Near-optimal multi-unit auctions with ordered bidders

Elias Koutsoupias

UNIVERSITY OF OXFORD

London 2013/10/18

ESRC Workshop on Algorithmic Game Theory
London School of Economics

Auctions for digital goods

- Digital good: an unlimited supply of identical items
- Aim: design a sealed-bid auction for n bidders to maximize revenue
- Bidder i has private valuation v_i
- The auction proposes a take-it-or-leave price p_i

Truthfulness and optimality

- An auction is *truthful* if and only if the price p_i does not depend on v_i
- Objective: design an optimal (truthful) auction.
- Optimal?
  - The optimal revenue is

$$\sum_i v_i$$

(unattainable, because of truthfulness)

  - Bayesian setting:
    - the value v_i is drawn from a publicly-known probability distribution F_i
    - optimal price p_i maximizes

$$p_i \cdot (1 - F_i(p_i))$$

(Myerson 1981)

  - **In the prior-free setting?**

Prior-free benchmarks

Prior-free auction	**Benchmark**
does not know private values arbitrary prices	knows everything restricted prices

Obvious benchmark: $\sum_{\mathbf{i}} \mathbf{v_i}$

hopeless

Optimal single-price $\mathcal{F}^{(1)}$: $\max_{\mathbf{k}} \mathbf{k} \cdot \mathbf{v_{(k)}}$

too strict; for the same reason that first-price auction is not truthful; $v_{(k)}$ is the k-th maximum value

Optimal single-price $\mathcal{F}^{(2)}$: $\max_{\mathbf{p}\,:\,\mathbf{v_{(2)}} \geq \mathbf{p}} \sum_{\mathbf{i}} \begin{cases} p & \text{if } v_i \geq p \\ 0 & \text{otherwise} \end{cases}$

it sells at least two items

Decreasing prices $\mathcal{M}^{(2)}$: $\max_{\mathbf{p}\,:\,\mathbf{v_{(2)}} \geq \mathbf{p_1} \geq \mathbf{p_2} \cdots \geq \mathbf{p_n}} \sum_{\mathbf{i}} \begin{cases} p_i & \text{if } v_i \geq p_i \\ 0 & \text{otherwise} \end{cases}$

The $\mathcal{F}^{(2)}$ benchmark

An auction is said to have competitive ratio c against $\mathcal{F}^{(2)}$ if it has revenue at least $\mathcal{F}^{(2)}(\mathbf{v})/\mathbf{c}$, **for every set of values v**.

Lower bound: There is no auction with competitive ratio less than 2.42 [Goldberg, Hartline, Karlin, Saks, and Wright, 2006]

Upper bound: There is an auction with competitive ratio 3.25 [Hartline and McGrew, 2005]

Selecting a good benchmark for i.i.d. values

- Consider **all optimal Bayesian auctions** for (publicly-known) i.i.d. values
- These are exactly the auctions with a fixed-price (the reserve price that maximizes $\mathbf{p} \cdot (\mathbf{1} - \mathbf{F}(\mathbf{p}))$)
- The value of the best such auction for every set of values v is $\mathcal{F}^{(1)}(v)$

A prior-free auction which is competitive against $\mathcal{F}^{(1)}$, is approximately optimal for all distributions.

No such auction exist, so we settle for $\mathcal{F}^{(2)}$ and we get almost the same guarantee.

Is there a similarly good benchmark for the asymmetric case (that is, when the distributions of bidders may differ) ?

Selecting a good benchmark for non-i.i.d.'s

- If we repeat the reasoning for non-identical distributions, we don't get a meaningful benchmark (because the set of optimal Bayesian auctions contains all auctions)
- But if we have (non-identical) distributions
  - with decreasing reserve price, or
  - each distribution stochastically dominates the next one,

  then the corresponding benchmark is $\mathcal{M}^{(1)}$
- Since $\mathcal{M}^{(1)}$ is unattainable, we settle for $\mathcal{M}^{(2)}$

Therefore if an auction is competitive against $\mathcal{M}^{(2)}$, it is near optimal against all of the above distributions.

$\mathcal{M}^{(2)}$ vs $\mathcal{F}^{(2)}$

For every set v of n values:

$$\mathcal{F}^{(2)}(v) \leq \mathcal{M}^{(2)}(v) \leq \Theta(\ln n) \cdot \mathcal{F}^{(2)}(v)$$

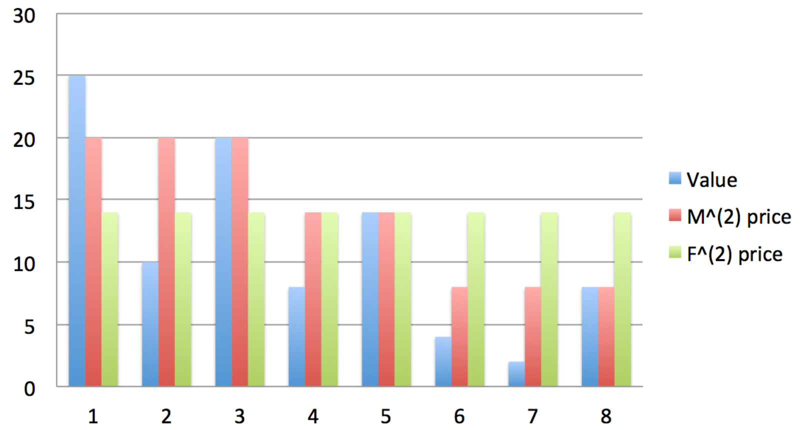

Competitive auctions against $\mathcal{M}^{(2)}$

- $\mathcal{M}^{(2)}$ was proposed in [Leonardi and Roughgarden, STOC 2012]
  They gave an auction with competitive ratio $\Theta(\log^* n)$.
- This talk, [**Bhattacharya, Koutsoupias, Kulkarni, Leonardi, Rughgarden, and Xu,** EC 2013]: we give an auction with constant competitive ratio

(In the paper, we also extended the result to limited supply auctions)

The Optimal Price Scaling auction

The main idea of the Optimal Price Scaling auction is:

1. Partition the bidders randomly into two parts A and B
2. Compute the optimal monotone prices for A and offer them to B (and vice versa)

For technical reasons, the actual auction is more complicated:

1. Use prices that are only powers of 2
2. With probability $1/2$, run the above scheme and with the remaining probability run a competitive ratio against $\mathcal{F}^{(\mathbf{2})}$

Performance

Theorem

The Optimal Price Scaling auction has constant competitive ratio against $\mathcal{M}^{(2)}$.

Our analysis proves a very high competitive ratio. It is an open problem to reduce it and find an almost matching lower bound.

Analysis

Let v be a set of values and let (v_A, v_B) be a random partition.
The analysis is based on

1. $\Pr\left[\mathcal{M}^{(2)}(v_A) \geq \frac{1}{3} \cdot \mathcal{M}^{(2)}(v)\right] \geq \frac{1}{16}$
2. The revenue extracted by the auction from B is $\Omega(\mathcal{M}^{(2)}(v_A))$

Price intervals

- Let p be a sequence of optimal monotone prices for v_A.
- Let J_k denote the interval of values in v for which the optimal price is 2^k. Call the values in J_k greater than 2^k, winning bids.

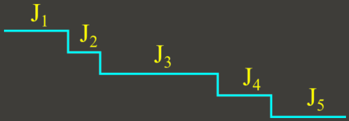

- It suffices to show that the winning bids of J_k are partitioned almost evenly between A and B (say with a ratio in $\left[\frac{1}{3}, \frac{2}{3}\right]$).

Intuitively, this seems true. But there are problems:

- The winning bids of some J_k's may be few (no concentration)
- A more subtle problem is that even when some J_k has many items, we cannot directly argue that the two parts have almost the same number of high values, because there is **a bias towards A,** by the way the levels were created.

Primal intervals

To resolve both issues, we use an intermediary set of some intervals, which we call *primal intervals*. They have the following properties:

- They are defined with respect to all values v

  (thus, no bias towards A or B)
- With constant probability, the winning bids of **all primal intervals** are almost evenly partitioned into A and B

  (therefore, we must have few primal intervals, and each one of them must be "long")
- There is a collection of primal intervals whose winning bids approximate the revenue of A

  (this can be achieved only if there are primal intervals of various sizes)

Despite these tradeoffs such sets of primal intervals exist.

Definition of primal intervals

Start with some primal interval. For every primal interval I, apply the following 5 rules to create new primal intervals.

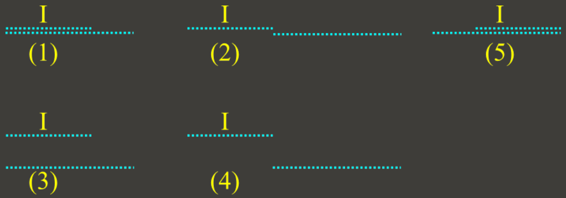

Matching and charging

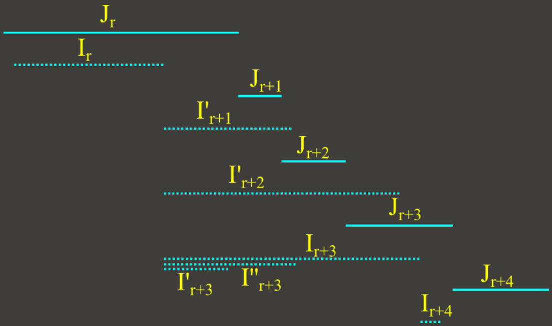

Open problems

- Tighten the analysis (lower and upper bounds)
- Find an (almost) optimal auction
- Extend to other asymmetric cases of bidders
- Extend to matroids

Thank you!