Yandex

Yandex

Reserve Pricing in Repeated Second-Price Auctions with Strategic Bidders

Alexey Drutsa

Setup

Second-Price (SP) Auction with Reserve Prices

Setting

- A good (e.g., an ad space) is offered for sale by a seller to M buyers
- Each buyer m holds a private valuation $v^m \in [0,1]$ for this good (v^m is unknown to the seller)

Actions

- The seller selects a reserve price p^m for each buyer m
- Each buyer m submits a bid b^m

Allocation and payments

- Determine actual buyer-participants: $\mathbb{M} = \{m \mid b^m \geq p^m\}$
- The good is received by the buyer $\bar{m} = \mathrm{argmax}_{m \in \mathbb{M}} b^m$ (that has the highest bid)
- This buyer pays $\overline{p^{\bar{m}}} = \max\{p^{\bar{m}}, \max_{m \in \mathbb{M} \setminus \{\bar{m}\}} b^m\}$

Repeated Second-Price Auctions with Reserve

Equal goods (e.g., ad spaces) are repeatedly offered for sale

- by a seller (e.g., RTB platform) to M buyers (e.g., advertisers)
- over T rounds (one good per round).

Each buyer m

- holds a private **fixed** valuation $v^m \in [0,1]$ for each of those goods,
- v^m is **unknown** to the seller.

At each round $t = 1, \dots, T$, the seller conducts SP auction with reserves:

- the seller selects a reserve price p_t^m for each buyer m
- and a bid b_t^m is submitted by each buyer m.

Seller's pricing algorithm

- The seller applies a pricing algorithm A that sets reserve prices $\{p_t^m\}_{t=1,m=1}^{T,M}$ in response to bids $\mathbf{b} = \{b_t^m\}_{t=1,m=1}^{T,M}$ of buyers $m = 1, \ldots, M$
- A price p_t^m can depend only on past bids $\{b_s^k\}_{s=1,k=1}^{t-1,M}$ and the horizon T.

Strategic buyers

- The seller announces her pricing algorithm A in advance

In each round t, each buyer m

- observes a history of previous rounds (available to this buyer) and
- chooses his bid b_t^m s.t. it maximizes his future γ_m-discounted surplus:

$$\mathrm{Sur}_t(A, v^m, \gamma_m, \{b_s^m\}) := \mathbb{E}\left[\sum_{s=t}^{T} \gamma_m^{s-1} \mathbb{I}_{\{m=\bar{m}_s\}} \left(v^m - \overline{p_s^m}\right)\right], \qquad \gamma_m \in (0,1],$$

where

$\mathbb{I}_{\{m=\bar{m}_s\}}$ is the indicator of the event when buyer m is the winner in round s

$\overline{p_s^m}$ is the payment of the buyer m in this case

Seller's goal

The seller's strategic regret:

$$\mathrm{SReg}(T, A, \{v^m\}_m, \{\gamma_m\}_m) := \sum_{t=1}^{T} \left(\max_m v^m - \mathbb{I}_{\{\mathbb{M}_t \neq \emptyset\}} p_t^{\overline{\bar{m}_t}}\right)$$

She seeks for a no-regret pricing for **worst-case** valuation:

$$\sup_{v^1,\dots,v^M \in [0,1]} \mathrm{SReg}(T, A, \{v^m\}_m, \{\gamma_m\}_m) = o(T)$$

Optimality: the lowest possible upper bound for the regret of the form $O(f(T))$.

Background,
Research question &
Main contribution

Background: 1-buyer case (posted-price auctions)

If one buyer ($M = 1$), a SP auction reduces to a posted-price auction:

- the buyer either accepts or rejects a currently offered price p_t^1
- the seller either gets payment equal to p_t^1 or nothing

[Kleinberg et al., FOCS'2003]	Optimal algorithm against **myopic** buyer with **truthful** regret $\Theta(\log\log T)$.
[Amin et al., NIPS'2013]	The strategic setting is introduced. $\nexists$ no-regret pricing for non-discount case $\gamma = 1$.
[Drutsa, WWW'2017]	Optimal algorithm against strategic buyer with regret $\Theta(\log\log T)$ for $\gamma < 1$.

Research question

The known optimal algorithms (PRRFES & prePRRFES) from posted-price auctions **cannot be directly applied** to set reserve prices in second-price auctions

- buyers in SP auctions have **incomplete information** due to presence of rivals
- the proofs of optimality of [pre]PRRFES strongly rely on complete information

In this study, I try to find an optimal algorithm for the multi-buyer setup

Main contribution

A novel algorithm for our strategic buyers
with regret upper bound of $\Theta(\log\log T)$ for $\gamma < 1$

A novel transformation that maps any pricing algorithm
designed for posted-price auctions to a multi-buyer setup

Main ideas

Two learning processes

$$\mathrm{SReg}(T, A, \{v^m\}_m, \{\gamma_m\}_m) := \sum_{t=1}^{T} \left(\max_m v^m - \mathbb{I}_{\{\mathbb{M}_t \neq \emptyset\}} \overline{p_t^{\overline{m}_t}}\right)$$

Find which buyer has the maximal valuation

Learning process #2

Find the buyers' valuations

Learning process #1

Learning proc.#1: an idea to localize a valuation

PRRFES is an optimal learner of a valuation in posted-price auctions.

However, its core localization technique relies on:

> The buyer completely knows the outcomes of current and all future rounds given their bids (due to absence of rivals)

Can we use PRRFES in the second-price scenario where each buyer does not know perfectly the outcomes of rounds?

Barrage pricing

- Reserve prices are personal (individual) in our setup
- Thus, we are able to “eliminate” particular buyers from particular rounds
- Namely, a buyer m will not bid above $1/(1-\gamma_m)$
- We call this price as “barrage” one and denote it by ∞

Let “eliminate” all buyers except some buyer m in a round t
Then the buyer m will have **complete information about outcome of this round** t

Learning proc.#2: an idea to find max valuation

The search algorithm works by maintaining a feasible interval $[u^m, w^m]$ that

- is aimed to localize the valuation v^m, i.e. $v^m \in [u^m, w^m]$
- shrinks as $t \to \infty$

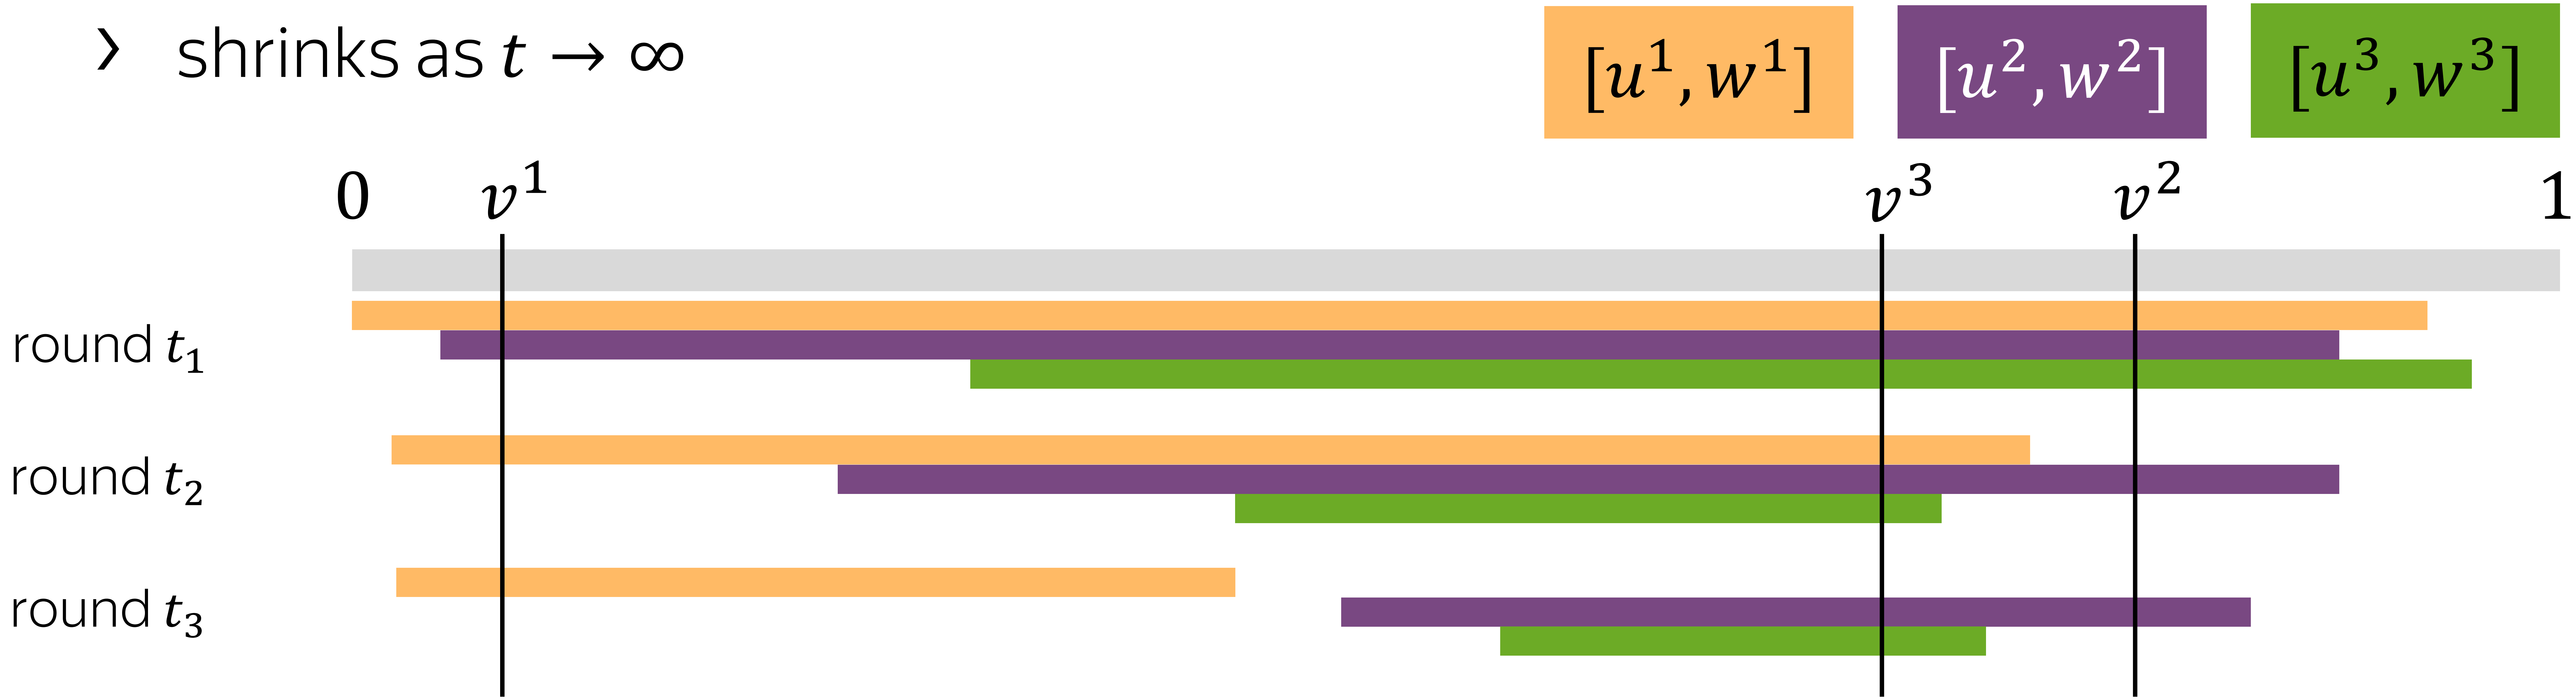

If, in a round t, it becomes that $w^m < u^n$ for some buyers m and n, then buyer m has non-maximal valuation which should not be searched anymore

Dividing algorithms

Key instrument that implements the ideas

transformation

div

Transformation **div**: cyclic elimination

Let A be an algorithm designed for repeated posted-price auctions

Its transformation $\mathbf{div}(A)$ is an algorithm for repeated SP auctions as follows

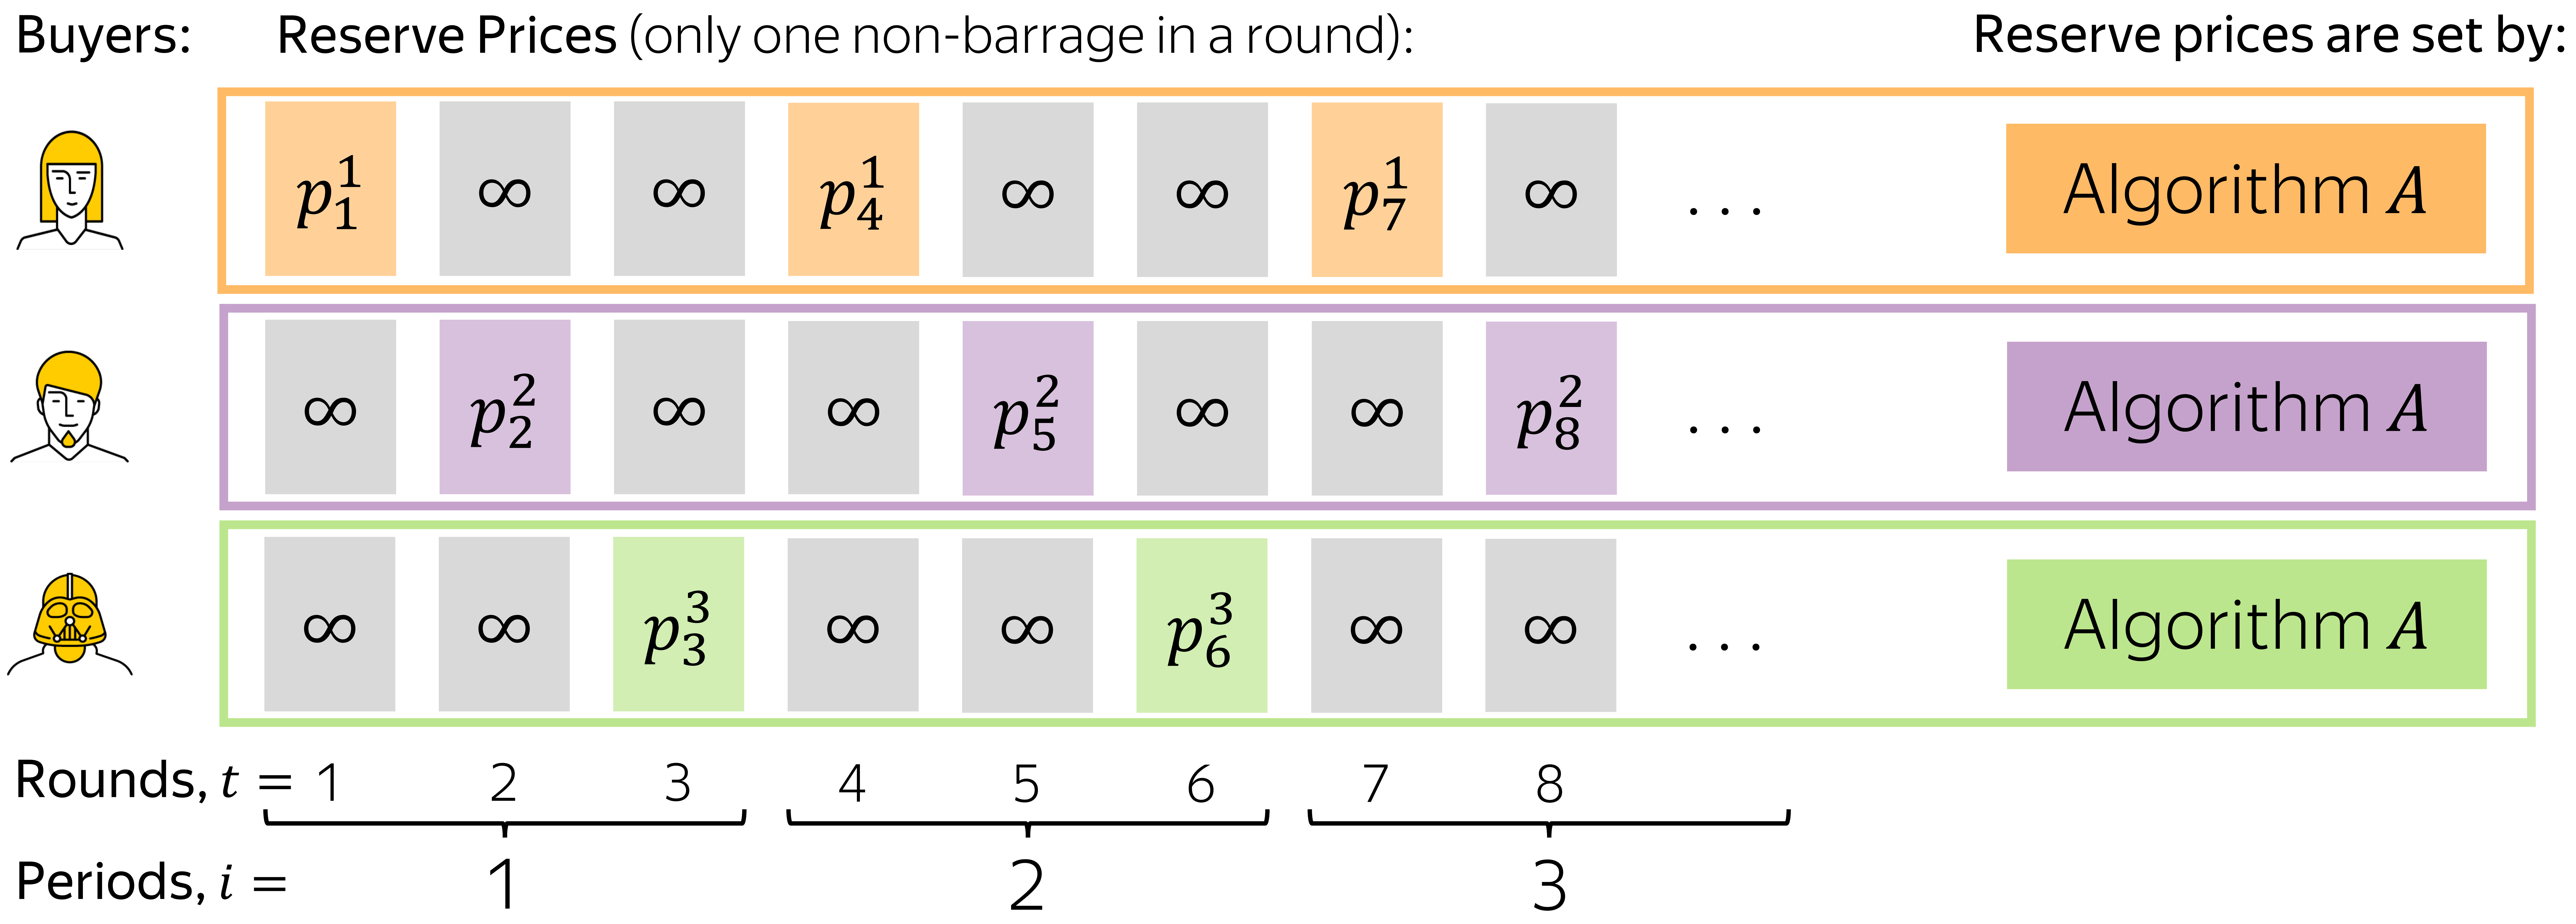

Transformation div: stopping rule

We stop considering a buyer m in periods when $w^m < u^n$ for some buyer n.

The number of periods with buyer m is referred to as subhorizon, I^m.

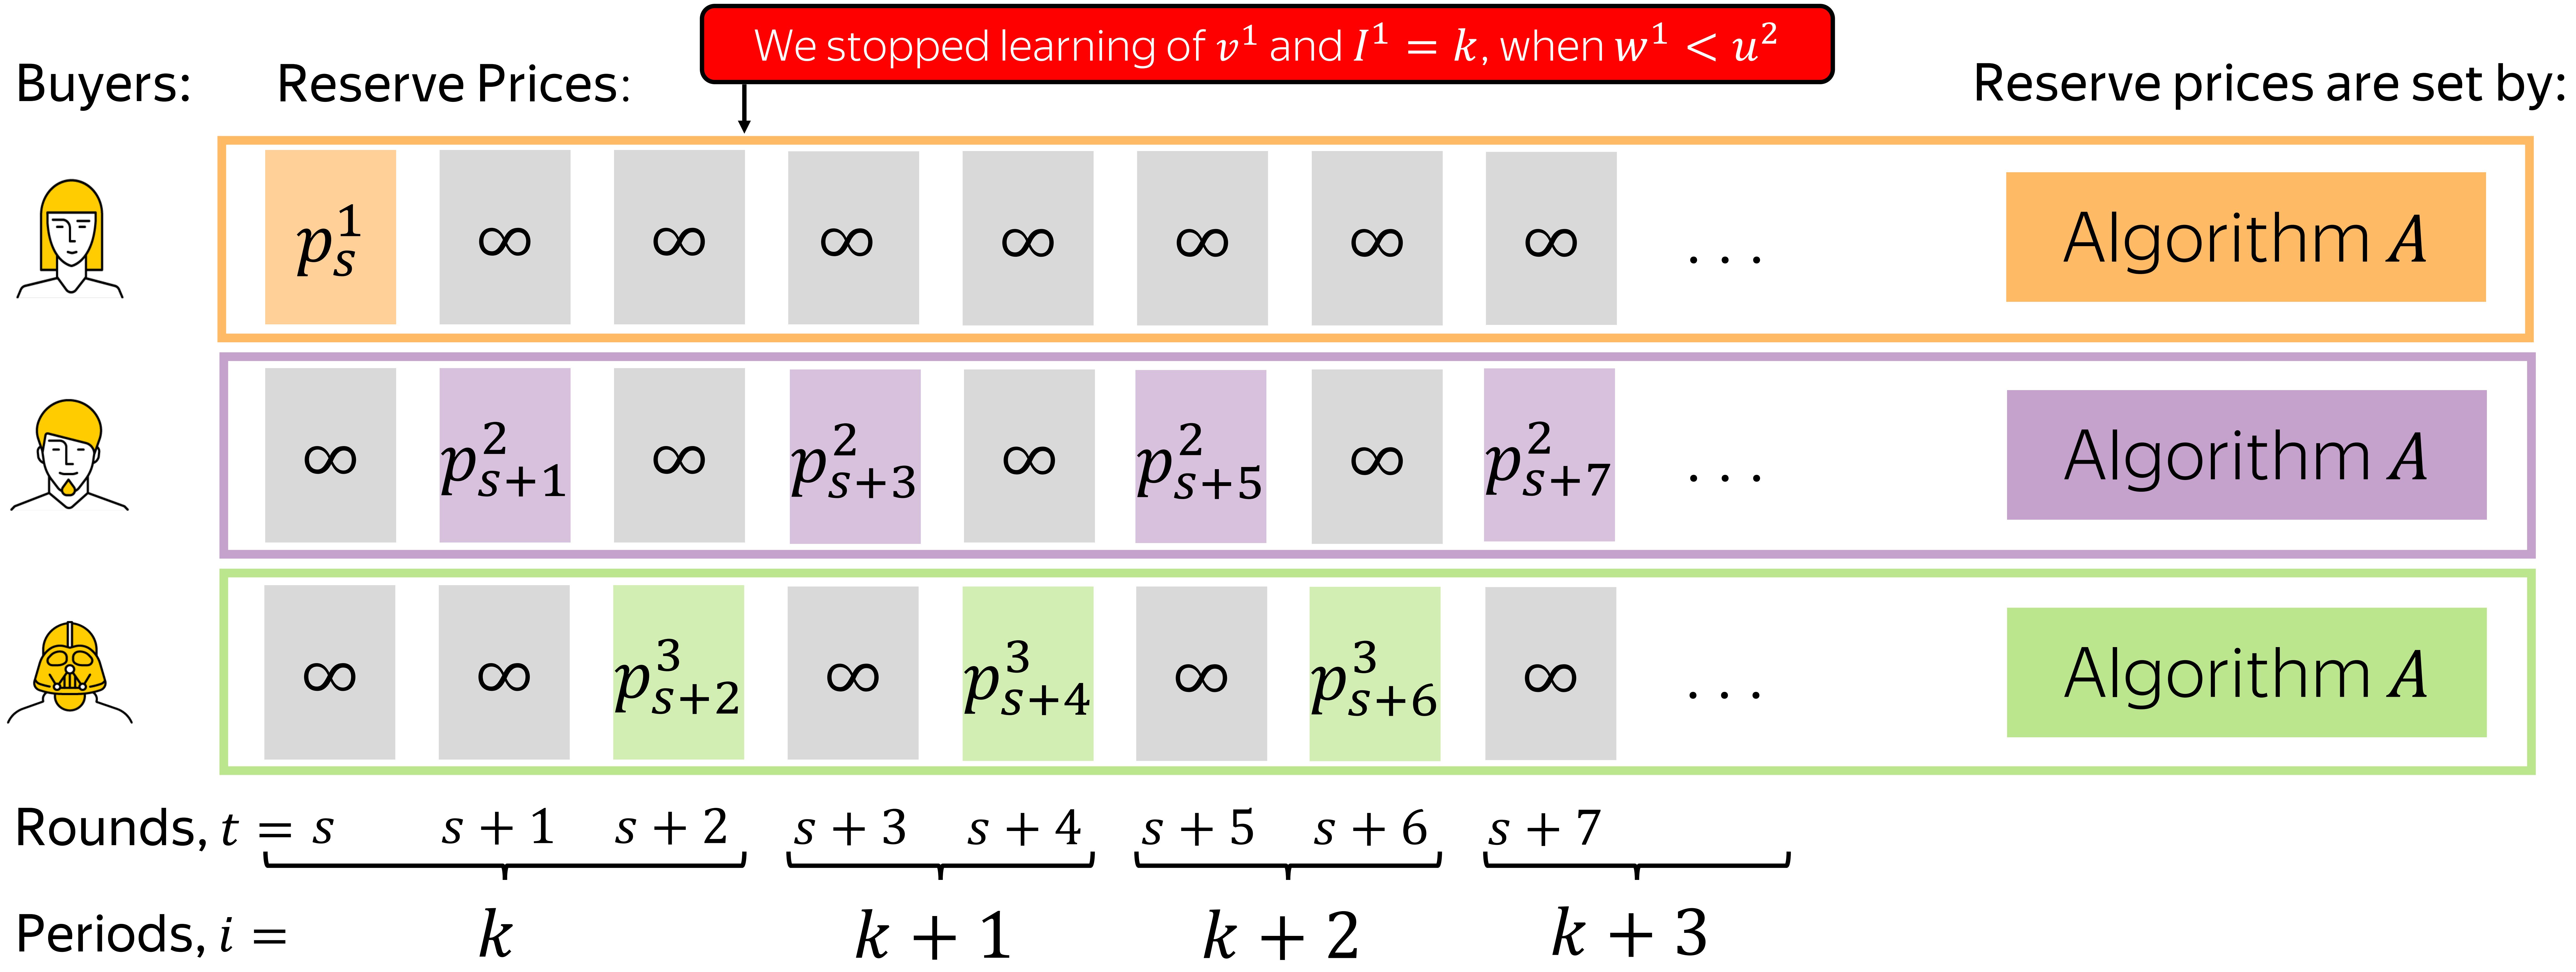

Transformation div: regret decomposition

Lemma 1. For the described transformation, strategic regret has decomposition:

$$\mathrm{SReg}(T, \mathbf{div}(A), \{v^m\}_m, \{\gamma_m\}_m) =$$

$$= \sum_m \mathrm{Reg}^m(T, A, v^m, \gamma^m) + \sum_m I^m(\max_n v^n - v^m)$$

Individual regrets

Measure how the algorithm A learns the valuation of each buyer

Deviation regret

Measures how fast we stop learning of non-maximal valuations

Key challenge against strategic buyer

Strategic buyer may lie and mislead algorithms, thus a good algorithm must
Extract correct information about a buyer’s valuation from his actions (bids)

> Dividing structure in a round allows to construct a tool to locate valuations:
> it is enough to make complete information situation in a round

Upper bound on valuation of strategic buyer

Let buyer m is the non-"eliminated " one in a round t.

- If the buyer accepts (bids above) the current reserve price p_t^m

$$\mathrm{Surplus}_t = \underbrace{\mathbb{E}\left[\gamma_m^{t-1}\mathbb{I}_{\{m=\overline{m}_t\}}\left(v^m - \overline{p_t^m}\right)\right]}_{\substack{\parallel \\ \gamma_m^{t-1}\mathbb{I}_{\{b_t^m \geq p_t^m\}}\left(v^m - p_t^m\right) = \gamma_m^{t-1}\left(v^m - p_t^m\right)}} + \underbrace{\mathbb{E}\left[\sum_{s=t+1}^{T}\gamma_m^{s-1}\mathbb{I}_{\{m=\overline{m}_s\}}\left(v^m - \overline{p_s^m}\right)\right]}_{\substack{\text{VI} \\ 0}}$$

- If the buyer rejects (bids below) the current reserve price p_t^m

$$\mathrm{Surplus}_t = \mathbb{E}\left[\sum_{s=t+r}^{T}\gamma_m^{s-1}\mathbb{I}_{\{m=\overline{m}_s\}}\left(v^m - \overline{p_s^m}\right)\right] \leq \frac{\gamma_m^{t+r-1}}{1-\gamma_m}\left(v^m - [\text{lowest\_price}]\right)$$

If we observe that a buyer rejects non-"barrage" reserve price, then:

$$v^m - p_t^m < \frac{\gamma_m^r}{1-\gamma_m-\gamma_m^r}\left(p_t^m - [\text{lowest\_price}]\right)$$

Optimal algorithm

Pricing algorithm divPRRFES

Apply the transformation **div**

to PRRFES algorithm

divPRRFES: individual and deviation regrets

Individual regrets

Our tool to locate valuations provides the upper bound (as in 1-buyer case):

$$\mathrm{Reg}^m(T, A, v^m, \gamma^m) = O(\log_2 \log_2 T)\ \forall m$$

Deviation regrets

- For each buyer m with non-maximal valuation (i.e., $v_m < \max_n v^n$)
- We can upper bound its subhorizon I^m:

$$I^m \leq \frac{C}{\max_n v^n - v_m}$$

divPRRFES is optimal

<u>Theorem.</u>

Let $\gamma_0 \in (0,1)$

Then for the pricing algorithm divPRRFES A with:

- the number of penalization rounds $r \geq \left\lceil \log_{\gamma_0} \frac{1-\gamma_0}{2} \right\rceil$ and
- the exploitation rate $g(l) = 2^{2^l}, l \in \mathbb{Z}_+$,

for any valuations $v^1, \ldots, v^M \in [0,1]$, any discounts $\gamma_1, \ldots, \gamma_M \in [0, \gamma_0]$, and $T \geq 2$, the strategic regret is upper bounded:

$$\mathrm{SReg}(T, A, \{v^m\}_m, \{\gamma_m\}_m) \leq C(\log_2 \log_2 T + 2) + B,$$

$$C := M\left(r \max_m v^m + 4\right), \qquad B := (24 + 5r)(M - 1).$$

Summary

Main contribution: reminding

A novel algorithm for setting reserve prices in second-price auctions with strategic buyers.
Its worst-case regret is optimal: $\Theta(\log\log T)$ for $\gamma < 1$

A novel transformation that maps any pricing algorithm designed for posted-price auction to a multi-buyer setups

Thank you!

Alexey Drutsa

Yandex

adrutsa@yandex.ru