Computing optimal single item auctions by local search

J.A.Kreuzberg

September 4, 2016

Abstract

This paper does research for a local search heuristic approach to compute a revenue maximizing single item auction, called an optimal auction. We try to find the maximal expected revenue of the auctioneer by maximizing the expected payments of the bidders. We show that this problem can be reduced to finding an optimal order of the possible types of the bidders. We use that insight to propose a simple local search algorithm for computing an optimal auction.

1 Introduction

An auctioneer wants to sell one item to a group of bidders where each bidder has private information about their value for the item. There already is an analytical approach of Myerson that computes the maximal expected revenue of the auctioneer for auctions in this setting. In this paper we search for a different approach to compute the maximal expected revenue. This approach appears to be a local search heuristic that has potential to be applied in auction settings where the analytical approach of myerson doesn't work anymore.

2 The optimal auction

2.1 The definition of an auction

To get a clear image how an auction is defined formally, there will first be introduced some basic notation, which will also help clarifying some of the model assumptions.

Assume an auctioneer wants to sell a single item to n bidders. Every bidder has a private valuation for this item, which is also known as a bidder's type. A bidder has a certain type, which is unknown for the other bidders and the auctioneer. Since a bidder's type is private information, the auctioneer doesn't know what type each bidder has. The auctioneer only knows that each bidder has a type space T_j consisting of all possible types bidder j could have: $T_j = \{t_1, \ldots, t_b\}$, where $j = 1, \ldots, n$ and t_b is the maximum type a bidder can have for the item with $b \in \mathbb{N}$. The assumption is made that all bidders are identical so the type space of each bidder is the same. Furthermore a bidders type is drawn from his distribution independently from the other bidders. Since all the n bidders have their own type space, the type space of the auction looks like $T^n = T_1 \cup \cdots \cup T_n = \{t_1, \ldots, t_b\}$. The type distribution vector φ that describes the probability to be a certain type from the type-space T, needs to satisfy $\sum \varphi_i = 1$ and $\varphi_i > 0$ for $i = 1, \ldots, b$. So the input for an auction is a stochastic type vector $(t_1, \ldots, t_n)$ that consists of the types each bidder has. This means that the auctioneer doesn't know what type each bidder has but he does know the distribution where it is drawn from. Besides the bidders, the auctioneer also has a value for the object, denoted as $t_\emptyset$, and the object has to be sold above or equal to this value. Such an auction as just described is known as a single item private value auction.

An auction always consists of two mappings: an allocation rule and a payment rule. The allocation rule $a(t)$ is a function that will determine which bidder will get the object for any given type vector t. This is a mapping from the type space $T^n = \{1, \ldots, b\}$ to the space of allocations $A = \{0, 1\}$:

$$(t_1, \ldots, t_n) \xrightarrow{a} (a_1(t), \ldots, a_n(t))$$

In a single item auction there can only be one good allocated among the bidders, so $\sum a_j(t) \leq 1$. An example of an allocation rule could be that the item will go to the bidder with the highest type or randomly assigned among all bidders with the highest type. This allocation rule is also used in a vickrey auction with reserve price [4]. In this example there are three possible scenario's for the allocation of the item.

$$a(t) = \begin{cases} (0, \ldots, 1, \ldots, 0) & \text{if one bidder has the highest type} \\ \left(0, \ldots, \frac{1}{m}, \ldots, \frac{1}{m}, \ldots, 0\right) & \text{if } m \text{ bidders have the highest type} \\ (0, \ldots, 0) & \text{if } t_i < t_\emptyset \text{ for all } i = 1, \ldots, n \end{cases}$$

As already said, the j-th index of the allocation vector corresponds to bidder j, where $j = 1, \ldots, n$. Since there is assumed that the bidder's type distribution is i.i.d., one can say without loss of generality that it is also possible to look at one bidder only. The allocation rule for a bidder, say bidder j, is:

$$a_j(t) = \begin{cases} 1 & \text{if bidder j has the highest type} \\ \frac{1}{m} & \text{if bidder } j \text{ and } m-1 \text{ other bidders have the highest type} \\ 0 & \text{if there exists a } t_i \text{ such that } t_j < t_i \text{ (where i} = \emptyset \text{ is allowed)} \end{cases}$$

Once the allocation rule has assigned a winner, the corresponding payment has to be made by this bidder. This payment follows from the payment rule. A payment rule also is a function that maps a type vector t to a vector of payments π, where $\pi_j(t) \in \mathbb{R}$ for $j = 1, \ldots, n$:

$$(t_1, \ldots, t_n) \xrightarrow{\pi} (\pi_1(t), \ldots, \pi_n(t))$$

An example of a payment rule is that the payment rule claims that the winning bidder has to pay an amount equal to his type. In that case $\pi(t)$ looks like:

$$\pi(t) = \begin{cases} (0, \ldots, t_j, \ldots, 0) & \text{if bidder j has the highest type} \\ (0, \ldots, 0) & t_i < t_\emptyset \text{ for all } i = 1, \ldots, j, \ldots, n \end{cases}$$

When the situation occurs that m of the n bidders have the highest type, it doesn't matter which bidder eventually wins the auction. This winning bidder has to pay the price of this highest type to the auctioneer. An auction with this allocation and payment rule is also known as a first price auction.

So now it is clear how an auction is defined. The input for an auction is a stochastic type vector t, depending on the reported types of the bidders. Applying the allocation rule and payment rule gives as output the winning bidder and the corresponding payment :

$$(t_1, \ldots, t_n) \xrightarrow{a} (a_1(t), \ldots, a_n(t))$$
$$(t_1, \ldots, t_n) \xrightarrow{\pi} (\pi_1(t), \ldots, \pi_n(t))$$

2.2 The expected allocation and payment

Since the input for an auction is stochastic, we compute the expected revenue of the auctioneer. Because the goal is to maximize this expected revenue, we are going to analyze the expected allocations and expected payments. The expected allocation is defined as the probability that bidder j wins the auction with type $t_i \in \{t_1, \dots, t_b\}$. Due to the stochastic input, we have to take the expectation over all possible types of other bidders and fix the type of bidder j. Therefore the type vector t^{-1} is defined as the types of all other bidders except the type of bidder j: $t^{-1} = (t_1, \dots, t_{j-1}, t_{j+1}, \dots, t_n)$. Herefore the expected allocation can be expressed as follows:

$$p_i = \sum_{t^{-1} \in T^{n-1}} a(t^{-1} | \text{ bidder } j \text{ has type } t_i) \cdot \varphi(t^{-1})$$

Here $a(t^{-1}|t_i)$ is the allocation of the object with input t^{-1} and the type of bidder j fixed to t_i. This allocation can only be made if the corresponding type vector t^{-1} occurs. This probability is given by $\varphi(t^{-1})$. Note that it can occur that the type t_i of bidder j is below the type of the auctioneer, $t_\emptyset$. In that case $p_i = 0$.
With this expected allocation, the probability that a bidder wins the auction can be described more in detail. Recalling the assumption that all bidders are independent and identical, every bidder has the same probability to win the auction. If there are n bidders who participate in the auction, the probability that bidder j would win the auction is lower or equal to $\frac{1}{n}$ because we have to take into account that there is a probabillity that the item won't be sold. Because it is unknown which type bidder j has, one must consider all possible types. This leads to the following:

$$\begin{aligned}
&P\{\text{bidder j wins the auction}\} \\
&= \sum_{i=1}^{b} P\{\text{Bidder j wins the auction} \mid \text{Bidder j has type } t_i\} \cdot P\{\text{Bidder j has type } t_i\} \\
&= \sum_{i=1}^{b} p_i \cdot \varphi_i \leq \frac{1}{n}
\end{aligned}$$

This expression is obtained by using the Law of Bayes to implement the condition that bidder j has type t_i.

The vector p has the following structure: $p = (p_1, \dots, p_b)$. There also is a probability that the type of a bidder is below the type of the auctioneer. The probability that the auctioneer won't sell the item, so the expected allocation of the auctioneer, is denoted as $p_\emptyset$. This probability consists of all probabilities that a bidder with a type below the auctioneer's type will win the auction:

$$p_\emptyset = \sum_{i: t_i < t_\emptyset} p_i$$

This makes sense because the auctioneer won't sell the object if the highest type is below her type.

The expected payment can be derived from the payment rule. The amount a bidder with type t_i expects to pay, Π_i, can be derived almost the same way as the expected allocation:

$$\Pi_i = \sum_{t^{-1} \in T^{n-1}} \pi(t^{-1} | \text{bidder } j \text{ has type } t_i) \cdot \varphi(t^{-1})$$

The expected payment of a bidder with type t_i is the payment he should pay according to the payment rule under the condition that this bidder has type t_i and the types of the other bidders are described by t^{-1}. Again the corresponding type vector t^{-1} has to occur.

2.3 Assumptions for the auction

In the auctions that are analyzed, two assumptions must hold. To derive these assumptions we introduce a new variable. The expected valuation of a bidder with true type t_i and reporting type

t_j , is denoted as $\nu(p_j|t_i)$. The expression for the expected valuation is as follows:

$$\nu(p_j|t_i) = t_i \cdot p_j \ \forall i, j \in \{1, \dots, b\}$$

The first assumption is that a bidder must obtain a non-negative surplus from participating in the auction. This means that the payment the auctioneer charges to a bidder is bounded by the expected valuation of a bidder.

$$\Pi_i \leq \nu(p_i|t_i) \ \forall i \in \{1, \dots, b\}$$

When this assumption doesn't hold, the auctioneer could charge infinite payments from the bidders.

The second assumption that must hold is that bidding truthful gives a higher surplus as bidding non-truthful. In this way bidders won't report another type as their true type. It could be that a choosen allocation and payment rule gives the incentive to report different than your true type. In 1981, the american economist named Roger B. Myerson, proved that for every allocation and payment rule that form an auction, there is a corresponding auction where these two assumptions hold. For this reason, without loss of generality, the assumption can be made that bidders are truthful in reporting their type when searching for the optimal auction. The incentive to bid your type truthfully is also known as the Lowercase Revelation Principle [3]:

$$\nu(p_i|t_i) - \Pi_i \geq \nu(p_j|t_i) - \Pi_j \ \forall t_i, t_j \in T$$

3 Mathematical foundations for the optimal auction

In section 2 the auction is formally described. This was neccessary to obtain a good understanding of the assumptions for the model. Now we can focus on optimizing the expected revenue of the auctioneer under the assumed conditions. The expected revenue of the auctioneer is denoted as $E[R]$. The expected revenue of the auctioneer constists of the the expected payments from all the n bidders. Recalling the argument from section 2 that all bidders are identical and independent from each other, it is sufficient to focus on the expected payment of one bidder only. Since the auctioneer doesn't know the type of a bidder, conditioning on the possible type is neccesary. So the expected payment of a bidder consists of the sum of the expected payments of a bidder with type t_i under the probability that the bidder has type t_i, for $i = 1, \dots, b$.

$$E[R] = \sum_{j=1}^{n} \Pi_j = \sum_{j=1}^{n} \sum_{i=1}^{b} \Pi_j^i \cdot \varphi_i = n \cdot \sum_{i=1}^{b} \Pi_i \cdot \varphi_i$$

To compute the expected payments, the assumption of truth telling will be rewritten. The lowercase revelation principle gives:

$$\nu(p_i|t_i) - \Pi_i \geq \nu(p_j|t_i) - \Pi_j \ \forall t_i, t_j \in T$$

Now take $j = i - 1$.

$$\begin{aligned} &\nu(p_i|t_i) - \Pi_i \geq \nu(p_{i-1}|t_i) - \Pi_{i-1} \\ \iff &\Pi_i - \Pi_{i-1} \leq \nu(p_i|t_i) - \nu(p_{i-1}|t_i) \\ \iff &\Pi_i - \Pi_{i-1} \leq t_i(p_i - p_{i-1}) \end{aligned}$$

Because the auction is seen from the perspective of the auctioneer, it is not difficult to understand that given some value for Π_{i-1}, the auctioneer will charge Π_i as large as possible. This yields

$$\Pi_i - \Pi_{i-1} = t_i(p_i - p_{i-1}) \ \forall i \in \{1, \dots, b\}$$

Because this holds for every type t_i, one can construct a formula for every expected payment of type t_i that only depends on the probabilities p_i. In figure 1 this is illustrated in a directed type graph. This graph consist of b vertices, where each vertex corresponds with a certain type. The weight of going from $t_{i-1} \to t_i$ is $t_i(p_i - p_{i-1})\ \forall i = 1, \dots, b$. For computing the expected payments we add a dummy vertex to the type graph. Hereby it is possible to compute the maximal expected payment from type t_i by finding a shortest path for $0 \to t_i$. Here $\Pi_0 = 0$ and $\nu(p_0|t_i) = 0$ because the dummy vertex is not linked to a type.

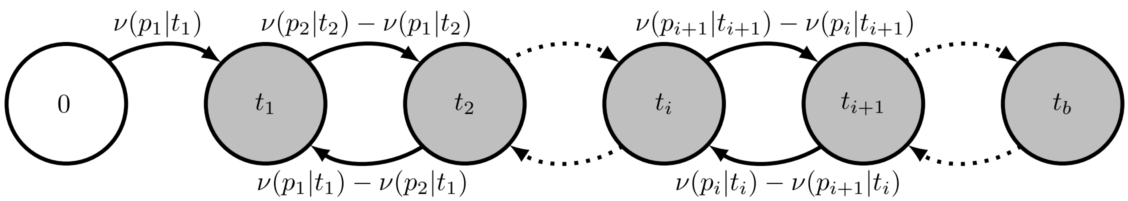

Figure 1: The directed type graph

In section 2 we stated the two assumptions that the surplus of a participating bidder is non-negative and that the bidders report their type truthfully. To make sure these assumptions still hold in the type graph, there may not exist cycles of negative length. Only when there are no cycles of negative length, the shortest path for $0 \to t_i$ corresponds with the expected payment Π_i. Consider the cycle from $t_i \to t_{i+1}$ and back.

$$\begin{aligned}
&[\nu(p_{i+1}|t_{i+1}) - \nu(p_i|t_{i+1})] + [\nu(p_i|t_i) - \nu(p_{i+1}|t_i)] \geq 0 \\
\Longleftrightarrow\ & t_{i+1}(p_{i+1} - p_i) + t_i(p_i - p_{i+1}) \geq 0 \\
\Longleftrightarrow\ & (t_{i+1} - t_i)(p_{i+1} - p_i) \geq 0 \\
\Longleftrightarrow\ & p_{i+1} - p_i \geq 0
\end{aligned}$$

Here we used the monotonicity of the types to derive the last step.
Since this holds for every $i \in \{1, \dots, b\}$, it follows that the expected allocation p is monotone in the types. This condition must hold while maximizing the expected revenue of the auctioneer.

In the directed type graph of figure 1 the edges from $t_i \to t_k, k > i + 1$ are not illustrated, but they do exist. The shortest path for the expected payment Π_i appears to be the directed path from $0 \to t_i$. Using the triangle inequality we can show that a directed path in the type graph of figure 1 is always shorter as skipping a vertex. To show this, take three random neigbour vertices, t_{i-1}, t_i, t_{i+1}, as is illustrated in figure 2.

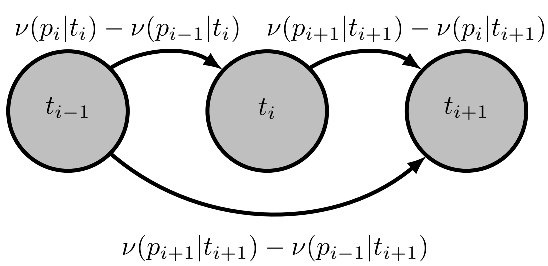

Figure 2: Triangle inequality

In the following derivation we show that the path $t_{i-1} \to t_i \to t_{i+1}$ is shorter as going from

$t_{i-1} \to t_{i+1}$ and skip the vertex t_i :

$$\begin{aligned}
&[\nu(p_i|t_i) - \nu(p_{i-1}|t_i)] + [\nu(p_{i+1}|t_{i+1}) - \nu(p_i|t_{i+1})] \\
=&t_i(p_i - p_{i-1}) + t_{i+1}(p_{i+1} - p_i) \\
=&t_{i+1}p_{i+1} + (t_i - t_{i+1})p_i - t_ip_{i-1} \\
\leq&t_{i+1}p_{i+1} + (t_i - t_{i+1})p_{i-1} - t_ip_{i-1} \\
=&t_{i+1}(p_{i+1} - p_{i-1}) \\
=&\nu(p_{i+1}|t_{i+1}) - \nu(p_{i-1}|t_{i+1})
\end{aligned}$$

In this proof we used the monotonicity of the types and the monotonicity of p such that $(t_i - t_{i+1})p_i \leq (t_i - t_{i+1})p_{i-1}$. So the maximal expected payment Π_i is indeed the directed path $0 \to t_i$ among all the vertices $t_1, \dots, t_{i-1}$. Therefore the maximal Π_i can be expressed as follows:

$$\begin{aligned}
\Pi_i &= \sum_{s=1}^{i} \nu(p_s|t_s) - \nu(p_{s-1}|t_s) \\
&= \sum_{s=1}^{i} t_sp_s - t_sp_{s-1} \\
&= \sum_{s=1}^{i} t_s\,(p_s - p_{s-1}) \\
&= t_ip_i - \sum_{s=1}^{i-1} (t_{s+1} - t_s)\,p_s
\end{aligned}$$

Recalling the formula to compute the expected revenue of the auctioneer:

$$\begin{aligned}
E[R] =&n \cdot \sum_{i=1}^{b} \varphi_i \cdot \Pi_i \\
=&n \cdot \sum_{i=1}^{b} \varphi_i \cdot \left(t_ip_i - \sum_{s=1}^{i-1} (t_{s+1} - t_s)\,p_s \right) \\
=&n \cdot \Big(\varphi_1t_1p_1 + \varphi_2(t_2p_2 - (t_2 - t_1)p_1) + \cdots + \varphi_b[t_bp_b - (t_b - t_{b-1})p_{b-1} - \cdots - (t_2 - t_1)p_1)]\Big) \\
=&n \cdot \Big([t_1\varphi_1 - (t_2 - t_1)\varphi_2 - \cdots - (t_2 - t_1)\varphi_b]p_1 \\
&+ [t_2\varphi_2 - (t_3 - t_2)\varphi_3 - \cdots - (t_3 - t_2)\varphi_b]p_2 + \cdots + [t_{b-1}\varphi_{b-1} - (t_b - t_{b-1})\varphi_b]p_{b-1} + \varphi_bp_b\Big) \\
=&n \cdot \sum_{i=1}^{b} \left(t_i\varphi_i - \left[(t_{i+1} - t_i) \sum_{s=i+1}^{b} \varphi_s \right] \right) p_i \\
=&n \cdot \sum_{i=1}^{b} \left(t_i - \frac{1 - \Phi(i)}{\varphi_i}(t_i - t_{i-1}) \right) \varphi_ip_i \\
=&n \cdot \sum_{i=1}^{b} v(i)\varphi_ip_i
\end{aligned}$$

Here $\Phi(i)$ is the cumulative distribution function such that: $\Phi(i) = \sum\limits_{s=1}^{i} \varphi_s$. The function $v(i)$ is called the virtual valuation according to type t_i:

$$v(i) = t_i - \frac{1 - \Phi(i)}{\varphi_i}(t_i - t_{i-1})$$

For simplicity assume that $v(i)$ is monotone in i. This is an important assumption because now only distibutions φ are allowed such that $v(i)$ is monotone. When $v(i)$ wouldn't be monotone, bidders could get the incentive to report a lower type as their true type to obtain a higher virtual valuation. So to assume the bidders report their type truthfully, $v(i)$ must be monotone [5].

The other part $\varphi_i p_i$ can be seen as the probability that a bidder has type t_i and wins the auction with this type. This can be transformed to the probability that a type t_i wins the auction, which is defined as x_i.
So the final expression for the expected revenue of the auctioneer is as follows:

$$\begin{aligned} E[R] &= n \cdot \sum_{i=1}^{b} v(i)\varphi_i p_i \\ &= n \cdot \sum_{i=1}^{b} v(i) x_i \end{aligned}$$

Since the goal is to find the maximum value for $E[R]$, it is clear to see that the x's with the highest virtual valuation has to be chosen as high as possible. Recalling the assumption $\sum\limits_{i=1}^{b} x_i \leq \frac{1}{n}$, there is a maximum of $\frac{1}{n}$ to be distributed over the vector $(x_1, \ldots, x_b)$.
One could think that the x_i with the highest virtual valuation equals $\frac{1}{n}$. However the composition of vector x, actually the composition of the vector p under the probability φ, has to be feasible. Feasibility means that the representation of the expected allocation can be mapped back to a real auction.
This feasibility can be obtained by using Border's Theorem [1]. Border's Theorem claims that the expected allocation p_i is feasible $\iff$

$$n \cdot \sum_{i \in S} \varphi_i p_i \leq 1 - (\sum_{i \notin S} \varphi_i)^n \ \forall S \subseteq \{1, 2, \ldots, b\}$$

The left hand side represents the expected distributed quantity of the good to the bidders with types in S. The right hand side represents the probability that one or more bidders have a type in S.
Rewriting this formula, we derive the following:

$$\begin{aligned} & n \cdot \sum_{i \in S} \varphi_i p_i \leq 1 - (\sum_{i \notin S} \varphi_i)^n \ \forall S \subseteq \{1, 2, \ldots, b\} \\ \iff & \sum_{i \in S} \varphi_i p_i \leq \frac{1 - (\sum\limits_{i \notin S} \varphi_i)^n}{n} \\ \iff & \sum_{i \in S} x_i \leq \frac{1 - (\sum\limits_{i \notin S} \varphi_i)^n}{n} \end{aligned}$$

By defining $G(S) = \frac{1 - (\sum\limits_{i \notin S} \varphi_i)^n}{n}$, the final condition for feasibility is:

$$\sum_{i \in S} x_i \leq G(S) \ \forall S \subseteq \{1, 2, \ldots, b\}$$

Note that we want to maximize $(x_1, \ldots, x_b)$ so that the expected revenue will be maximal. We know that the virtual valuations are montone in the types so we maximize $(x_1, \ldots, x_b)$ according to their virtual valuation. Since $v(1) \leq v(2) \leq \cdots \leq v(b)$, we maximize $(x_1, \ldots, x_b)$ in the same order such that $x_1 \leq x_2 \leq \cdots \leq x_b$. But this solution of $(x_1, \ldots, x_b)$ needs to be feasible. So the following inequalities must hold:

$$\begin{aligned}
&x_b \leq G(\{b\}) \\
&x_b + x_{b-1} \leq G(\{b, b-1\}) \\
&\dots \\
&x_b + x_{b-1} + \dots + x_1 \leq G(\{b, b-1, \dots, 1\})
\end{aligned}$$

In each condition we can isolate x_i such that:

$$\begin{aligned}
&x_b \leq G(\{b\}) \\
&x_{b-1} \leq G(\{b, b-1\}) - x_b \\
&\dots \\
&x_1 \leq G(\{b, b-1, \dots, 1\}) - x_b - x_{b-1} - \dots - x_2
\end{aligned}$$

Here we can apply the algorithm of Jack Edmonds. With Edmonds' Greedy algorithm we first maximize x_b, after that x_{b-1}, untill we reach x_1. So the optimal feasible solution x can be computed by replacing the inequalities through equalities.

$$\begin{aligned}
&x_b = G(\{b\}) \\
&x_{b-1} = G(\{b, b-1\}) - x_b = G(\{b, b-1\}) - G(\{b\}) \\
&\dots \\
&x_1 = G(\{b, b-1, \dots, 1\}) - x_b - x_{b-1} - \dots - x_2
\end{aligned}$$

Now we know how to obtain the optimal solution $(x_1, \dots, x_b)$ by this greedy algorithm, we need to prove the feasibility of this greedy solution. We do this by using two equivalent properties of the function $G(S)$. It appears that the function $G(S)$ is a submodular set function which is non-decreasing and non-negative. Such a submodular set function G(S) has the following equivalent properties:

1. $G(S \cap R) + G(S \cup R) \leq G(S) + G(R)\ \forall S, R \subseteq \{1, \dots, b\}$
2. $G(S \cup \{l\}) - G(S)$ is non-increasing in S $\forall S \subseteq \{1, \dots, b\}, l \notin S$

Due to the submodularity of $G(S)$, we can prove the feasibility of the solution $(x_1, \dots, x_b)$.

$$\text{To show: } \sum_{i \in S} x_i \leq G(S)\ \forall S \subseteq \{1, 2, \dots, b\}$$

The idea is to use induction to show this holds for every cardinality of the subset S. So the first step is to verify this for $|S| = 1$. Choose an arbitrary $j \in \{1, \dots, b\}$, then

$$x_j := G(\{1, 2, \dots, j\}) - G(\{1, 2, \dots, j-1\}) \leq G(\{j\}) \text{ for } S = \{j\}$$

Here the property of submodularity is used for computing x_j such that $R = \{1, \dots, j-1\}$ and $S = \{j\}$ for $j \in \{1. \dots, b\}$.

Now assume $\sum_{i \in S} x_i \leq G(S)\ \forall S : |S| \leq l$, and prove this also holds for $|S| = l + 1$.

Let $(i_1, \ldots, i_b)$ be an order of the type set $\{1, \ldots, b\}$ such that $i_k \in \{1, \ldots, b\}$ for $k = 1, \ldots, b$

$$\begin{aligned} \text{Let } S &= \{i_1, \ldots, i_{l+1}\} \\ &= S' \cup \{i_1\} \text{ , where type } i_1 \text{ is the minimum of this set.} \end{aligned}$$

$$\begin{aligned} \text{Since } x_{i_1} &= G(\{i_1\} \cap S') + G(\{i_1\} \cup S') - G(S') \\ &= G(S) - G(S') \end{aligned}$$

$$\begin{aligned} \implies \sum_{i \in S} x_i &= \sum_{i \in S'} x_i + x_{i_1} \\ &\leq G(S') + x_{i_1} \\ &\leq G(S') + [G(S) - G(S')] \\ &= G(S) \end{aligned}$$

Now we proved feasibility for the solution x, the maximization problem to obtain the highest expected revenue is reduced to:

$$\begin{aligned} \max_x \ & n \cdot \sum_{i=1}^{b} v(i) x_i \\ \text{s.t. } & \sum_{i \in S} x_i \leq G(S) \ \forall S \subseteq \{1, \ldots, b\} \\ & p_i \leq p_{i+1} \ \forall i \in \{1, \ldots, b\} \end{aligned}$$

It is also possible to formally prove optimality for the solution x. From the primal problem of maximizing the expected revenue, one can construct a dual problem. By complementary slackness it can be shown that x is an optimal solution for both the primal and dual problem so indeed is an optimal solution [2].

4 Local Search

In the previous section we derived the optimal solution by first maximizing x_b, then x_{b-1}, until x_1. This actually is an order of types in which the corresponding x will be maximized.
By applying the Edmonds' Greedy Algorithm, the strategy is to look for which type i_1 the virtual valuation $v(i_1)$ has the highest value and first maximize this x_{i_1}. After that x_{i_2} , for which the virtual valuation $v(i_2)$ has the second highest value, will be maximized. This keeps going untill the type with the lowest priority i_b is reached such that x_{i_b} will be maximized. One can see that there actually has to be found an optimal order $(i_1, \ldots, i_b)$ to optimize the x vector such that $\sum_{i=1}^{b} v(i) x_i$ will be maximal. Since the virtual valuations are monotone in the types, the order of of virtual valuations is the same order of possible types. Therefore the following theorem can be stated:

Theorem 1 *Searching for the auctioneer's revenue maximizing auction can be reduced to finding an optimal order of possible types* $(i_1, \ldots, i_b)$.

It occurs that types have a negative virtual valuation. If this happens, the corresponding x values are set to zero. Those corresponding x and so the expected allocation p have no influence on the expected revenue anymore. Those probabilities belong to the probability that the auctioneer won't sell the item. So we can rewrite the formula for $p_\emptyset$:

$$p_\emptyset = \sum_{i: v(i) < 0} p_i$$

4.1 Finding the optimal order of types

The optimal order $(i_1, \ldots, i_b)$ can be found by a local search heuristic on the possible types. Choose a random type and analyze at which priority in the order it gives the highest expected revenue for the auctioneer. Doing this for all possible types, eventually it will give the optimal order of types. The local search heuristic is illustrated in figure 3.

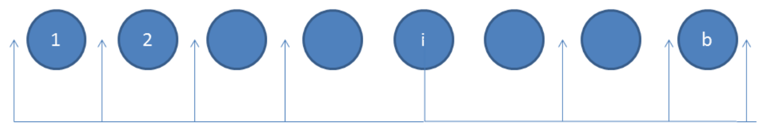

Figure 3: Finding the optimal order of types

Actually finding the optimal order of types comes down to finding the order of virtual valuations. Due to the monotonicity of the virtual valuations we know that the optimal order for auctions this paper looks at, is $(b, b-1, \ldots, 1)$.

5 Computational results

After achieving the result that it is possible to obtain the optimal auction by a local search heuristic theoratically, it is interesting to verify this by computational results. There are two main questions to investigate:

1. Does the local search heuristic give the same expected revenue as the analytical approach of Myerson?
2. What influence has the type distribution on the expected revenue?

Furthermore it is interesting how the exected revenue reacts when the amount of bidders increases.

5.1 Does the local search heuristic give the same expected revenue as the analytical approach of Myerson?

The straight answer to the question is that the local search heuristic indeed gives the same expected revenue as the analytical approach of Myerson. This can be seen by the following example:

Example 1 *Assume the following input* $n = 10$, $T = \{1, 2, \ldots, 14\}$, $\varphi \sim uniform(0, 14)$

Using this input, we can compute the virtual valuations:

$$v(T) = [\text{-12.0000 -10.0000 -8.0000 -6.0000 -4.0000 -2.0000 -0.0000 2.0000 4.0000 6.0000 8.0000 10.0000 12.0000 14.0000}]$$

We can compute each x_i *by computing* $G(S) = \frac{1-(\sum_{i \notin S} \varphi_i)^{10}}{10}$ *for the corresponding* S*. Every type that has a negative virtual valuation, doesn't have a probability to win the auction. Recalling the probability that the auctioneer won't sell the item is* $p_\emptyset = \sum_{i:v(i)<0} p_i$

According to the virtual valuations, the optimal order of types is $(14, 13, \ldots, 1)$ *so the expected allocation* p *under the probability* φ *is*

$$(x_1, x_2, \ldots, x_{14}) = [0\ 0\ 0\ 0\ 0\ 0\ 0\ 0.0003\ 0.0008\ 0.0023\ 0.0055\ 0.0124\ 0.0263\ 0.0523]$$

Therefore the expected allocation can be computed:

$$(p_1, p_2, \ldots, p_{14}) = \textit{[0 0 0 0 0 0 0 0.0038 0.0117 0.0315 0.0771 0.1741 0.3676 0.7328]}$$

Here $p_\emptyset = \sum_{i=1}^{7} p_i = 0.0014$.

From the virtual valuation and the solution x, we can compute the expected revenue of the auctioneer. So when there are 10 participating bidders, each having a type space $\{1, 2, \ldots, 14\}$ *and the type distribution is uniform, the expected revenue of the auctioneer is :*

$$E[R] = 10 \cdot \sum_{i=1}^{14} v(i) x_i = 12.3367$$

Using the same input, the analytical approach of Myerson gives the same output as the greedy algorithm.

5.2 What influence has the type distribution on the expected revenue?

Recalling the assumption that the type distribution ensures montone virtual valuations. So not all type distributions are feasible. When the probability of reporting a type changes, the expected revenue will be influnced by that. In example 1 the type distribution φ is uniform.
So $\varphi_1 = \varphi_2 = \cdots = \varphi_{14} = \frac{1}{14}$
The following example has an exponential type distribution. First we need to check if an exponential distribution gives monotone virtual valuations. In this example it holds that $\varphi_1 \leq \varphi_2 \leq \cdots \leq \varphi_b$, so the virtual valuations $v(i) = t_i - \frac{1-\Phi(i)}{\varphi_i}(t_i - t_{i-1})$ will be monotone in the types. The part $\frac{1-\Phi(i)}{\varphi_i}(t_i - t_{i-1})$ will decrease if the types increases so the virtual valuations $v(i)$ won't decrease in the types t_i.

Example 2 *The input for the auction is* $n = 10$, $T = \{1, 2, \ldots, 14\}$,
$(\varphi_1, \varphi_2, \ldots, \varphi_{14})$ *=(0.0000 0.0000 0.0000 0.0000 0.0001 0.0002 0.0006 0.0016 0.0043 0.0116 0.0315 0.0855 0.2325 0.6321)*

Using this input, we can calculate the virtual valuations:

$$v(T) = 10000 \cdot \textit{(-6.9989 -2.5747 -0.9471 -0.3484 -0.1281 -0.0471 -0.0173}$$
$$\textit{-0.0063 -0.0022 -0.0007 -0.0002 0.0000 0.0001 0.0001)}$$

Here the optimal order of types is also $(14, 13, \ldots, 1)$ *so the expected allocation p under the probability φ is*

$$(x_1, x_2, \ldots, x_{14}) = \textit{[0 0 0 0 0 0 0 0 0 0 0 0.0000 0.0000 0.1000]}$$

Therefore the expected allocation can be computed:

$$(p_1, p_2, \ldots, p_{14}) = \textit{[0 0 0 0 0 0 0 0 0 0 0 0.0000 0.0000 0.1582]}$$

Here $p_\emptyset = \sum_{i=1}^{11} p_i = 2.9731 \cdot 10^{-13} \approx 0$.
So when there are 10 participating bidders, each having the type space $\{1, 2, \ldots, 14\}$ *and the type distribution is exponential, the expected revenue of the auctioneer is :*

$$E[R] = 10 \cdot \sum_{i=1}^{14} v(i) x_i = 13.9998$$

When we compare example 1 and 2, the results are very intuitive. The expected allocations of example 1 are more devided over the types. This makes sense because when a bidder has type 12 he still has a probability to win the auction since the probability that other bidders have a higher type is not that very high due to the uniform type distribution. Comparing this with example 2, we see that bidders with a type below 14 nearly have no probability to win the auction. This also makes sense because due to the exponential type distribution there is a very high probability that another bidder has type 14.
Furthermore we can see that the expected revenue is higher under an exponenital type distribution than a uniform type distribution. This also follows from the type distribution since in example 2 there is a high probability that a bidder has type 14.

It is also interesting to investigate what happens with the expected revenue if the amount of bidders increases. Figure 4 illustrates how the expected revenue is plotted against the amount of bidders.

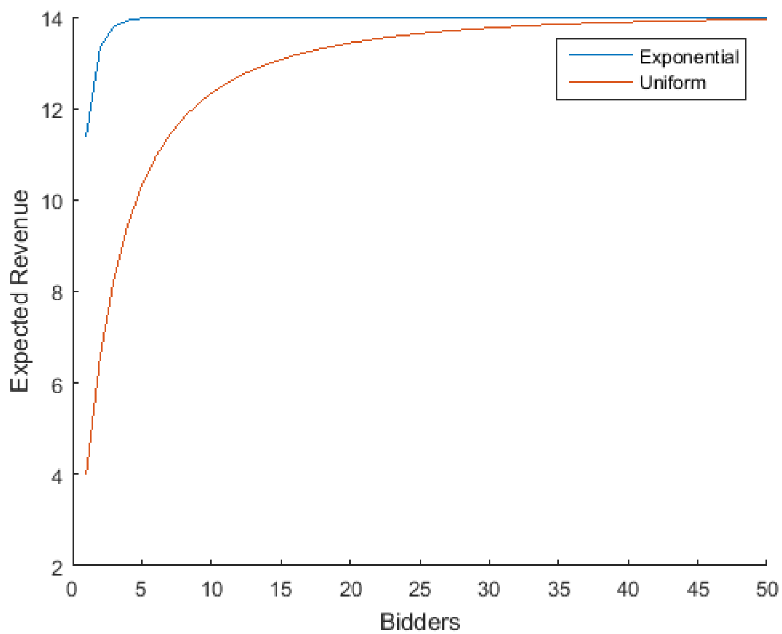

Figure 4: The expected revenue plotted against the amount of bidders. Here $T = \{1, 2, \ldots, 14\}$

6 Conclusion

The main conclusion of this paper is Theorem 1 that is stated in section 4. We are able to reduce the problem of computing the revenue maximizing auction to finding an optimal order of possible types.
We started by deriving a formula for the expected revenue that only depends on the expected payments of the bidders. It appeared that eventually the expected revenue only depends on the type distribution and the probability that a type wins the auction. So by finding an optimal order of possible types, those probabilities can be maximized according to the priority of the type. Among this way we were able to reduce the problem to finding an optimal order of possible types. Since the allowed type distributions ensures monotonicity in the virtual valuations, we already know what the optimal order of types is. But when it appears that the type distribution results in virtual valuations that are not monotone, the optimal order of types follows from the local search algorithm.

7 Future work

In this paper, we look at single item private value auctions. For these auctions we proved that under certain circumstances, the optimal auction can be derived from finding an optimal order of possible types. For these auctions we could also compute the expected revenue by the analytical approach of Myerson. But when the auctions consist of a more complex setting, it is almost never possible to compute the expected revenue by the analytical approach of Myerson. For this reason it is very interesting to investigate if this local search heuristic also gives the optimal auction if the auction consists of a more complex setting, like bidders with multi-dimensional types. To investigate this, the insight of this paper is very useful.

References

[1] Kim C Border. Reduced form auctions revisited. *Economic Theory*, 31(1):167–181, 2007.

[2] Bodo Manthey. Optimization modeling. pages 14–17.

[3] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[4] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[5] Rakesh V Vohra. Optimization and mechanism design. *Mathematical programming*, 134(1):283–303, 2012.