

A Mechanism for Fair Distribution of Resources with Application to Sponsored Search

Evgenia Christoforou^{1,2}, Antonio Fernández Anta¹, and Agustín Santos¹

¹ Institute IMDEA Networks, Madrid, Spain

² Universidad Carlos III, Madrid, Spain

1 Introduction

We have designed a general mechanism for the fair distribution of resources in an online fashion among a set of users. We refer to this mechanism as Fair and Efficient Distribution of Resources (FEDoR). While the FEDoR mechanism is general, and applicable to many scenarios, we apply it and present it in the context of sponsored search; where advertisers pay the search engine to show their content, usually in order to get traffic to their own websites. A large amount of the search engine's income derives from sponsored search. In sponsored search, a number of advertisers are competing for a limited number of slots in each specific keyword search in the search engine. In order to distribute the available slots among the advertisers, search engines are starting to hold keyword auctions.

On the one hand, we have the advertisers (bidders/players), that have to carefully select the search engine, keywords, time frame, and geographical location of the ads to be placed. The goal of the advertisers is to have as much quality traffic to their websites as possible, while staying on budget. On the other hand, we have the search engine (seller), that wants to maximize her revenue without compromising the user experience with too many or irrelevant ads. Thus, the search engine wants to have a mechanism to auction ad spaces (slots) that is easy to understand by the advertisers, and that provides her high revenues. A desirable property of this mechanism is to be incentive compatible (i.e., the best strategy for the advertisers is to bid their true values). Advertisers may value a slot based on different metrics, like the cost per mille impressions (CPM, the cost for the advertisement appearing a thousand time in a specific keyword search), the cost per click (CPC, how much she values a click of a user), or the cost per action (CPA, how much she values a user actually making a transaction on her website). Hence, depending on the goal of the advertiser, an auction mechanism may be more or less appealing to her.

In general terms, an online auction mechanism addresses two questions in its design: (1) how to allocate bidders to auctioned goods (e.g., slots) and (2) what price each bidder pays for the good. The most straightforward type of auction is the generalized first price auction (GFP), where the i th highest bidder gets the i th ‘‘best good’’, and pays the amount of her bid to the seller. In the generalized second price auction (GSP) [4] the i th highest bidder gets the i th best good, but pays to the seller the bid of the $(i + 1)$ st highest bidder. In the Vickrey-Clark-Groves (VCG) [3] auction the i th highest bidder gets the i th best good, but pays the *externality* that she imposes on the other bidders by winning that slot. The auction mechanism of Google [1] is a GSP auction that also takes into account a ‘‘quality score’’ [2]. Part of that quality score is the click through rate (CTR) that a specific advertiser will obtain if it gets a specific slot. In this case, the allocation mechanism and the payment of the bidder that wins a slot is obtained by multiplying her bid by the CTR of that slot. In this work we have modeled the CTR as a weight (w_k where k is the position/slot) of the importance of each slot or good and we assume that is the same for all bidders.

Contributions This work proposes a generic auction mechanism where advertisers express how much they value a slot in a specific keyword search. We assume that the bids of the advertisers are expressed as uniform values between 0 and 1. Hence the issue of the different valuations (i.e., CPM, CPC, CPA) is alleviated.

The goal of our mechanism is to fairly allocate the slots to the bidders over the course of the online sponsored auction for the specific keyword (in a certain time interval). The mechanism uses a flat fee as the payment scheme and allocates the first slot to the bidder with the highest bid, the second slot to the bidder with the second highest bid, etc. Under the constraint of fairness the mechanism is also social efficient, for this reason we call our mechanism *Fair and Efficient Distribution of Resources* (FEDoR).

FEDoR is a mechanism that not only can be easily understood by the bidder (e.g., compared to VCG), it provides also a sense of fairness to the bidders, making them more prone to participate in the sponsored auction. Due to its unique design, the bidder has no incentives to apply any strategic behaviour. From the sellers side

FEDoR is also a good match since its revenue can be adjusted while providing a friendly mechanism to the bidder, giving them incentives to advertise in the specific search engine.

This work proposes a novel approach where distributing the slots among the advertisers is based on how much each advertiser values appearing in a keyword search slot at a specific time. The proposed approach makes this value independent of her true payment to the search engine, which can take the form of a flat fee. For this purpose we have designed a new auction mechanism that fairly distributes resources (or goods, e.g., slots) in online fashion, based on the users' declared preferences, while being socially efficient. This work assumes that instead of one, k non-identical goods are assigned to the n players each time. FEDoR can be used even when the auction is done in a distributed fashion (i.e., without central authority), and it provides fairness, social efficiency and incentive compatibility.

More specifically our contribution is summarized in the following results:

- We analytically show that the FEDoR mechanism guarantees *fairness*, which means that every player will receive the i th best slot, for every i , the same number of times in expectation.
- We prove that the mechanism is *socially efficient*, in the sense that the social utility (the total valuation of the goods assigned) is maximized, subject to the fairness property.
- We show that being honest and announcing the real valuation of the goods at every round maximizes the expected utility (value of goods received) of a player. This means that the mechanism is *truthful*.
- We experimentally compare FEDoR with VCG and GSP, under the assumption that the private values of the players are uniform between 0 and 1. Results show that by appropriately adjusting the flat fee, the mechanism designer can tune the desired utility for the seller and advertiser, opposed to the static approach of VCG and GSP.
- Additionally, experimental results go hand by hand with our analytical results asserting the truthfulness property.

Related Work This work has been inspired by, and forms an extension of, the work of Santos et al. [6]. In that work a fair collaborative system was designed and analyzed. More precisely, in that work a set of players has a common interest in the execution of a set of tasks. Executing a task comes with a different cost for each player. Every time a task has to be executed, players declare their cost for computing the task. The authors assume that selfish players may exist, that try to avoid executing tasks in order to increase their benefit. The proposed mechanism is called QPQ. It uses the concept of *linking mechanisms* [5], where instead of offering incentives or payments to the players to declare their true values, it limits the players responses to a probability distribution known to the mechanism designer. However, unlike the linking mechanisms, QPQ does not assume that the probability distribution of the players costs is known. In addition the presence of non-rational players is allowed.

2 The FEDoR Mechanism.

The FEDoR mechanism works as follows, in every instance of the game (every auction), a player i observes its type θ_i . Then, applying its strategy σ_i , the player chooses a type $\hat{\theta}_i$ that will be reported as the player's bid. These bids are gossiped among all players, so that all of them end with the same vector of reported types $\hat{\theta}$. As mentioned, we will assume the existence of a perfect goodness-of-fit test $GoF_Test(\cdot)$, which is used to determine whether the reported type $\hat{\theta}_x$ is a uniform random sample in $[0, 1]$, for every player x . If not, the reported type is replaced in the vector v of bids by a pseudorandom value uniform in $[0, 1]$ derived from the rest of reported types $\hat{\theta}_{-x}$. Since all players apply the same goodness-of-fit test to the same value $\hat{\theta}_x$, and the same pseudorandom function to the same vector $\hat{\theta}_{-x}$, they all assign the same value to v_x . As a result, the same vector v is obtained by each player, which contains the bids of the players that passed the test and the generated values for those that did not. The distribution of slots (outcome of the mechanism) is then done by choosing the k highest values of vector v , $v_{i_1} \geq v_{i_2} \geq \dots \geq v_{i_k}$, for slots 1 to k .

3 Results

Analytical Results. As mentioned above we analytically show that the FEDoR mechanism guarantees fairness, which means that every player i (honest or not) has the same probability of getting the j th slot for every j . Additionally we have shown that the strategy that maximizes the utility of a player i is to be honest, i.e. truthfulness. Finally we have proved social efficiency, that is if all players are honest, the social utility is maximized.

Experimental Results We have complemented our theoretical results with experiments, in order to visually show the potential of FEDoR. Our results were obtained from 100 experiments of 10000 rounds (auctions) each. The number of participating players was nine and the number of slots three. The weights where $w_1 = 3, w_2 = 2, w_3 = 1$. Honest players are considered the ones that reveal their true valuation for the set of slots. Cheater players are the ones declaring a different valuation from their actual valuation on the set of slots. All players' private valuations follow a uniform distribution.

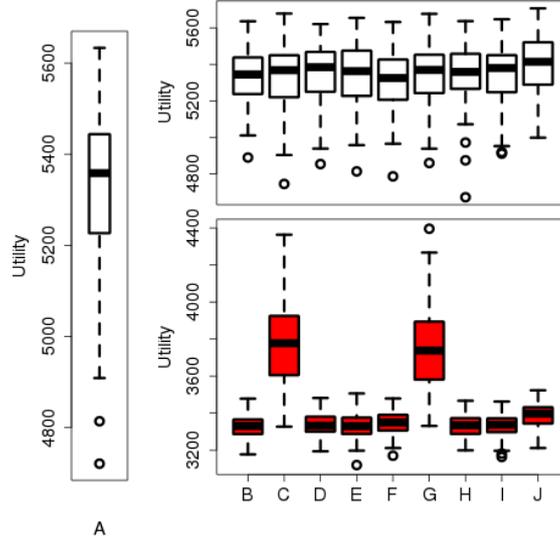


Fig. 1. The utility of an honest player compared to a cheating player in different scenarios using the FEDoR mechanism, for $p.value < 0.1$ and history length 1000. Left graph represents the distribution of the utility of an honest player in a scenario where all 9 players are honest following a uniform. Right top graph represents the utility of an honest player with a uniform distribution in scenarios B-J. Left bottom graph represents the utility of a cheating player with distribution analogous to scenarios B-J. Scenarios are as follows, B:8 players with uniform and represented cheater with normal ($\mu = 0.5, \sigma = 0.15$), C:8 players with uniform and represented cheater with beta distribution with $\beta = 0.9$, D:8 players with uniform and represented cheater with beta distribution with $\beta = 0.7$, E:8 players with uniform and represented cheater with random distribution, F:6 players with uniform 3 with random distribution including the represented cheater, G:6 players with uniform 3 with beta ($\beta = 0.9$) distribution including the represented cheater, H:6 players with uniform 3 with beta ($\beta = 0.7$) distribution including the represented cheater, I:6 players with uniform 3 with normal distribution ($\mu = 0.5, \sigma = 0.15$) including the represented cheater, J:5 players with uniform 1 with random, 1 with beta ($\beta = 0.9$), 1 with beta ($\beta = 0.7$) and represented cheater with normal distribution ($\mu = 0.5, \sigma = 0.15$).

We investigate the utility of the honest and cheaters in a variety of different scenarios, while having the KS test as our GoF test. We have shown analytically (assuming a perfect GoF test) that the best strategy for each player is to be honest. Our experimental results come to assert this even in the case where the GoF is not perfect. From Figure 1, the utility of the honest player is greater than the cheater in all scenarios considered. Independently of the behavior of the rest of the players, the mean utility of the honest player is around the same value. There is no cheating strategy that will give a higher utility. Especially in scenarios C and G where the player cheats with a beta (where parameter is $\beta = 0.9$) indeed the cheaters increase their mean utility by roughly around 500 units but still the mean utility of the honest player is around 1500 units higher. This also proves the efficiency of the KS test, for $p.value < 0.1$ and history length 1000.

We have conducted experiments that compare the utility of the seller and the players produced by the FEDoR mechanism with the classical mechanisms GSP, and VCG with externality. To do so, we have simulated the three mechanisms. The results obtained are presented in Figure 2. From Figure 2(a) we can notice that the utility of the seller and every player are the same for VCG and GSP in the case of one slot. When the number of slots increases, the utility of the seller is greater with GSP than with VCG, while on the other hand the utility of the player is greater in VCG than in GSP. The advantage of FEDoR, as it is shown in Figure 2, is that the flat fee that the players pay to participate provides a tradeoff between seller and player utilities, and allows to choose any point in the lines shown in the figure. From Figure 2(b) we derive the same conclusions. In addition we notice that while the number of players increases, the utility decreases in all three mechanism. On the contrary, the utility of the

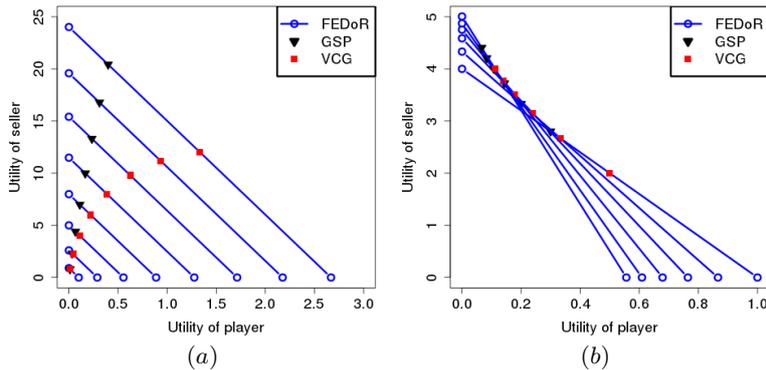


Fig. 2. Comparing FEDoR with VCG and GSP. The values for FEDoR form a line in each experiment, since the utilities of seller and players can be tuned with the flat fee. The values for VCG and GSP always lie on top of the corresponding FEDoR line. Plots represent the mean utility of the seller and one player over 10000 rounds of execution and 100 experiments. (a) Scenarios with 9 player and the number of slots increasing from 1 (leftmost line) to 8 (rightmost line). (b) Scenarios with 3 slots and the players decreasing from 9 (line with largest slope) to 4 (line with smallest slope).

seller increases in all three mechanism when the number of player increases. The results of Figure 2 are performed assuming a perfect GoF test, for comparison reasons with the other mechanisms and following the analysis.

4 Discussion and Future Work

We assume that the bids of the players follow a uniform distribution. Although this might seem as a constrain, if the real bids do not follow such a distribution, they can be transformed by using the PIT transformation, as proposed by Santos et al. [6]. Additionally, having all players paying the same flat fee, while players with a higher budget could claim a larger amount of goods may look like a limitation of the FEDoR. An easy approach to solve this possible constrain is by allowing a player with a higher budget to participate with more than one identities. As a consequence they will be increasing the amount of goods that they will receive proportional to their identities.

As this work proposes an alternative way in viewing auction mechanism, several possible directions appear. One question is what happens if not all players are interested in the set of goods/slots auctioned in a round. With the current mechanism, the players are bidding for receiving a slot, they have no means of indicating their disinterest in certain goods/slots. Another aspect that we are not considering here is that some players might have correlated preferences between certain sets of goods/slots in the same or different auction held.

Finally, an interesting future direction is on the scenario of sponsored search. Through out the paper we have assume that the weight (clickthrough rate) for each slot is the same for all players. In practical applications such as in Adworks of Google [2] each advertiser has a personal clickthrough rate estimation based on the slot and the quality of his advertisement (i.e. according to the relevance between keyword and advertisement). Thus our current mechanism should be adjusted appropriately to deal with this extension. A particularly challenging part of this extension comes with FEDoR being applied in a distributed setting. In that event a scheme must be designed in order for the clickthrough rate of the players to be kept private from the rest of the players.

Acknowledgments: This work has been supported in part by the Comunidad de Madrid grant Cloud4BigData-CM, the MINECO grant TEC2011-29688-C02-01, the National Natural Science Foundation of China grant 61020106002, and European Unions Seventh Framework Programme, CROWD project, grant number 318115.

References

1. Google. "google adwords". www.adwords.google.com/.
2. Google. "how are ads ranked". <https://support.google.com/adwords/answer/1722122?topic=115>, 2005.
3. L. M. Ausubel and P. Milgrom. The lovely but lonely vickrey auction. *Combinatorial auctions*, pages 17–40, 2006.
4. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, National Bureau of Economic Research, 2005.

5. M. O. Jackson and H. F. Sonnenschein. The linking of collective decisions and efficiency. *manuscript, California Institute of Technology*, 2003.
6. A. Santos, A. Fernández Anta, and L. López Fernández. Quid Pro Quo: a mechanism for fair collaboration in networked systems. *PloS one*, 8(9):e66575, 2013.