CS 573: Algorithmic Game Theory	Lecture date: April 9, 2008
Instructor: Chandra Chekuri	Scribe: Shuyi Chen

Contents

1 Digital Goods Auction – *follow-up*

Definition 1.1 *(Digital Goods auction) In a digital goods auction, there are,*

- *n bidders with valuations $v_1, v_2, ..., v_n$ and bids $b_1, b_2, ...b_n$.*
- *n units of an item for sale.*

Each bidder is interested in one unit of the item.

If VCG is used to to allocate items to all players, all players will be charged 0 and the revenue in VCG is 0.
In Bayesian setting, if all players share a common distribution function F and it is known, then the optimal revenue maximizing auction is to set price $p^* = argmax_p p(1 - F(p))$. In asymmetric setting where each bidder has a different distribution function F_i for his valuation, we can use Myerson's optimal mechanism to get the optimal auction.
In non-Bayesian setting, we can employ the *empirical* Myerson mechanism as follow.

Definition 1.2 *(Deterministic Optimal Price auction)DOP for digital goods,*

- *Let $b = (b_1, ..., b_n)$ be the bids of the n players, and $b_{-i} = (b_1, ..., b_{i-1}, b_{i+1}, ..., b_n)$.*
- *For each i, compute the best single price p_i that maximizes profit for b_{-i},*
  $p_i = argmax_p \; p|\{j | j \neq i, b_j \geq p\}|$

.

Claim 1.3 *If all bidders' valuations are i.i.d. according to a common distribution function F with support from a bounded interval, then as $n \to \infty$, the expected revenue of DOP approximates that of the optimal (Myerson) auction.*

Given the DOP auction, we want to know how well it does in revenue-maximizing compared to the optimal revenue $OPT = max_p \; p|\{j | v_j \geq p\}|$ associated with a single price if valuations of players are known.

Theorem 1.4 *No deterministic anonymous auction for bidders with valuation in* $[1, h]$ *can achieve revenue* $\frac{OPT}{\alpha} - \gamma h, \forall \alpha, \gamma$

The above theorem shows that no constant competitive deterministic auction exists. However, it turns out that randomness helps here. We introduce a random sampling auction as follow.

Definition 1.5 *(Random Sampling Optimal Price auction)In RSOP,*

- *we partition bidders into two sets,* N_1 *and* N_2 *with each bidder placed in either set with equal probability(0.5).*
- *find the best price* p_1 *for bids in* N_1 *and* p_2 *for bids in* N_2*.*
- *offer items at price* p_1 *to bidders in* N_2 *and price* p_2 *to bidders in* N_1*.*

Theorem 1.6 *For bids in* $[1, h]$*, the expected profit of RSOP approximates* OPT *as* $n \to \infty$*, or,* $\forall v_1, ..., v_n$*,* $E[RSOP] \geq (1 - \epsilon_n) OPT(v_1, ..., v_n) - h$*, where* $\epsilon_n \to \infty$ *as* $n \to \infty$*.*

2 Competitive Auctions for Digital Goods

Recall that when the bidders were drawn from a known distribution there was a very natural notion of what it meant to be the optimal auction. For a given distribution there exists a truthful auction that obtains the highest expected profit. This auction is optimal! The same is not true in worst case settings.

Claim 2.1 *There does not exist an auction that is best on every input.*

This impossibility leaves us with the question of how we can arrive at a rigorous theoretical framework in which we can determine some auction to be optimal. The key to resolving this issue is in moving from absolute optimality to relative optimality. The general approach will be to try to design an auction that is always (in worst case) within a small constant factor of some *profit measure.*

Definition 2.2 *A profit measure is a function* $\phi : \Re^n \to \Re$ *which maps a vector of valuations to a target profit. In terms of digital goods auction, the function is* $\phi : [1, h]^n \to \Re_+$

For digital goods auction, we define the following *profit measure.*

- $(\mathcal{T})$ *"Sell to all bidders at their valuation."*
  $Profit = \mathcal{T}(\mathbf{v}) = \sum_i v_i$.
- $(\mathcal{F})$ *"Sell at optimal single sale price."*
  $Profit = \mathcal{F}(\mathbf{v}) = max_p \, pn_p, n_p = |\{i | b_i \geq p\}|$.
- $(\mathcal{F}^2)$ The *optimal single priced profit with at least two winners* is $\mathcal{F}^2 = max_{p, n_p \geq 2} \, pn_p, n_p = |\{i | b_i \geq p\}|$.

Definition 2.3 *An auction is* α*-competitive w.r.t. a profit measure* P *if* $\forall \bar{v} \in [1, h]^n$*, the expected revenue of the auction* $E[\bar{v}] \geq \frac{P(\bar{v})}{\alpha}$*.*

Theorem 2.4 *$\forall n$ and $\bar{v} \in [1, h]^n$, no (deterministic or randomized) auction is $\mathcal{O}(logh)$-competitive w.r.t. $\mathcal{T}(\bar{v})$.*

Proof: Exercise. □

Corollary 2.5 *For $n = 2$, no auction is $\mathcal{O}(logh)$-competitive w.r.t. $\mathcal{F}$.*

Theorem 2.6 *$\forall n$ and $\bar{v} \in [1, h]^n$, no deterministic anonymous auction can achieve $\mathcal{O}(lgh)$-competitive w.r.t. $\mathcal{F}^2$.*

Theorem 2.7 *RSOP is 15-competitive w.r.t. $\mathcal{F}^2$.*

In fact, RSOP can be derandomized to yield a deterministic (non-anonymous) auction that is constant competitive *w.r.t.* $\mathcal{F}^2$. Also,one can also show that,

Theorem 2.8 *No auction, even with randomization, can achieve better than 2.42 competitive w.r.t. $\mathcal{F}^2$.*

3 Sponsored Search Auction

Web search engines like Google and Yahoo! monetize their service by auctioning off advertising space next to their standard algorithmic search results. Typically, the sponsored results are displayed in a format similar to algorithmic results: as a list of items each containing a title, a text description, and a hyperlink to the advertiser's Webpage. We call each position in the list a *slot*. Generally, advertisements that appear in a higher ranked slot (higher on the page) gain more attention and more clicks to lower ranked slot. Generally speaking, the advertisers bid on keywords that are relevant to them. Suppose the user only enters a single keyword each time, then, when the user enters the keyword §, all advertisers that bid on § are taken into the auction and the winners will be assigned a *slot* on the result page. Finally, the search engine will get paid only if a user clicks on a *slot*.
In sponsored search auctions, there are several practical issues.

- The same auction repeats every day for the same keyword, so it is a repeated game with many players. Further, each advertiser bids on hundreds of keywords and it's difficult to take all into account in a model.
- In addition, bidders have budget constraits on their advertising.
- What is more, information is really not symmetric, since, usually, the auctioneer(search engine) knows more about the information, i.e. the user click behavior and the click statistic for each keyword.

For the above reason, it is only an approximation to treat the auction for a single keyword as an isolated static auction. We give the definition for an isolated static auction with a single keyword as follow.

Definition 3.1 *A static auction for a single keyword consists of*

- *n bidder(advertisers) who are interested in a single keyword.*

- *k slots (k is assumed to be fixed and $k < n$).*
- *for each bidder i and slot j, there is a click through rate (CTR) α_{ij} which is the probability that slot j is clicked when player i is in it.*

Here are the assumptions we make.

Assumption 3.2
- *CTRs are known and CTRs do not depend on identity of the bidder.*
- *Bidder expected payoff is the same for clicks at any slot*
- $\alpha_{ij} \geq \alpha_{ij+1}, \forall i, j < k$

Under these assumptions, we can think of the problem as a single-parameter problem. Each bidder i has a private value v_i for a click irrespective of where the click came from.