Secretaries and Prophets

Algorithmic Game Theory

Winter 2019/20

Online Auctions

Secretaries and Matching

Matroid Secretary

Prophet Inequalities

Online Auctions

In many applications potential customers are not all simultaneously in the market, they come and go over time. We study mechanisms to sell a single item online, i.e., in a scenario where bidders arrive and depart sequentially one by one.

Online Single-Item Auction

- In round $t = 1, 2 \ldots$ bidder i_t arrives and reports a value v_{i_t} for the item
- Mechanism decides immediately (without knowning future bidders $i_{t+1}, i_{t+2}, \ldots$) whether i_t gets the item or not, and his payment.
- All decisions about item allocation and payments are *final* and *irrevocable*.

Without knowledge about the number and the maximum value of the bidders it is in the worst-case impossible to assign the item to the highest bidder. In terms of approximation of social welfare every mechanism is extremely bad in the worst case. Thus, worst-case analysis is not informative. Instead, we focus on two stochastic models for analysis with slightly different assumptions on arrival times and values of the bidders.

Online Auctions

Secretaries and Matching

Matroid Secretary

Prophet Inequalities

Random-Order Model

Random-Order Model

- Values of bidders are unknown
- Number n of bidders is known
- Bidders arrive in uniform random order

We strive to maximize social welfare, i.e., assign the item to the highest bidder. This is the standard secretary problem: Find the bidder with highest value in a uniform random arrival order.

Secretary Algorithm

Let $r \in \{1, \ldots, n\}$ be the sample size

Sample: In round $t = 1, \ldots, r$ bidder i_t reports value v_{i_t}. He is rejected.

Acceptance: In round $t = r+1, \ldots, n$ bidder i_t reports value v_{i_t}. If item is still available and i_t highest bidder so far, assign item to i_t.

Competitive Ratio

We evaluate online algorithms using the competitive ratio:

- S^* is an optimal feasible subset of bidders that maximizes $\sum_{i \in S^*} v_i$.
- Online algorithm is c-competitive: Outputs feasible set T of bidders with

$$\mathbb{E}\left[\sum_{i \in T} v_i\right] \geq \frac{1}{c} \cdot \sum_{i \in S^*} v_i$$

  c is competitive ratio of the algorithm.
- Note: T is a **random subset**. Randomization based on random arrival order of bidders and (possibly) internal randomization of the algorithm.

Consider the algorithm above. For $r = \lfloor n/e \rfloor$ the algorithm assigns the item to the **highest bidder** with a **probability of at least** $1/e$. Hence:

Proposition

The Secretary algorithm with $r = \lfloor n/e \rfloor$ is e-competitive.

Incentive Compatibility

A bidder shall have no incentive to lie, even if he knows

- all values of other bidders,
- the arrival order of all bidders,
- the mechanism to choose allocation and payments, and
- possible internal random bits for internal randomization of the mechanism.

Hence, the bidder shall have at the end in hindsight and for every possible outcome of the random choices an incentive to report his value truthfully.

IC here means that the mechanism is (ex-post universally) truthful.

Consider the Secretary algorithm. Are there payments to turns it into an IC mechanism?

The algorithm must be monotone. Also, we must assign payments directly at the time of allocation, without knowing values of future bidders...

An IC Single-Item Online Auction

Let $\tau = \max_{t=1,\ldots,r} v_{i_t}$ be the best value in the sample phase. If the mechanism assigns the item, then let i_{t^*} be the receiving bidder. Assign payments $p_{i_{t^*}}(v) = \tau$, and $p_{i_t}(v) = 0$ für all $t \neq t^*$.

Proposition

The Secretary mechanism is incentive compatible and e-competitive for the single-item auction in the random-order model.

Proof:
Consider the bidders in hindsight in the order they arrived.

- No bidder i_t with $t > t^*$ or $t \leq r$ can make the algorithm assign him the item by unilaterally changing his reported value.
- For bidder i_t with $t = r+1,\ldots,t^*$ the mechanism assigns the item to the earliest bidder with $v_{i_t} \geq \tau$ at price τ.

Bidder i_{t^*} with $v_{i_{t^*}} \geq \tau$ gets the item, pays τ and is happy with true bid. Every bidder i_t with $t = r,\ldots,t^*-1$ has $v_{i_t} \leq \tau$, gets and pays nothing, and is happy with true bid. □

Multiple Items

Multi-Item Markets and Unit-Demand Bidders

- Set I of n bidders
- Set J of m items
- Bidder $i \in I$ has private values $v_{ij} \geq 0$ for every $j \in J$.
- It is known that bidder i has unit demand, i.e., he values only the best received item

$$v_i(S) = \max_{j \in S} v_{ij} \quad .$$

The allocation with maximum social welfare can be given by a max-weight matching M^* in the complete bipartite graph G with vertex set $I \cup J$.

Online Mechanisms for Unit-Demand Auctions

Online Mechanisms

- Bidders arrive sequentially in random order
- All items available upfront, number n of bidders known upfront
- Bidders arrive in uniform random order, report private values upon arrival
- Online mechanism decides: Which item (if any) should be assigned to the bidder? What does he pay?
- Allocation and payment must be decided before next bidder arrives
- Every decision is final and irrevocable

Mechanism Objectives:

1. ex-post universally incentive compatible
2. approximate social welfare $\sum_{i \in I} v_i(S)$ as good as possible
3. polynomial-time computation

Online VCG Mechanism for Unit-Demand Auctions

1. Init assignment $M \leftarrow \emptyset$ and payments $p_i(v) \leftarrow 0$ for all $i \in I$
2. Reject the first r bidders
3. In round $t = r+1, \ldots, n$ do:
4. Bidder i_t arrives
5. I_t is set of all bidders arrived so far. J_t is set of unassigned items.
6. $M_t \leftarrow$ max-weight matching in induced subgraph $G_t = G \cap (I_t \cup J_t)$
7. $v(M_t) \leftarrow \sum_{(k,j)\in M_t} v_{kj}$ the social welfare of M_t.
8. **if** i_t matched to $j \in J_t$ in M_t **then**
9. Assign j to i_t and add (i_t, j) to M.
10. i_t gets VCG payment

$$p_{i_t}(v) = v(M_t') - v(M_t \setminus (i_t, j)),$$

where M_t' is max-weight matching in $G_t' = G \cap ((I_t \setminus \{i_t\}) \cup J_t)$.

Example

M at the end of round 3:

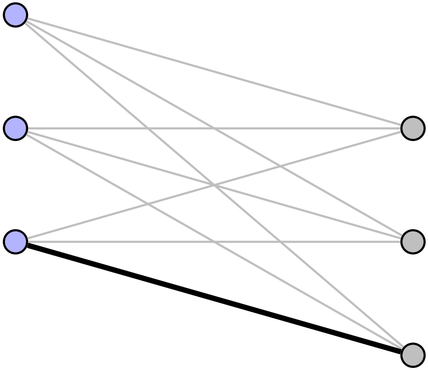

Example

Bidder i_4 arrives and reports values for all items:

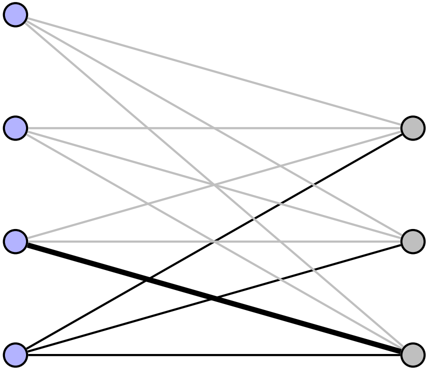

Example

Bipartite graph G_4:

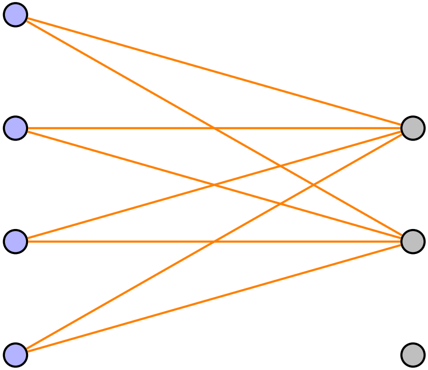

Example

Max-weight matching M_4 in G_4 leaves i_4 unmatched:

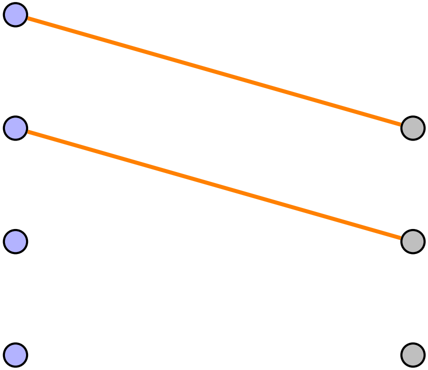

Example

M at the end of round 4:

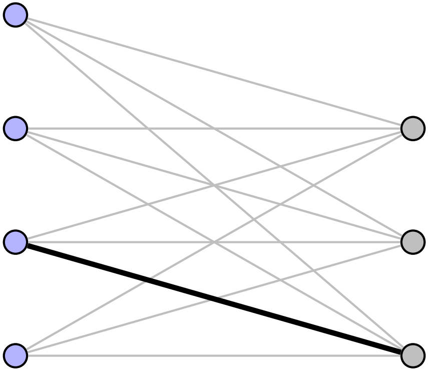

Example

Bidder i_5 arrives and reports values for all items:

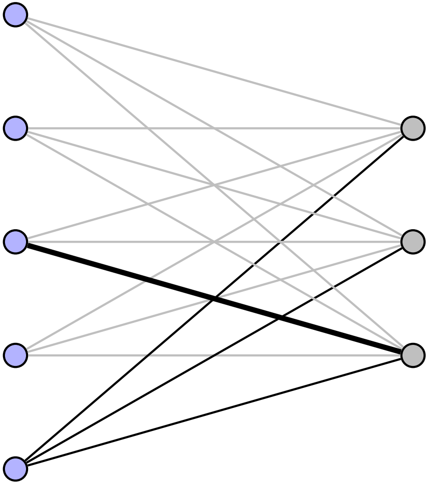

Example

Bipartite graph G_5:

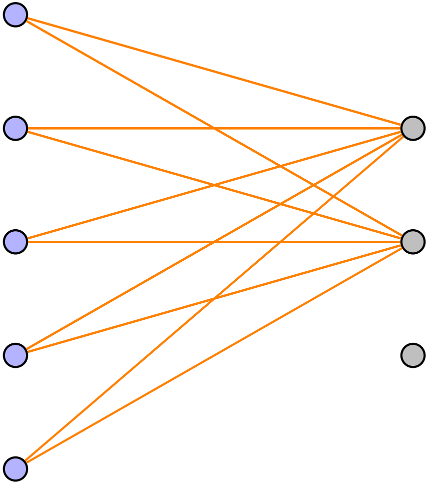

Example

Max-weight matching M_5 in G_5 matches i_5 to item 1:

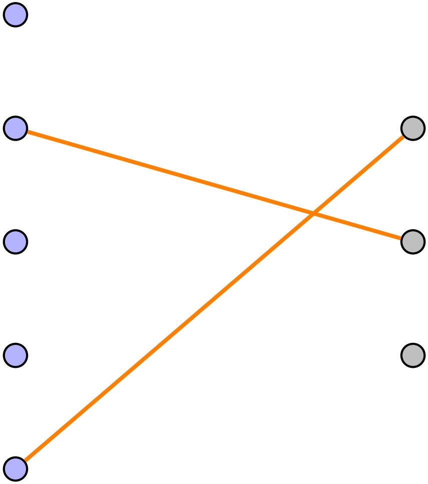

Example

M at the end of round 5:

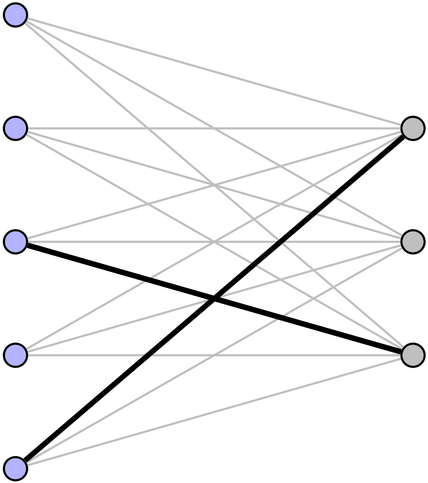

Online VCG

Proposition

The Online VCG mechanism is incentive compatible and can be implemented in polynomial time.

Proof:
Bidder i_t cannot change the set I_t of arrived bidders or the set J_t of remaining unassigned items. When we determine his assignment and payments in round t, he faces a standard VCG mechanism that assigns items in J_t to unit-demand bidders in I_t in an IC fashion.

Computationally most demanding are steps 6 and 9, where we need to compute max-weight bipartite matchings. This can be done in polynomial time with linear programming (c.f. "Theoretische Informatik I"). □

Theorem (Reiffenhäuser, 2019)

The Online VCG mechanism with $r = \lfloor n/e \rfloor$ is e-competitive for multi-item unit-demand auctions in the random-order model.

Online VCG is e-competitive

All sets I_t and J_t are random variables.

If we fix the set I_t of bidders and their order of arrival over the first t rounds, then the set J_{t+1} of unassigned items at the end of round t is fully determined, since VCG uses no randomness.

The **key technical idea** of the theorem is the following:

If we fix the set I_t of bidders but assume random arrival order of bidders in I_t over the first t rounds, then the every set J_{t+1} of k items has the same probability to remain. Moreover, this probability is stochastically independent of I_t.

Formally, consider $\Pr[I_t, J_{t+1}]$ the probability that a given set of bidders I_t arrives in the first t rounds and a set J_{t+1} remains in the end of round t.

Lemma (Independence Lemma)

Consider any round $t = 1, \ldots, n$, any two subsets $I_t, I'_t \subseteq I$ with $|I_t| = |I'_t| = t$ and any two subsets $J_{t+1}, J'_{t+1} \subseteq J$ with $|J_{t+1}| = |J'_{t+1}|$. Then

$$\Pr[I_t, J_{t+1}] \quad = \quad \Pr[I'_t, J'_{t+1}] \ .$$

Online VCG is e-competitive

The proof is easy for a single item with $|J| = 1$, in the standard secretary case (Exercise).

For several items, the statement of the lemma can be proven via induction over the rounds t. We omit the proof.

The independence lemma can be used to show the following:

Lemma

For any given item $j \in J$, the probability that j remains unassigned at the end of round t is $\Pr[j \in J_{t+1}] = r/t$.

Proof:
Suppose we fix the two sets I_t, J_t, and an item $j \in J_t$.
Now M_t is determined, and so is the bidder $i_{M_t}(j)$ matched to j in M_t (if any).
Note that

$$\Pr[j \text{ gets assigned in round } t \mid I_t, J_t, j] \;=\; \Pr[i_{M_t}(j) \text{ arrives in round } t \mid I_t, J_t, j]$$

Online VCG is e-competitive

For any two bidders i^1, i^2, consider the event that i^1 arrives in round t and the subset $I^1_{t-1} = I_t \setminus \{i_1\}$ arrived in the first $t-1$ rounds. Similarly, the event i^2 comes in round t and $I^2_{t-1} = I_t \setminus \{i_2\}$ before.

Due to the independence lemma

$$\Pr[I^1_{t-1}, J_t] = \Pr[I^2_{t-1}, J_t] \quad .$$

No matter which bidder of I_t is i_t, the item set J_t has the same probability to remain. In reverse, if we assume J_t is the set of remaining items before the start of round t, then each bidder $i \in I_t$ has the same probability (of $1/t$) to be i_t. This holds, in particular, for $i_{M_t}(j)$. Hence:

$$\Pr[i_{M_t}(j) \text{ arrives in round } t \mid I_t, J_t, j] = 1/t$$

Note that this is the same quantity for all sets I_t, J_t and items $j \in J_t$.

Online VCG is e-competitive

Each $j \in J_t$ can be assigned in at most one round. The probability that j survives in a given round t is always $1 - 1/t$. Hence, the probability that it survives from the beginning until the end of t rounds is

$$\Pr[j \in J_{t+1}] = \prod_{k=r+1}^{t} \left(1 - \frac{1}{t}\right) = \frac{r}{r+1} \cdot \frac{r+1}{r+2} \cdots \frac{t-1}{t} = \frac{r}{t} \ .$$

□

Consequently, the expected number of remaining items before start of round t is

$$\mathbb{E}[|J_t|] = \sum_{j \in J} \Pr[j \in J_t] = \frac{mr}{t-1}.$$

Online VCG is e-competitive

Consider edge $e_t = (i_t, j)$ added to the matching in round t and let $v(e_t) = v_{i_t,j}$ be the value. If no edge is added in round t, we assume $e_t = \emptyset$ and $v(\emptyset) = 0$.

Clearly, e_t is a random variable.

Let M^* be an offline optimum, i.e., a max-weight matching in G.

Lemma

The expected value of e_t is

$$\mathbb{E}[v(e_t)] \quad \geq \quad \frac{r}{n(t-1)} \cdot v(M^*) \ .$$

Online VCG is e-competitive

Proof:
Suppose the set of remaining items J_t has cardinality k and consider the graph G_t. I_t is a uniform random size-t-subset of I.

Due to the independence lemma, J_t is a uniform random size-k-subset of J.

Hence, each edge $e \in M^*$ has a probability of $\frac{t}{n} \cdot \frac{k}{m}$ to be included into G_t. This implies

$$\begin{aligned} \mathbb{E}[v(M_t) \mid k = |J_t|] &\geq \mathbb{E}[v(M^* \cap I_t \times J_t) \mid k = |J_t|] \\ &= \mathbb{E}[\sum_{(i,j) \in M^* \cap I_t \times J_t} v_{ij} \mid k = |J_t|] \\ &= \sum_{(i,j) \in M^*} v_{ij} \cdot \Pr[i \in I_t, j \in J_t \mid k = |J_t|] \\ &= v(M^*) \cdot \frac{t}{n} \cdot \frac{k}{m} \quad . \end{aligned}$$

Online VCG is e-competitive

We saw in the previous lemma that any edge $(i, j) \in M_t$ is added to M in round t if $i = i_t$, which happens with probability $1/t$, even if we condition on I_t and J_t. Therefore

$$\mathbb{E}[v(e_t) \mid k = |J_t|] \quad = \quad \frac{1}{t} \cdot \mathbb{E}[v(M_t) \mid k = |J_t|] \quad \geq \quad v(M^*) \cdot \frac{1}{n} \cdot \frac{k}{m} \ ,$$

and this implies

$$\begin{aligned}
\mathbb{E}[v(e_t)] &= \sum_{k=0}^{n} \Pr[J_t = k] \cdot \mathbb{E}[v(e_t) \mid k = |J_t|] \\
&\geq v(M^*) \cdot \frac{1}{nm} \cdot \sum_{k=0}^{m} \Pr[|J_t| = k] \cdot k \\
&= v(M^*) \cdot \frac{1}{nm} \cdot \mathbb{E}[|J_t|] \quad = \quad \frac{r}{n(t-1)} \cdot v(M^*) \ .
\end{aligned}$$

□

Online VCG is e-competitive

Finally, we prove the theorem by applying the last lemma.

Proof (Theorem):
The overall assignment M of the algorithm is composed of all edges added in rounds $t = r+1, \ldots, n$. Due to linearity of expectation, we see

$$\begin{aligned}
\mathbb{E}[v(M)] &= \mathbb{E}\left[\sum_{t=r+1}^{n} v(e_t)\right] = \sum_{t=r+1}^{n} \mathbb{E}[v(e_t)] \\
&\geq \sum_{t=r+1}^{n} \frac{r}{n(t-1)} \cdot v(M^*) \\
&= \frac{\lfloor n/e \rfloor}{n} \cdot v(M^*) \cdot \sum_{t=\lfloor n/e \rfloor+1}^{n} \frac{1}{t-1} \\
&\approx \frac{1}{e} \cdot v(M^*) \cdot (\ln(n) - \ln(n/e)) = \frac{1}{e} \cdot v(M^*) \cdot \ln \frac{n}{n/e} \\
&= \frac{1}{e} \cdot v(M^*) \ .
\end{aligned}$$

□

The estimation in line 4 deteriorates the ratio by at most an additive term $1/n$.

Online Auctions

Secretaries and Matching

Matroid Secretary

Prophet Inequalities

Matroid Auction

Packing Bidders

- Set I of n bidders
- Bidder $i \in I$ has single private value $v_i \geq 0$
- Every outcome $S \in A \subseteq 2^I$ is a subset of bidders.
- Single-parameter domain, binary amount of stuff:
  $x_i(S) = 1$ if $i \in S$ and 0 otherwise.
- A yields a matroid packing problem, i.e. the pair (I, A) is a matroid (definition of "matroid" later).

Some Examples:

- Single-Item Auction: $A = \{S \mid 1 \geq |S|\} = I \cup \{\emptyset\}$
- k-Item Auction: $A = \{S \mid k \geq |S|\}$
- Forest Auction: Every bidder is an edge in a known graph G. The auction can choose an acyclic subset of edges/bidders
  $A = \{S \mid S$ is acyclic set of edges in $G\}$

Online Mechanisms for Matroid Auctions

Online Mechanisms

- I and A are known upfront
- Bidders arrive in uniform random order, report private value upon arrival
- Online mechanism decides: Accept or reject the bidder? If accepted, what does he pay? If rejected, payment 0.
- Every bidder is decided before next bidder arrives
- Decision on accept/reject and payment is final and irrevocable

Mechanism Objectives;

1. ex-post universally incentive compatible
2. approximate social welfare $\sum_{i \in S} v_i$ as good as possible
3. polynomial-time computation

Forest Auction

We first consider online mechanisms for graphic matroids:

Graphic Matroid

- Bidders are edges of a connected graph $G = (K, E)$
- Outcome set A composed of all forests (acyclic sets of edges).
- Every $S \in A$ with maximum cardinality is spanning tree of G
- Every spanning tree consists of $k = |K| - 1$ edges
- Bidder i has private value v_i, we denote $v(S) = \sum_{i \in S} v_i$

Kruskal greedy algorithm: Computes an optimal spanning tree with maximum total value

Exchange property: S spanning tree, S^* optimal spanning tree. There are pairs of bidders $(i_1, i_1^*), \ldots, (i_k, i_k^*) \in S \times S^*$ such that for all $1 \leq t \leq k$

- $S_t = (S \cup \{i_1^*, \ldots, i_t^*\}) \setminus \{i_1, \ldots, i_t\}$ is spanning tree
- $S^* = S_k$
- $v(S_t) \geq v(S_{t-1})$

Random Threshold Mechanism

1. $p_i(v) \leftarrow 0$ for all bidders
2. Reject the first $n/2$ bidders, we denote this set by Y.
3. Choose $j \in \{0, 1, 2, \ldots, \lceil \log k \rceil\}$ uniformly at random
4. Set threshold $\tau \leftarrow \max_{x \in Y} v_x / 2^j$, set $S \leftarrow \emptyset$
5. In round $t = n/2 + 1, \ldots, n$ do:
6. Consider arriving bidder i_t
7. **if** $v_{i_t} \geq \tau$ and $(S \cup \{i_t\}) \in A$ **then**
8. $S \leftarrow S \cup \{i_t\}$ and $p_{i_t}(v) \leftarrow \tau$.

Theorem

The Random Threshold mechanism is incentive compatible and can be implemented in polynomial time. It is $O(\log k)$-competitive for forest auctions.

Random Threshold Mechanism

Proof:
IC can be shown similarly to the secretary mechanism (Exercise). Polynomial running time is straightforward. We only show the competitive ratio.

We concentrate on bidders from S^* with significant value.

- Let S^* be an optimal spanning tree. We number the k bidders from S^* by $1, \ldots, k$ in non-increasing order of value $v_1 \geq \ldots \geq v_k$.
- Note: v_1 is the highest Bidder from I (c.f. Kruskal algorithm), but the other are **not necessarily the next $k-1$ bidders** with highest values from I.
- Choose q such that: $(v_q \geq v_1/k) \wedge ((q = k) \vee (v_{q+1} < v_1/k))$.
- Note:

$$\sum_{i=q+1}^{k} v_i < v_1, \quad \text{and hence} \quad \sum_{i=1}^{q} v_i \geq v(S^*)/2.$$

- The q highest bidders in S^* yield at least half of the optimal social welfare.

Logarithmic Ratio

The analyis relies on classes of values based on bidders $i = 1, \ldots, k$ from S^*. W.l.o.g. we assume that $v_1 > v_2 > \ldots > v_k$.

- There are exactly i bidders in S^* with value at least v_i.
- Let $m_i(T)$ be the number of bidders in $T \subset I$ with value at least $v_i/2$.
- We see that

$$\sum_{i=1}^{q} v_i \quad = \quad \left[\sum_{i=1}^{q-1} (v_i - v_{i+1}) \cdot i\right] + v_q \cdot q.$$

- For every set T

$$v(T) \quad \geq \quad \frac{1}{2} \cdot \left[\sum_{i=1}^{q-1} (v_i - v_{i+1}) \cdot m_i(T)\right] + \frac{1}{2} \cdot v_q \cdot m_q(T).$$

Proof

Lemma

Suppose S is the outcome of Random Threshold. For every $i = 1, \ldots, q$ we have

$$\mathbb{E}[m_i(S)] \quad \geq \quad \frac{1}{8(\lceil \log k \rceil + 1)} \cdot i.$$

The theorem follows from the lemma:

$$\begin{aligned}
\mathbb{E}[v(S)] &\geq \frac{1}{2} \cdot \left[\sum_{i=1}^{q-1} (v_i - v_{i+1}) \cdot \mathbb{E}[m_i(S)] \right] + \frac{1}{2} \cdot v_q \cdot \mathbb{E}[m_q(S)] \\
&\geq \frac{1}{16(\lceil \log k \rceil + 1)} \left[\sum_{i=1}^{q-1} (v_i - v_{i+1}) \cdot i \right] + \frac{1}{16(\lceil \log k \rceil + 1)} \cdot v_q \cdot q \\
&= \frac{1}{16(\lceil \log k \rceil + 1)} \cdot \sum_{i=1}^{q} v_i \\
&\geq \frac{1}{32(\lceil \log k \rceil + 1)} \cdot v(S^*) \qquad \square
\end{aligned}$$

Proof of Lemma

Proof (Lemma):
We show the lemma inductively for every value of i. The start with $i = 1$ is simple (Exercise). Consider $1 < i \leq q$.

- Suppose a^* is the bidder with highest value.
- We condition on an event E that two assumptions hold simultaneously:
  (1) The highest bidder is in the sample $a^* \in Y$, and
  (2) for τ we choose j such that $v_i \geq v_{a^*}/2^j \geq v_i/2$.

- We can compute S^* with the Kruskal algorithm. Hence, $v_1 = v_{a^*}$ and $v_q \geq v_1/k \geq v_{a^*}/2^{\lceil \log k \rceil}$. This implies there is a suitable j to fulfill assumption (2) in every case $1 < i \leq q$.

- The algorithm picks the suitable j with prob. $1/(\lceil \log k \rceil + 1)$.
- Overall, the prob. of event E is $\Pr[E] = 1/(2(\lceil \log k \rceil + 1))$.

Proof of Lemma

Let us bound the competitive ratio conditioned on event E.

- The highest i bidders from S^* are a cycle-free set $S' = \{1, \ldots, i\}$. With assumption (2) all values $v_1 \geq \ldots \geq v_i \geq \tau = v_{a^*}/2^j$.
- With assumption (1) we know $a^* = 1$ in Y. Hence, in expectation at least $(i-1)/2 \geq i/4$ bidders of S' are not in Y and can be chosen by our algorithm. This implies $\mathbb{E}[|S' \setminus Y| \mid E] \;\geq\; i/4$.
- In this case, because of the exchange property, the algorithm will pick at least $S' \setminus Y$ many bidders. The expected size of the output S conditioned on event E is

$$\mathbb{E}[|S| \mid E] \;\geq\; \mathbb{E}[|S' \setminus Y| \mid E] \;\geq\; i/4.$$

- Since $\tau \geq v_i/2$ and every chosen bidder has a value at least τ:

$$\mathbb{E}[m_i(S) \mid E] \;=\; \mathbb{E}[|S| \mid E] \;\geq\; i/4.$$

Finally, the conditioning on event E can be removed by multiplying with $\Pr[E]$.

□

Matroids – Definition

Algorithm and analysis can be applied directly to arbitrary matroid auctions.

Definition (Matroid)

A tuple $M = (I, A)$ is a matroid if $I = \{1, \ldots, n\}$ is a finite set of bidders and A a non-empty family of subsets of I such that:

- If $S \in A$ and $T \subseteq S$, then $T \in A$ as well, and
- if $S, T \in A$ and $|T| < |S|$, then there is an $s \in S \backslash T$ mit $T \cup \{s\} \in A$.

Notation

- A set $S \in A$ is called independent.
- A maximal independent set $B \in A$ is called basis.
- The cardinality of every basis is the same and is called rank $rk(M)$ of the matroid. Notation: $k = rk(M)$.

Matroids – Weights and Exchange Property

Definition (Weighted Matroid)

- A matroid with weights $v_i \in \mathbb{R}$ for every $i \in I$ is called weighted.
- The weight of an independent set S is $v(S) = \sum_{i \in S} v_i$.
- An optimal basis is a basis with maximal weight.

Kruskal greedy algorithm: Computes an optimal basis with maximum total weight

Exchange property: B basis, B^* optimal basis. There are pairs $(b_1, b_1^*), \ldots, (b_k, b_k^*) \in B \times B^*$ such that for all $1 \leq t \leq k$

- $B_t = (B \cup \{b_1^*, \ldots, b_t^*\}) \setminus \{b_1, \ldots, b_t\}$ is basis
- $B^* = B_k$
- $v(B_t) \geq v(B_{t-1})$

Is there a constant-competitive algorithm?

Membership oracle for matroid (I, A):
Given any set $S \subseteq I$ decides whether $S \in A$ or not.

Theorem (Babaioff, Immorlica, Kleinberg 2007)

The Random-Threshold mechanism

- *is incentive compatible,*
- *runs in polynomial time with a poly-time membership oracle,*
- *is $O(\log k)$-competitive for arbitrary matroid auctions.*

Is this a good result? Can we obtain constant-competitive algorithms? Can they be augmented by payments to yield IC mechanisms?

Is there a constant-competitive algorithm?

Matroid Secretary Conjecture

For every matroid there is an algorithm in the random-order model that is ...

Weak: ... constant-competitive

Strong: ... e-competitive

Over the last decade, the conjecture was proven for many subclasses of matroids. In general, however, the conjecture remains open. The current best algorithms are due to Lachish (2014) as well as Feldman, Svensson, Zenklusen (2018). They obtain a competitive ratio of $O(\log \log k)$, but they do not necessarily extend to incentive-compatible mechanisms.

For graphic matroids we prove the weak conjecture.

Constant-Competitive Mechanism for Forest Auctions

Parallel Secretaries

1. Fix an arbitrary order $k_1, k_2, k_3, \ldots$ of *vertices of the graph*
2. Choose $X \in \{0, 1\}$ uniformly at random
3. **if** $X = 1$ **then** orient every edge $e \in E$ to vertex with lower index
4. **else** orient every edge to vertex with higher index
5. For every vertex k_i in parallel:
   Execute Secretary algorithm on the incoming edges and choose at most one incoming edge of k_i. Assign payments as in the Secretary mechanism for single-item auctions.

Theorem (Korula, Pal 2009)

The Parallel Secretaries mechanism is incentive compatible and can be implemented in polynomial time. It is $2e$-competitive for forest auctions.

Proof

Proof:
Incentive compatibility can be shown similarly as for the Secretary mechanism (Exercise). Polynomial running time is straightforward. We only show the competitive ratio.

- Depending on the orientation of edges, the (direct or inverse) order of vertices becomes a topological order, i.e., G becomes a directed acyclic graph. Even if every vertex adds an arbitrary incoming edge to S, then S stays cycle-free.
- We only have to bound $\mathbb{E}[v(S)]$.
- Suppose G_X is the directed graph for $X \in \{0, 1\}$.
- Suppose $h_X(k_i)$ is an incoming edge of k_i in G_X with highest value
- Suppose $S_X = \{h_X(k_i) \mid k_i \in K\}$, and S^* is an optimal forest in G.

Proposition

$$v(S^*) \quad \leq \quad \sum_{k_i \in K} v_{h_0(k_i)} + v_{h_1(k_i)} \quad = \quad v(S_0) + v(S_1) \;.$$

Proof

Conditioned on the choice of X the algorithm obtains for every vertex k_i in expectation at least a $1/e$-fraction of the highest value of any incoming edge of k_i. Thus, we obtain for $x = 0$ as well as $x = 1$

$$\mathbb{E}[v(S) \mid X = x] \quad \geq \quad 1/e \cdot v(S_x) \ .$$

Using the proposition we see

$$\begin{aligned}\mathbb{E}[v(S)] &= \frac{1}{2} \cdot (\mathbb{E}[v(S) \mid X = 0] + \mathbb{E}[v(S) \mid X = 1]) \\ &\geq \frac{1}{2e} \cdot (v(S_0) + v(S_1)) \\ &\geq \frac{1}{2e} \cdot v(S^*) \ .\end{aligned}$$

□

Online Auctions

Secretaries and Matching

Matroid Secretary

Prophet Inequalities

Bidders with Distributions

Here we draw the private value v_i for every bidder i independently from a known distribution $\mathcal{V}_i$ with non-negative values and finite expectation.

Distribution Model

- Values of bidders and arrival order are unknown
- Number n of bidders is known
- Value v_i drawn independently from known distribution $\mathcal{V}_i$ for bidder i

We first study the single-item auction and strive to maximize social welfare. Afterwards we also consider revenue maximization.

Prophet

1. $p_i(v) \leftarrow 0$ for all bidders $i \in I$
2. $\tau \leftarrow \mathbb{E}_{x_j \sim \mathcal{V}_j}[\max_j x_j]/2 = \mathbb{E}[v_{\max}]/2$
3. In round $t = 1, 2, \dots, n$:
4.   **if** $v_{i_t} \geq \tau$ and item unassigned **then**
5.   Assign item to bidder i_t and set $p_{i_t}(v) \leftarrow \tau$.

2-Approximation to the Prophet

Theorem (Krengel, Sucheston 1978)

The Prophet mechanism is incentive compatible and can be implemented in polynomial time. It is 2-competitive for the single-item auction in the distribution model.

Proof:
IC follows similarly as for the Secretary mechanism (Exercise). Polynomial running time is straightforward. We prove the competitive ratio.

The algorithm assigns the item in the first round t^* in which a bidder with value $v_{i_{t^*}} \geq \tau$ arrives. If no such bidder exists, then the item is never assigned and we define $v_{i_{t^*}} = 0$.
We will see that the expected value of our algorithm is $\mathbb{E}[v_{i_{t^*}}] \geq \tau$. Since all values are drawn from distributions, the optimum $v_{\max}$ is also a random variable. Hence, to bound the competitive ratio we will show:

$$\mathbb{E}[v_{i_{t^*}}] \geq \tau = \frac{1}{2} \cdot \mathbb{E}_{x_j \sim \mathcal{V}_j}[\max_j x_j] = \frac{1}{2} \cdot \mathbb{E}[v_{\max}].$$

Analysis

- Let $q(x) = \Pr[v_{\max} \geq x]$ the prob. that at least one bidder has value at least x.
- If no bidder has value at least τ, then the item is never assigned. This happens with prob. $(1 - q(\tau))$.
- Hence, for every bidder i_t: With prob. at least $(1 - q(\tau))$ the item is still available in round t.
- With prob. $q_t(x) = \Pr[v_{i_t} > x]$ bidder i_t has value $v_{i_t} > x \geq \tau$. Then the algorithm assigns the item to i_t.
- Consider the event $E_t(x)$ that (1) the algorithm assigns the item to i_t and (2) his value $v_{i_t} > x$. It holds

$$\Pr[E_t(x)] \quad \geq \quad (1 - q(\tau)) \cdot q_t(x)$$

.

Analysis

- For every $x \geq \tau$ the item can be assigned to at most one bidder with value at least x. Thus, at most one of the events $E_t(x)$ can occur. This implies:

$$\Pr[v_{i_{t^*}} > x] \quad = \quad \sum_{t=1}^{n} \Pr[E_t(x)] \quad \geq \quad (1 - q(\tau)) \cdot \sum_{t=1}^{n} q_t(x) \ .$$

- Several bidders can simultaneously have a value of at least x. Hence, for the prob. that *at least one bidder* has value at least x we see: (Union Bound)

$$\Pr[v_{\max} > x] \quad \leq \quad \sum_{i=1}^{n} q_i(x) \ .$$

Thus, for $x \geq \tau$ there is a relation between the choice of our algorithm and the optimum:

$$\Pr[v_{i_{t^*}} > x] \quad \geq \quad (1 - q(\tau)) \cdot \Pr[v_{\max} > x] \ .$$

Analysis

- For the expectation we use the integral definition for distributions over non-negative values:

$$\begin{aligned}
\mathbb{E}[v_{i_{t^*}}] &= \int_{x=0}^{\infty} \Pr[v_{i_{t^*}} > x] \, dx \\
&= q(\tau) \cdot \tau + \int_{x=\tau}^{\infty} \Pr[v_{i_{t^*}} > x] \, dx \\
&\geq q(\tau) \cdot \tau + (1 - q(\tau)) \int_{x=\tau}^{\infty} \Pr[v_{\max} > x] \, dx \\
&= q(\tau) \cdot \tau + (1 - q(\tau)) \left(\mathbb{E}[v_{\max}] - \int_{x=0}^{\tau} \Pr[v_{\max} > x] \, dx \right) \\
&\geq q(\tau) \cdot \tau + (1 - q(\tau)) \left(\mathbb{E}[v_{\max}] - \tau \right)
\end{aligned}$$

- The natural trade-off becomes obvious: For larger τ, we have higher value when assigning the item, but we have less probability $q(\tau)$ that there is anyone with value at least τ for the item.
- For $\tau = \mathbb{E}[v_{\max}]/2$ the $q(\tau)$-terms cancel, and we obtain $\mathbb{E}[v_{i_{t^*}}] \geq 1/2 \cdot \mathbb{E}[v_{\max}]$ as desired.

Prophets

The inequality

$$\mathbb{E}[v_{i_{t^*}}] \quad \geq \quad \frac{1}{2} \cdot \mathbb{E}[v_{\max}]$$

is called prophet inequality, because it bounds the value in comparison to a *prophet* who knows all values in advance and picks the best bidder.

If we apply the algorithm with **virtual values** instead of original values, we obtain a prophet inequality for the **expected virtual welfare**. This yields a 2-approximation to the **expected revenue of the optimal incentive-compatible (offline) auction**.

Revenue Maximization

Revenue Prophet

1. $p_i(v) \leftarrow 0$ for all bidders $i \in I$
2. We use the shorthand $(x)^+ = \max(0, x)$
3. $\tau \leftarrow \frac{1}{2} \cdot \mathbb{E}_{v \sim \mathcal{V}}[\max_i(\varphi_i(v_i))^+]$
4. In round $t = 1, 2, \ldots, n$:
5.   **if** $\varphi_{i_t}(v_{i_t}) \geq \tau$ and item unassigned **then**
6.   Assign item to bidder i_t and set $p_{i_t}(v) \leftarrow \varphi_{i_t}^{-1}(\tau)$.

Theorem

The Revenue Prophet mechanism is incentive compatible for regular distributions and can be implemented in polynomial time. It is 2-competitive w.r.t. the optimal single-item auction.

Proof:

IC follows as above, since for regular distributions $\varphi_t(v_t)$ is monotone in v_i. The analysis of the ratio above can be transfered with minor changes to virtual values and virtual welfare.

Revenue Maximization

This implies

$$\mathbb{E}[\varphi_{i_{t^*}}(v_{i_{t^*}})] \quad \geq \quad \frac{1}{2} \cdot \mathbb{E}[\max_i(\varphi_i(v_i))^+] \ .$$

Note: $\mathbb{E}[\max_i(\varphi_i(v_i))^+]$ is the virtual value obtained if the item is always given to the bidder with highest non-negative virtual value. This is the **virtual value of the optimal (offline) auction**. Expected payments are expected virtual value, and thus

$$\mathbb{E}\left[p_{i_t^*}(v)\right] \quad \geq \quad \frac{1}{2} \cdot \mathbb{E}\left[\sum_i p_i^*(v)\right] \ ,$$

where p_i^* are the payments of the optimal (offline) auction. □

For **revenue maximization in the offline case** this represents a **simple IC mechanism**: Compute τ and offer the item sequentially to bidders in arbitrary order at a fixed price of τ, until a bidder decides to buy. In this way, we obtain **at least half of the expected revenue** of every other IC mechanism.

Matroid Auctions

We can extend the scenario to matroid auctions in the natural way::

- Known: Set I of bidders, set A of outcomes, matroid $M = (I, A)$
- Bidder $i \in I$ has unknown private value $v_i \sim \mathcal{V}_i$, distribution $\mathcal{V}_i$ known
- Bidders arrive in unknown order. Online mechanism decides directly: accept/reject and payments. Decisions are final
- Goal: Construct an independent set $S \in A$ with highest total value. Mechanism shall be IC.

Here we analyze only a class of mechanism that – instead of a global threshold τ – set a deterministic threshols τ_i for every bidder $i \in I$. We define $\tau_i = \infty$ if $S \cup \{i\} \notin A$. The algorithm accepts a bidder if and only if $v_i \geq \tau_i$. In this case bidder i pays τ_i.

For such an approach we need to make careful choices for the τ_i. In particular, we want the thresholds to (1) avoid accepting too many bidders with small value and (2) avoid rejecting too many bidders with high value. The notion of α-balanced thresholds formalizes these conditions.

α-balanced Thresholds

- For every bidder $i \in I$, let v_i be the private value and v'_i an arbitrary sampled value. Both are drawn independently from $\mathcal{V}_i$.
- The arrival sequence is $\sigma = (a_1, v_{a_1}), \ldots, (a_n, v_{a_n})$. In the following definition of thresholds we fix the arrival sequence (and, hence, all v_i's) and postulate a condition for every such sequence.
- Let $S = S(\sigma)$ be the chosen set of bidders.
- Optimal basis for sampled values v' is B'.
- By the exchange property there is at least one partition of B' in B_c and B_r such that $S \cup B_r$ is a basis of M.
- From all such partitions, let $(B_c(S), B_r(S))$ be the one that maximizes the sampled welfare $v'(B_r(S))$.

Definition

Definition

For any $\alpha > 0$, a mechanism has α-balanced thresholds if for every sequence σ and every X disjoint from $S = S(\sigma)$ with $S \cup X \in A$ it holds that the deterministic thresholds $\tau_i = \tau_i(\sigma)$ satisfy

$$\sum_{i \in S} \tau_i \geq \left(\frac{1}{\alpha}\right) \cdot \mathbb{E}_{v' \sim \mathcal{V}}[v'(B_c(S))] \qquad (1)$$

$$\sum_{i \in X} \tau_i \leq \left(1 - \frac{1}{\alpha}\right) \cdot \mathbb{E}_{v' \sim \mathcal{V}}[v'(B_r(S))] \quad . \qquad (2)$$

Proposition

If a mechanism has α-balanced thresholds, then it is incentive compatible and α-competitive for matroid auctions.

Mechanism for the Matroid Auction

Expected-Margin-Thresholds

1. $p_i(v) \leftarrow 0$ for all $i \in I$, $S \leftarrow \emptyset$
2. In round $t = 1, 2, \ldots n$:
3. **if** $(S \cup \{i_t\}) \notin A$ **then** $\tau_t \leftarrow \infty$; **else**

$$\tau_t \leftarrow \frac{1}{2} \cdot \mathbb{E}_{v' \sim \mathcal{V}}[v'(B_r(S)) - v'(B_r(S \cup \{i_t\}))] \ ,$$

   where all $v'_i \sim \mathcal{V}_i$ drawn independently.
4. **if** $v_{i_t} \geq \tau_t$ **then** $S \leftarrow S \cup \{i_t\}$ and $p_{i_t}(v) \leftarrow \tau_t$

Theorem (Kleinberg, Weinberg 2019)

Expected-Margin-Thresholds has 2-balanced thresholds, is incentive compatible and 2-competitive for matroid auctions.

Discussion

- The distribution model seems to have "nicer" properties than the random-order model, e.g., a better ratio for the single-item auction (2 vs. e).
- For matroid auctions the improvement is substantial (2 vs. $o(1)$). A final evaluation, however, depends on the resolution of the matroid secretary conjecture.
- In recent years, algorithms in the random-order and the distribution model have been proposed for a number of additional online packing problems.
- There exist approaches for online variants of knapsack, matching, independent set, packing integer programs and further problems.
- The properties of algorithms for stochastic online optimization are generally only poorly understood and a topical area of research in theoretical computer science and beyond.

Literature

- M. Babaioff, N. Immorlica, R. Kleinberg. Matroids, Secretary Problems, and Online Mechanisms. SODA 2007.
- N. Korula, M. Pal. Algorithms for Secretary Problems on Graphs and Hypergraphs. ICALP 2009.
- O. Lachish. $O(\log \log \text{rank})$-Competitive Ratio for the Matroid Secretary Problem. FOCS 2014.
- M. Feldman, O. Svensson, R. Zenklusen. A Simple $O(\log \log \text{rank})$-Competitive Algorithm for the Matroid Secretary Problem. Math. Oper. Res. 43(2):638-650, 2018.
- P. Freeman. The Secretary Problem and its Extensions: A Review. Intl. Stat. Rev. 51(2):189-206, 1983.

Literature

- R. Kleinberg, M. Weinberg. Matroid Prophet Inequalities and Applications to Multi-Dimensional Mechanism Design. Games Econom. Behav. 113:97-115, 2019.
- U. Krengel, L. Sucheston. On semiamarts, amarts, and processes with finite value. Adv. in Prob. Related Topics 4:197-266, 1978.