Online Learning in Online Auctions

Avrim Blum[1], Vijay Kumar[2], Atri Rudra[3] and Felix Wu[4]

1: *CMU*, 2: *Amazon.com*, 3: *UT Austin*, 4: *UC Berkeley*

Overview

- Auction Model
- Previous Work
- Online Learning
- Weighted Majority in Auctions
- Lower Bound
- Dynamic Posted Price Mechanism

Online Auction

- Bidders arrive one-by-one.
- Each bidder i wants a copy of the item.
- Bidder values the item at v_i, places a bid b_i.
- Assume bidder valuations fall within a fixed range, normalized to $[1, h]$.
- Auction function: for each bidder, given bids, determine:
  - ◊ allocation: whether bidder i wins the item (x_i)
  - ◊ payment: amount bidder pays if she wins (p_i)
- Function might be randomized
- Goal: maximize revenue

Digital Goods

- Examples: digital music, pay-per-view TV, online journal articles.
- Key feature: negligible marginal cost
  - ◊ Unlimited supply, no constraint on allocation.
- Digital goods sold continuously, not in a single round
- Auction is alternative to current “two-price" model:
  - ◊ Peer-to-peer downloading vs. standard sale.

Incentive Compatibleness

- Determine (x_i, p_i) for each bidder such that max profit for bidder happens at $b_i = v_i$.
- Example 1: Set $p_i = b_i$.
- Exmaple 2: Vickery Auction (offline setting)
  - ◊ Winner pays the second highest bidders bid.
- [BHW02] Incentive Compatible $\Leftrightarrow$ Auction function is bid-independent.
  - ◊ $p_i = f(b_1, \cdots, b_{i-1})$.
- Assumed that $b_i = v_i$.

Competitive Analysis

- Economists assume: bids are drawn from some population.
- Computer Scientists assume: bids chosen by adversary.
  - ◊ Compare outcome with appropriate benchmark.
- Let $\vec{b}$ denote a sequence of bidder valuations.
- Benchmark: *optimal fixed price revenue* $\mathcal{F}(\vec{b}) = \max_{f_i \in [h]}(n_i \cdot f_i)$.
- Proportional to the revenue of the optimal one-round auction.
- Competitve ratio: $\max_{\vec{b}} \frac{\mathcal{F}(\vec{b})}{R(\vec{b})}$.

Competitive Analysis (contd.)

	2	3	1	2	4
4					4
3		3			3
2	2	2		2	2
1	1	1	1	1	1

- [BHW02] shows that *any* deterministic auction is $\Omega(h)$ comeptitive.
- Consider only Randomized auctions.

Prior Work

- This model:
  - [Bar-Yossef, Hildrum, Wu 2002]
- Offline Auctions for Digital goods
  - [Goldberg, Hartline, Wright 2001]
  - [Fiat et. al. 2002]
- Online Auctions with limited supply:
  - [Lavi and Nisan 2000]
  - [Bagchi et. al. 2001]

Online Learning

- Time steps and experts
- At each step:
  - ◊ Learning algorithm choses an expert.
  - ◊ Each expert achieves some gain (chosen by the adversary).
  - ◊ Learning algortihm achieves the gain of its chosen expert.
- Goal: achieve a total gain close to that of the best single expert.
- Weighted Majority algorithm ([Littlestone and Warmuth 1994].
- Auction setting: fixed prices = experts.

WM algorithm

Parameters: A real number $\alpha \in (0, 1]$.

Initialization: Initialize $r_k(\vec{b}) = 0,\ w_k(0) = 1$ for each expert k.

Repeat for $i = 1, 2, \cdots$ until no more bids

Select for bid i the price s_i to be u_k with probability

$$p_k(i) = \Pr[s_i = u_k] = \frac{w_k(i-1)}{\sum_{j=1}^{\ell} w_j(i-1)}$$

Observe b_i.

Update $r_k(\vec{b})$ and $w_k(i) = (1+\alpha)^{r_k(b_1,\ldots,b_i)/h}$ for each expert k.

Results

- WM is $(1+\epsilon)$-competitive when $\mathcal{F} \geq c(\epsilon) h \log\log h$. (Auer et al)
- Theorem: WM is $\omega(1)$ competitive for $\mathcal{F} = o(h \log\log h)$.
- Proof idea: Consider bid sequence of n bids of value 1.
  - ◇ $\frac{h-1}{(1+\alpha)^{\frac{n}{h}}}$ is constant.
  - ◇ Note $P_1(i) = \frac{(1+\alpha)^{\frac{i}{h}}}{h-1+(1+\alpha)^{\frac{i}{h}}}$.

Lower Bound

- Theorem: Any constant-competitive auction has $\mathcal{F} = \Omega(h)$.
- No auction can predict the first bid well.

Comparison with Previous Results

- This paper:
  - ◊ Weight function : Exponential.
  - ◊ Competitive ratio : $O(1)$.
  - ◊ Lower bound on $\mathcal{F}$: $\Omega(h \log \log h)$.
- [BHW02]:
  - ◊ Weight function : weighted buckets.
  - ◊ Competitive ratio : $O(\sqrt{\log h})$.
  - ◊ Lower bound on $\mathcal{F}$: $4h$.
- Constant ratio, $\mathcal{F} = \Omega(h)$.

Dynamic Posted Price Mechanism

- What if price is revealed to bidder ?
- Bidder either accepts or rejects, giving much less information
- Corresponds to the “bandit" problem in learning theory.
- Result: $(1+\epsilon)$-competitive, $\mathcal{F} = \Theta(h \log h \log \log h)$.

Open Questions

- Can we reduce the $\log \log h$ factor between the upper and lower bound for the full information game ?
- Is $\log h$ factor-loss necessary when going from full information to the partial information auctions ?