter Environments

Vangelis Markakis
markakis@gmail.com

Mechanism Design

- What is mechanism design?
- It can be seen as reverse game theory
- Main goal: design the rules of a game so as to
  - avoid strategic behavior by the players
  - and more generally, to enforce a certain behavior for the players or other desirable properties
- Applied to problems where a “social choice” needs to be made
  - i.e., an aggregation of individual preferences to a single joint decision
- strategic behavior = declaring false preferences in order to gain a higher utility

Examples

- Elections
  - Parliamentary elections, committee elections, council elections, etc
  - A set of voters
  - A set of candidates
  - Each voter expresses preferences according to the election rules
    - E.g., by specifying his single top choice, or by specifying his first few choices, or by submitting a full ranking of the candidates
  - Social choice: can be a single candidate (single-winner election) or a set of candidates (multi-winner election) or a ranking of the candidates

Examples

- Auctions
  - An auctioneer with some items for sale
  - A set of bidders express preferences (offers) over items
    - Or combinations of items
  - Preferences are submitted either through a valuation function, or according to some bidding language
  - Social choice: allocation of items to the bidders

Examples

- Government policy making and referenda
  - A municipality is considering implementing a public project
  - Q1: Should we build a new road, a library or a tennis court?
  - Q2: If we build a library where shall we build it?
  - Citizens can express their preferences in an online survey or a referendum
  - Social choice: the decision of the municipality on what and where to implement

Specifying preferences

- In all the examples, the players need to submit their preferences in some form
- Representation of preferences can be done by
  - A valuation function (specifying a value for each possible outcome)
  - A ranking (an ordering on possible outcomes)
  - An approval set (which outcomes are approved)
- Possible conflict between increased expressiveness vs complexity of decision problem

Single-item Auctions

Auctions

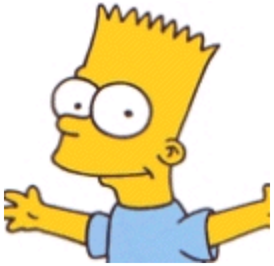

Set of players
N = {1, 2, …, n}

1 indivisible good

Auctions

- A means of conducting transactions since antiquity
  - First references of auctions date back to ancient Athens and Babylon
- Modern applications:
  - Art works
  - Stamps
  - Flowers (Netherlands)
  - Spectrum licences
  - Other govermental licences
  - Pollution rights
  - Google ads
  - eBay
  - Bonds
  - ...

Auctions

- Earlier, the most popular types of auctions were
  - The English auction
    - The price keeps increasing in small increments
    - Gradually bidders drop out till there is only one winner left
  - The Dutch auction
    - The price starts at $+\infty$ (i.e., at some very high price) and keeps decreasing
    - Until there exists someone willing to offer the current price
  - There exist also many variants regarding their practical implementation
- These correspond to ascending or descending price trajectories

Sealed bid auctions

- Sealed bid: We think of every bidder submitting his bid in an envelope, without other players seeing it
  - It does not really have to be an envelope, bids can be submitted electronically
  - The main assumption is that it is submitted in a way that other bidders cannot see it
- After collecting the bids, the auctioneer needs to decide:
  - Who wins the item?
    - Easy! Should be the guy with the highest bid
      - Yes in most cases, but not always
  - How much should the winner pay?
    - Not so clear

Sealed bid auctions

Why do we view auctions as games?

- We assume every player has a valuation v_i for obtaining the good
- Available strategies: each bidder is asked to submit a bid b_i
  - $b_i \in [0, \infty)$
  - Infinite number of strategies
- The submitted bid b_i may differ from the real value v_i of bidder i

First price auction

Auction rules

- Let $\mathbf{b} = (b_1, b_2, \ldots, b_n)$ the vector of all the offers
- Winner: The bidder with the highest offer
  - In case of ties: We assume the winner is the bidder with the lowest index (not important for the analysis)
  - E.g. if there is a tie among bidder 2 and bidder 4, the winner is bidder 2
- Winner's payment: the bid declared by the winner
- Utility function of bidder i,

$$u_i(\mathbf{b}) = \begin{cases} v_i - b_i\,, & \text{if i is the winner} \\ 0, & \text{otherwise} \end{cases}$$

Incentives in the first price auction

Analysis of first price auctions

•There are too many Nash equilibria

•Can we predict bidding behavior? Is some equilibrium more likely to occur?

•Hard to tell what exactly will happen in practice but we can still make some conclusions for first price auctions

Observation: Suppose that $v_1 \geq v_2 \geq v_3 \ldots \geq v_n$. Then the profile $(v_2, v_2, v_3, \ldots, v_n)$ is a Nash equilibrium

Corollary: The first price auction provides incentives to bidders to hide their true value

•This is highly undesirable when $v_1 - v_2$ is large

Auction mechanisms

We would like to explore alternative payment rules with better properties

Definition: For the single-item setting, an auction mechanism receives as input the bidding vector **b** = $(b_1, b_2, \ldots, b_n)$ and consists of

- an allocation algorithm (who wins the item)
- a payment algorithm (how much does the winner pay)

Most mechanisms satisfy individual rationality:

- Non-winners do not pay anything
- If the winner is bidder i, her payment will not exceed b_i (it is guaranteed that no-one will pay more than what she declared)

Auction mechanisms

Aligning Incentives

•Ideally, we would like mechanisms that do not provide incentives for strategic behavior

•How do we even define this mathematically?

An attempt:

Definition: A mechanism is called truthful (or strategyproof, or incentive compatible) if for every bidder i, and for every profile $\mathbf{b}_{-i}$ of the other bidders, it is a dominant strategy for i to declare her real value v_i, i.e., it holds that

$$u_i(v_i, \mathbf{b}_{-i}) \geq u_i(b', \mathbf{b}_{-i}) \text{ for every } b' \neq v_i$$

Auction mechanisms

- In a truthful mechanism, every rational agent knows what to play, independently of what the other bidders are doing
- It is a win-win situation:
  - The auctioneer knows that players should not strategize
  - The bidders also know that they should not spend time on trying to find a different strategy
- Very powerful property for a mechanism
- Fact: The first-price mechanism is not truthful

Are there truthful mechanisms?

The 2nd price mechanism (Vickrey auction)

[Vickrey '61]

•Allocation algorithm: same as before, the bidder with the highest offer

- In case of ties: we assume the winner is the bidder with the lowest index

•Payment algorithm: the winner pays the 2nd highest bid

•Hence, the auctioneer offers a discount to the winner

Observation: the payment does not depend on the winner's bid!

- The bid of each player determines if he wins or not, but not what he will pay

The 2nd price mechanism (Vickrey auction)

[Vickrey '61] (Nobel prize in economics, 1996)

- **Theorem:** The 2nd price auction is a truthful mechanism

Proof sketch:

- Fix a bidder i, and let $\mathbf{b}_{-i}$ be an arbitrary bidding profile for the rest of the players
- Let $b^* = \max_{j \neq i} b_j$
- Consider now all possible cases for the final utility of bidder i, if he plays v_i
  - $v_i < b^*$
  - $v_i > b^*$
  - $v_i = b^*$
  - In all these different cases, we can prove that bidder i does not become better off by deviating to another strategy

Optimization objectives

What do we want to optimize in an auction?

Usual objectives:

- Social welfare (the total welfare produced for the involved entities)
- Revenue (the payment received by the auctioneer)

We will focus first on social welfare

Optimization objectives

What do we want to optimize in an auction?

Definition: The utilitarian social welfare produced by a bidding vector **b** is SW(**b**) = $\Sigma_i\ u_i(\mathbf{b})$

•The summation includes the auctioneer's utility (= the auctioneer's payment)

•The auctioneer's payment cancels out with the winner's payment

➢For the single-item setting, SW(**b**) = the value of the winner for the item

➢An auction is welfare maximizing if it produces an allocation with optimal social welfare when bidders are truthful

Vickrey auction: an ideal auction format

Summing up:

Theorem: The 2^{nd} price auction is

- truthful [incentive guarantees]
- welfare maximizing [economic performance guarantees]
- implementable in polynomial time [computational performance guarantees]

Even though the valuations are private information to the bidders, the Vickrey auction solves the welfare maximization problem as if the valuations were known

Generalizations to single-parameter environments

Single-parameter mechanisms

- In many cases, we do not have a single item to sell, but multiple items
- But still, the valuation of a bidder could be determined by a single number (e.g., value per unit)
- Note: the valuation function may depend on various other parameters, but we assume only a single parameter is private information to the bidder
  - The other parameters may be publicly known information
- We can treat all these settings in a unified manner
- Our focus: Direct revelation mechanisms
  - The mechanism asks each bidder to submit the parameter that completely determines her valuation function

Examples of single-parameter environments

- Single-item auctions:
  - One item for sale
  - each bidder is asked to submit his value for acquiring the item
- k-item unit-demand auctions
  - k identical items for sale
  - each bidder submits his value per unit and can win at most one unit
- Knapsack auctions
  - k identical items, each bidder has a value for obtaining a certain number of units
- Single-minded auctions
  - a set of (non-identical) items for sale
  - each bidder is interested in acquiring a specific subset of items (known to the mechanism)
  - Each bidder submits his value for the set she desires

Examples of single-parameter environments

- Sponsored search auctions
  - multiple advertising slots available, arranged from top to bottom
  - each bidder interested in acquiring as high a slot as possible
  - each bidder submits his value per click
- Public project mechanisms
  - deciding whether to build a public project (e.g., a park)
  - each bidder submits his value for having the project built

In all these settings, we can have multiple winners in the auction

Some Notation

- Suppose we have n players
- Let v_i be the parameter that is private information to player i
  - Usually v_i corresponds to value per unit, or in general maximum willingness to pay per unit received
  - Or v_i can be the value derived by the bidder when she is a winner (e.g., in public project problems)

General form of direct-revelation mechanisms for single-parameter problems:

- Input: The bidding vector $\mathbf{b} = (b_1, \ldots, b_n)$ by the players
  - each b_i may differ from v_i
- Allocation rule: Choose an allocation $\mathbf{x}(\mathbf{b}) = (x_1(\mathbf{b}), x_2(\mathbf{b}), \ldots, x_n(\mathbf{b}))$
  - $x_i(\mathbf{b})$ = number of units received by pl. i or generally the allocation to i
- Payment rule: $\mathbf{p}(\mathbf{b}) = (p_1(\mathbf{b}), p_2(\mathbf{b}), \ldots, p_n(\mathbf{b}))$
  - $p_i(\mathbf{b})$ = payment for bidder i

Some Notation

- We will use (**x**, **p**) to refer to a mechanism with allocation function **x**, and payment function **p**
- Final utility of bidder i in a mechanism M = (**x**, **p**):
  - $u_i(\mathbf{b}) = v_i\, x_i(\mathbf{b}) - p_i(\mathbf{b})$
  - Quasi-linear form of utility functions
- For simplicity, we often write $(x_1, x_2,\ldots, x_n)$ instead of $(x_1(\mathbf{b}), x_2(\mathbf{b}),\ldots, x_n(\mathbf{b}))$
- We focus on mechanisms that satisfy Individual Rationality:
  - If a bidder i is a non-winner ($x_i(\mathbf{b}) = 0$), then $p_i(\mathbf{b}) = 0$
  - For winners, the payment rule satisfies $p_i(\mathbf{b}) \in [0, b_i\, x_i(\mathbf{b})]$ for every bidding vector **b** and every i
  - The auctioneer can never ask a bidder for a payment higher than her declared total value for what she won

Examples of single-parameter environments

Describing the feasible allocations

- Single-item auctions:
  - $x_i \in \{0, 1\}$ for every i, and $\Sigma_i x_i = 1$
- k-item unit-demand auctions
  - k identical items for sale
  - $x_i \in \{0, 1\}$, $\Sigma_i x_i <= k$
- Knapsack auctions
  - k identical items for sale
  - For each bidder, demand w_i
  - $x_i \in \{0, 1\}$ for every i, $\Sigma_i w_i x_i <= k$
- Public project mechanisms
  - Deciding whether to build a public project (e.g., a park)
  - Only 2 feasible allocations: (0, 0, ..., 0) or (1, 1, ..., 1)

Allocation rules and truthful mechanisms

- Can we understand how to derive truthful mechanisms?
- Actually, we can rephrase this as:
  - Suppose we are given an allocation rule **x**
  - Can we tell if **x** can be combined with a pricing rule **p**, so that (**x**, **p**) is a truthful mechanism?
- This would allow us to focus only on designing the allocation algorithm appropriately
- Consider the single-item auction
  - Allocation rule 1: Give the item to the highest bidder
  - Allocation rule 2: Give the item to the 2nd highest bidder
- For rule 1, we have seen how to turn it into a truthful mechanism (Vickrey auction)
- For rule 2?
  - We have not seen how to do this, but we have also not proved that it cannot be done

Allocation rules and truthful mechanisms

- Consider a mechanism with allocation rule **x**
- Fix a player i, and fix a profile $\mathbf{b}_{-i}$ for the other players
- Allocation to player i at a profile $\mathbf{b} = (z, \mathbf{b}_{-i})$ is given by $x_i(\mathbf{b})$
- Keeping $\mathbf{b}_{-i}$ fixed, we can view the allocation to player i as a function of his bid
  - $x_i = x_i(z, \mathbf{b}_{-i})$, if bidder i bids z

- Definition: An allocation rule is monotone if for every bidder 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

Monotonicity of allocation rules

Examples

- Back to the single-item auction
- The allocation rule that gives the item to the highest bidder is monotone
  - If a bidder wins at profile **b**, she continues to be a winner if she raises her own bid (keeping $\mathbf{b}_{-i}$ fixed)
  - If she was not a winner at **b**, then by raising her bid, she will either remain a non-winner or she will become a winner
- The allocation rule that gives the item to the 2^{nd} highest bidder is not monotone
  - If I am a winner and raise my bid, I may become the highest bidder and will stop being a winner

Myerson’s lemma

[Myerson ’81]

- Theorem: For every single-parameter environment,
  - An allocation rule **x** can be turned into a truthful mechanism if and only if it is monotone
  - If **x** is monotone, then there is a unique payment rule **p**, so that (**x**, **p**) is a truthful mechanism
    - Subject to the constraint that if $b_i = 0$, then $p_i = 0$

- One of the classic results in mechanism design
- In fact, in many cases we can also compute the payments by a simple formula

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 $b^* = \max_{j \neq i} b_j$

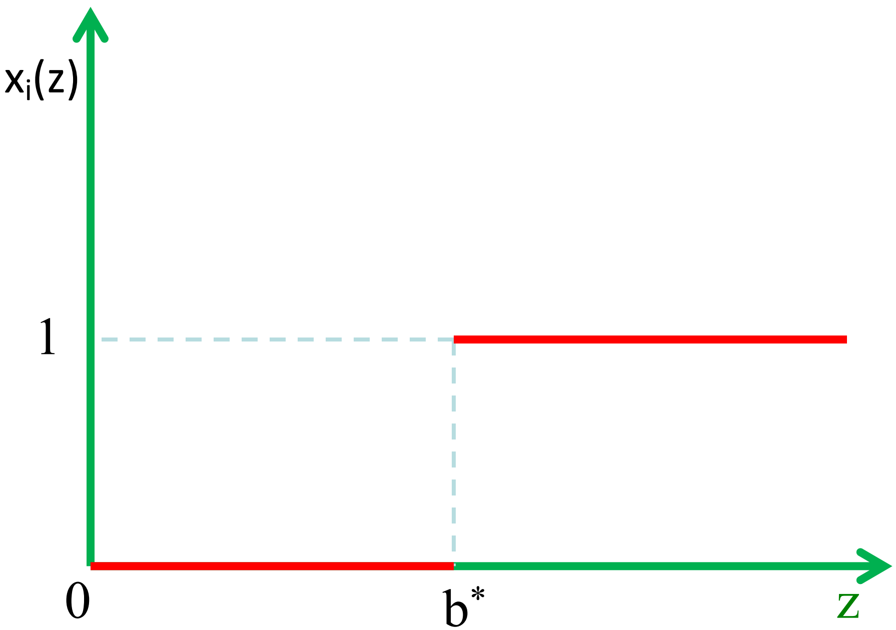

Facts:

- 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

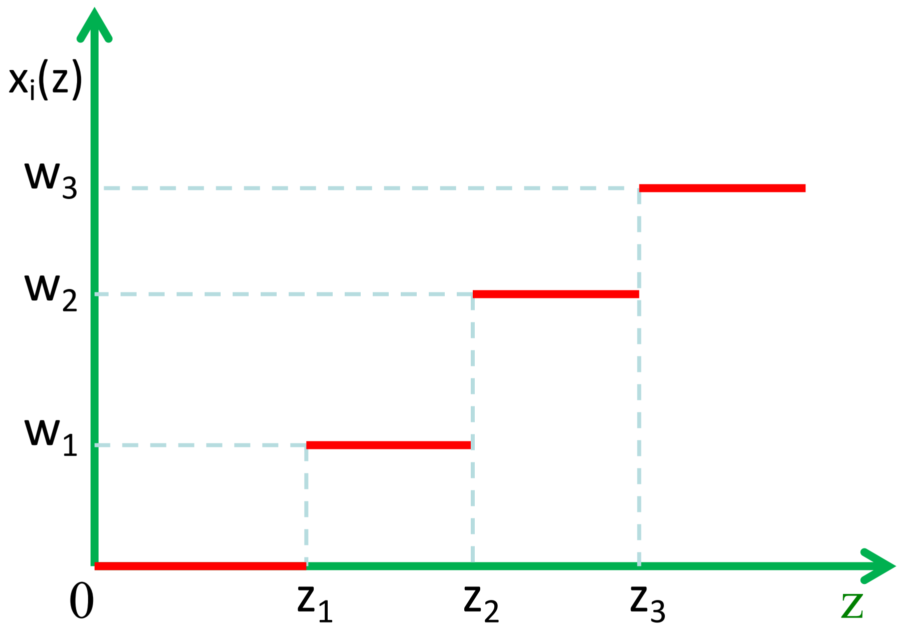

- Suppose bidder i bids b_i
- Look at the jumps of $x_i(z, b_{-i})$ in the interval $[0, b_i]$
- Suppose we have k jumps
- Jump at $z_1 = w_1$
- Jump at $z_2 = w_2 - w_1$
- Jump at $z_3 = w_3 - w_2$
- ...
- Jump at $z_k = w_k - w_{k-1}$

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

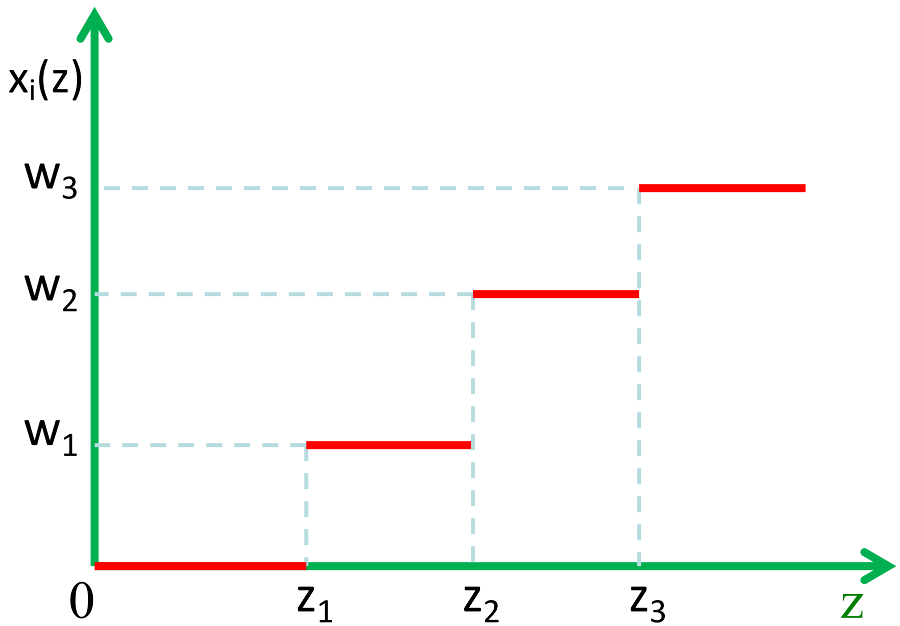

Payment formula

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

Applying Myerson's lemma

- Single-item auctions
- The allocation rule of giving the item to the highest bidder is monotone
- The payment rule of Myerson gives us precisely the Vickrey auction
  - Non-winners pay nothing: If a bidder i is not a winner, there is no jump within $[0, b_i]$ in the function $x_i(z, \mathbf{b}_{-i})$
  - The winner pays (2^{nd} highest bid) · [jump at 2^{nd} highest bid] = 2^{nd} highest bid

- Corollary: The Vickrey auction is the only truthful mechanism for single-item auctions, when the winner is the highest bidder

Sponsored Search Auctions

What is sponsored search?

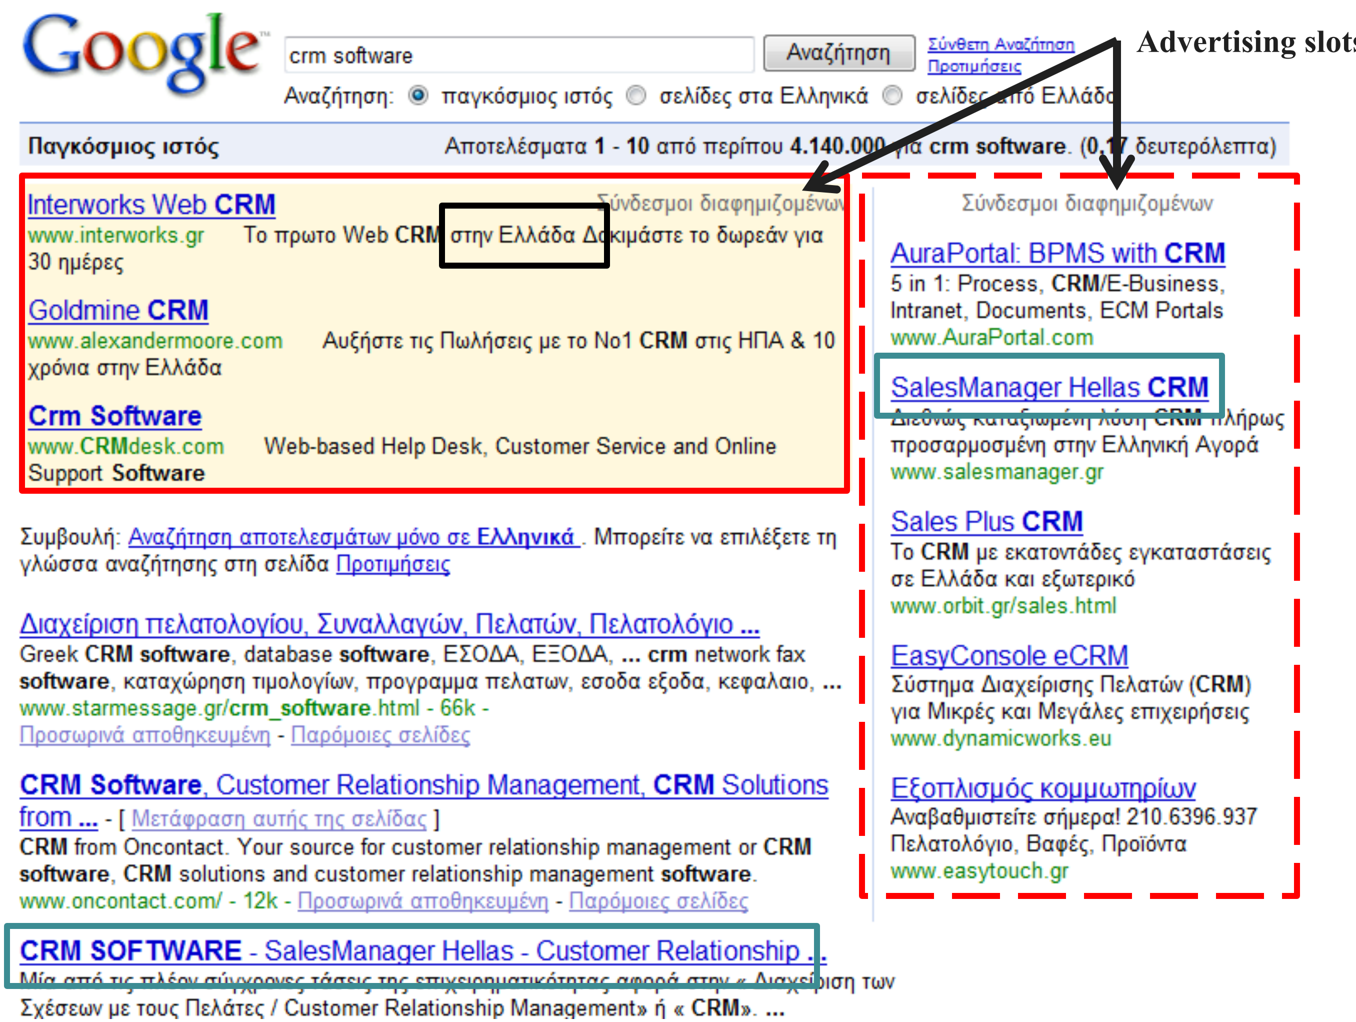

What is sponsored search?

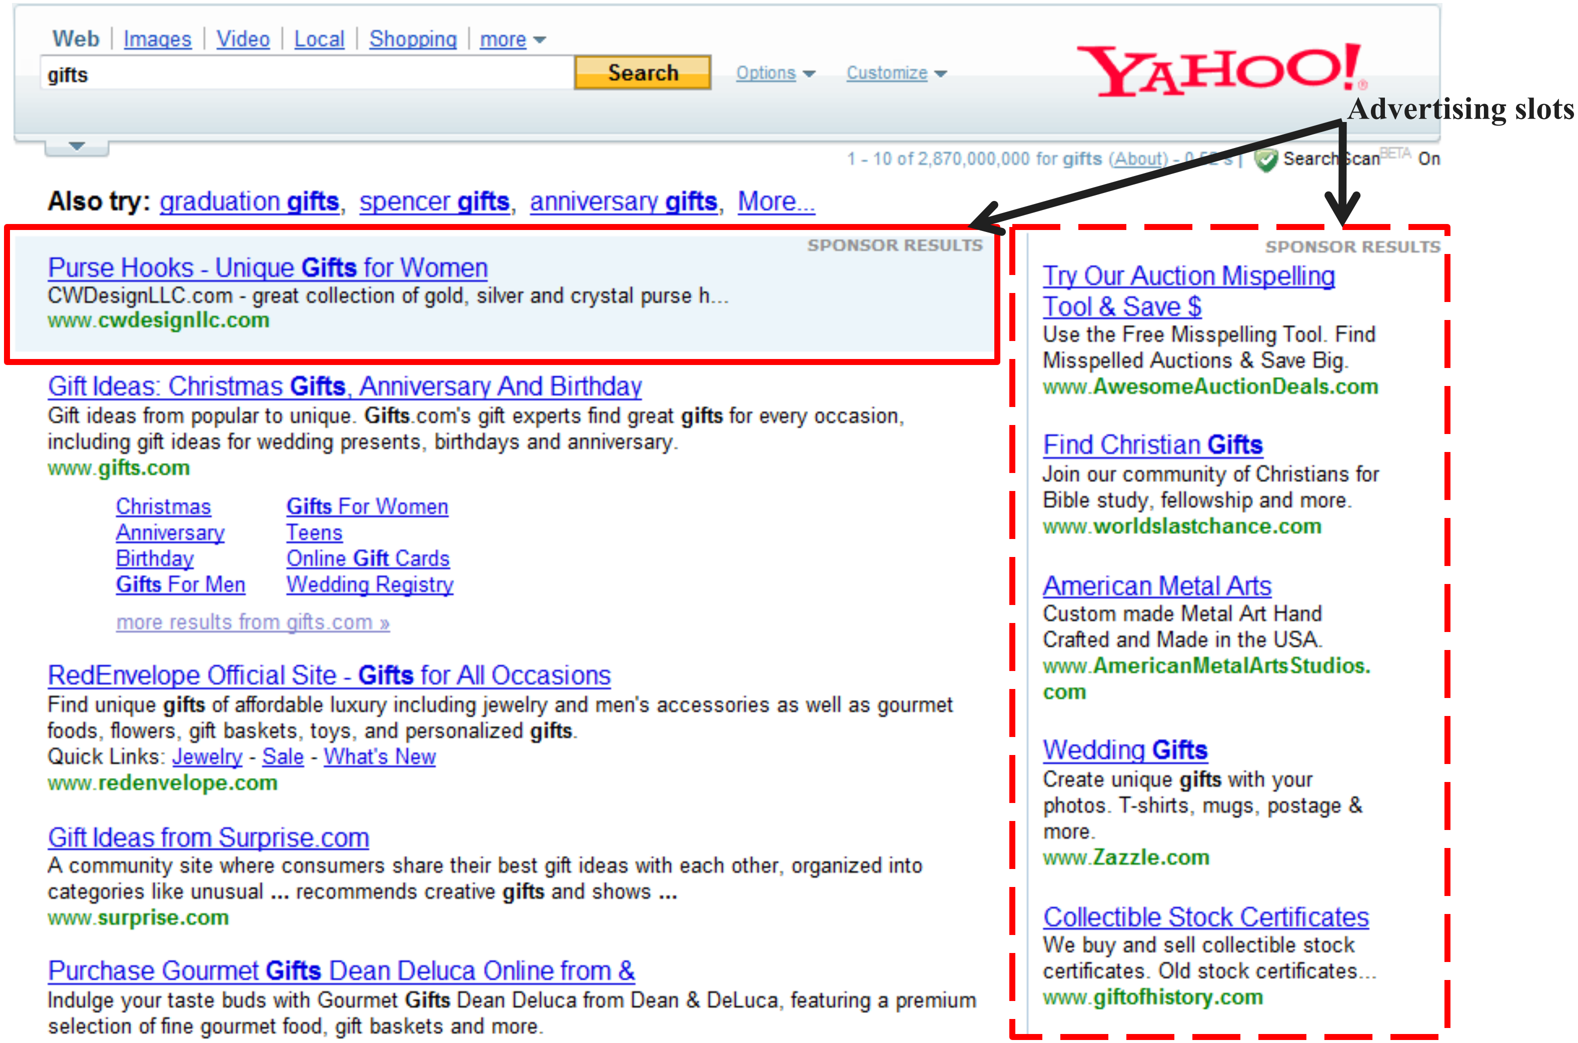

How does it work?

- For a fixed search term (e.g. *ipod*)
  - n advertisers
  - k slots (typically k << 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 click-through rates and statistics show that

$$\alpha_1 \geq \alpha_2 \geq \alpha_3 \geq \ldots \geq \alpha_k$$

  - 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, $x_i(b_i, b_{-i}) \in \{0, \alpha_k, \alpha_{k-1}, \dots, \alpha_1\}$

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

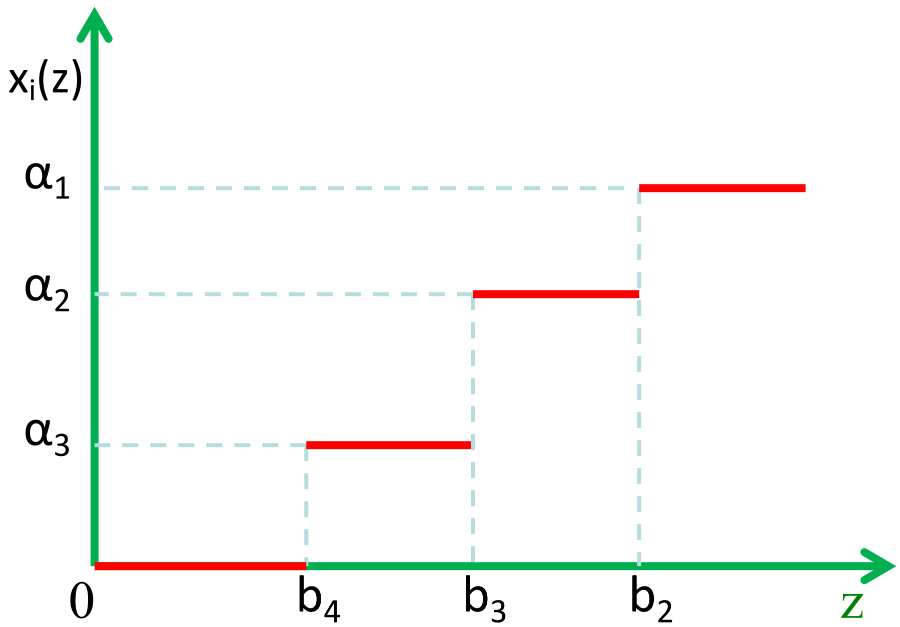

- Look at the jumps of x_i in the interval $[0, b_1]$
- Jump at $b_4 = \alpha_3$
- Jump at $b_3 = \alpha_2 - \alpha_3$
- Jump at $b_2 = \alpha_1 - \alpha_2$

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

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

- This would have been the payment if bidders cared for impressions and not for clicks
- 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: $b_1 \geq b_2 \geq \ldots \geq b_n$
  - 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}

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 $q_i \times b_i$
    $q_1 \times b_1 \geq q_2 \times b_2 \geq \ldots \geq q_n \times b_n$
  - 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 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
- GSP was employed probably by accident
  - As an attempt to 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
  - Theoretical analysis: the Nash equilibria of GSP have revenue at least as high as the revenue of truthful bidding
  - Further connections also exist between GSP outcomes and the outcome of the truthful mechanism