Online Competitive auctions

Elias Koutsoupias

Joint work with George Pierrakos

Outline

1. The model
2. The state of the art
3. The online question
4. On stochastic input and randomized algorithms

Digital goods auctions

- We want to sell a digital good (with no replication cost)
- There are n bidders who have a **private valuation** for the good
- Objective: Design an auction to maximize the profit

- Offline
  All bidders are present
- Online
  Bidders appear online

How to model uncertainty?

- Adversarial:
  The input is designed by a powerful adversary who knows the algorithm and tailors the set of bids to defeat it
- Stochastic:
  There is a known or unknown probability distribution.
  - Independent bids: Each bid is selected independently from the distribution
  - Correlated bids: The probability distribution is for sets of bids and not for each bid separately
- Random-order (online)

  The adversary selects the set of bids and they are presented in a random order, as in the **secretary problem**

Some truthful offline auctions

An auction is **truthful** if and only if the price offered to a bidder is independent of his bid

- DOP (offline)
  - To every bidder offer the optimal single price for the remaining bidders
- RSOP (offline)
  - Partition the randomly bidders into two sets
  - Find the optimal single price for each set and offer it to the bidders of the other set
- BPSF (online)
  - To every bidder offer the optimal single price for the revealed bids (the online version of DOP)

How to evaluate an auctios?

- Let $b_1 > b_2 > \ldots > b_n$ be the bids
  Compare a mechanism against
- Sum of all bids: $\sum_i b_i$ (unrealistic)
- Optimal single-price profit: $\max_i i{*}b_i$ (problem: highest bid impossible to get)
- A reasonable benchmark: $F^{(2)} = \max_{i>=2} i{*}b_i$
  The **optimal** profit of
  - a single-price auction
  - which sells the good to at least 2 bidders

  This is the benchmark we adopt
- We call an algorithm ρ-**competitive** if its profit is at least $F^{(2)}/\rho$

Questions for benchmark $F^{(2)}$

- Optimal competitive ratio for the **adversarial offline** case?
  - Symmetric deterministic: unbounded
  - Randomized: $\in [2.42, 3.24]$
  - RSOP is 4.64 competitive
  - Conjecture: RSOP is 4-competitive

(Goldberg-Hartline-Karlin-Wright-Saks, Hartline-McGrew)

Question for benchmark $F^{(2)}$

- Optimal competitive ratio for the stochastic case?
  - Again $\in [2.42, 3.24]$
  - Why the same? Because of Yao's lemma
  - Theorem: For bid-independent distributions the answer is 2.42
- Optimal online competitive ratio for the random-order case?
  - Theorem: There is a generic transformation of offline auctions to online auctions, with only a loss of a factor of 2 in the competitive ratio.
  - Competitive ratio $\in [4, 6.48]$
  - Conjecture: The BPSF auction is 4-competitive

(Previous work: Majiaghayi-Kleinberg-Parkes, in 2004 showed a very high competitive ratio)

The online question

$$b_{\pi_1}, \ldots, b_{\pi_{t-1}} \rightarrow b_{\pi_t}$$

- π is a random permutation
- What is the best price to offer to b_{π_t}?
- We assume that the past bids are known
- A learning question?

The online setting

$$b_{\pi_1}, \ldots, b_{\pi_{t-1}} \rightarrow b_{\pi_t}$$

- Min, Mean, Median: unbounded competitive ratio
- Max: competitive ratio approx. $k/(H_k - 1)$, where $F^(2) = kb_k$. No bad for small values of k (less than 4 for k<=5)
- SCS is a variant of RSOP with offline competitive 4. Its online version has competitive ratio less than 4 for k>=5

Transforming an offline mechanism to online

$$b_{\pi_1}, \ldots, b_{\pi_{t-1}} \rightarrow b_{\pi_t}$$

- Simply run the offline algorithm for the set of revealed bids and the current (unrevealed bid)
- Is it good? We compare we $\mathrm{F}^{(2)}$ of **all** bids
- Theorem:
  We loose a factor of 2 at most. In fact, only k/(k-1) where $F^(2) = kb_k$.

Proof

- Let ρ be the offline competitive ratio
- Expected online profit at step $t = 1/t * 1/\rho$ * expected offline profit of the first t bids
- with probability $\binom{t}{m}\binom{n-t}{k-m}/\binom{n}{k}$ the first t bids have m of the highest k bids which contribute to the optimum.
- offline profit >= m b_k, **when m>=2**
- Putting everything together

$$\text{online profit} >= \sum_{t=2}^{n} \sum_{m=2}^{\min\{t,k\}} \frac{\binom{t}{m}\binom{n-t}{k-m}}{\binom{n}{k}} \frac{1}{t\rho} m b_k = (k-1)/\rho b_k$$

$$= (k-1)/(k\rho)F^(2)$$

How to prove lower bounds for randomized algorithms?

- Find a bad distribution of bids and show that no deterministic mechanism can fair well against it (**Minmax** / **Yao's Lemma**)
- What is the worst distribution?
- Theorem: For distributions which select the bids independently, the distribution with the highest competitive ratio has cumulative distribution P[x]=1-1/x
- Why?

How to prove lower bounds for randomized algorithms?

- Lemma:

Let D_1, D_2 be two probability distributions with cumulative distributions F_1, F_2 such that $F_1(x) \leq F_2(x)$ for every x. Let also $G : R^n \to R$ be a function which is non-decreasing in all its variables. Then

$$E_{b \in D_1^n}[G(b)] \geq E_{b \in D_2^n}[G(b)]$$

- The important condition in the proof is that the values in $b \in \mathrm{R}_+^n$ are independent.

How to prove lower bounds for randomized algorithms?

- Let $G = F^{(2)}$, which is non-decreasing in every bid
- Fix a distribution F(x) of the bids
- By scaling, assume that the online profit is 1
- Let

$$F_1(x) = \begin{cases} 0 & x < 1 \\ 1 - \frac{1}{x} & x \geq 1 \end{cases} \qquad F_2(x) = F(x)$$

- The lemma gives that F_1 is the worst distribution