2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology

Ad Exchange: Intention Driven Auction Mechanisms for Mediating Between Publishers and Advertisers

Rina Azoulay
Dept. of Computer Science
Jerusalem College of Technology
Jerusalem, Israel
rrinaa@gmail.com

Esther David
Dept. of Computer Science
Ashkelon Academic College
Ashkelon, Israel
astrdod@acad.ash-college.ac.il

***Abstract*—This paper concerns a scenario where sellers (i.e., publishers), who are willing to dedicate space on their Website to ads, and buyers (i.e., advertisers), are brought to a common, automatic marketplace. State of the art mechanisms exist for this ad-exchange scenario, however none of the previously proposed solutions fully take into account the preferences of the publishers. In this paper, we developed solutions for the case of multiple advertisers and multiple publishers, while considering the publishers' preferences. We propose three truthful mechanisms: (i) the Hungarian VCG, (ii) the Simultaneous English Auction, and (iii) the Distributed Relocation Protocol. Each mechanism includes an allocation rule and a payment scheme. The Hungarian VCG achieves the optimal allocation, but it is not budget balanced, and it causes a deficit on the part of the ad exchange auctioneer. The other two mechanisms are heuristics, budget balanced, decentralized, avoid manipulations on the advertisers' side, and simulations show that they reach near optimal solutions.**

Keywords- Agent-Based Marketplaces; Auction Markets; Ad exchange

I. Introduction

This paper concerns a scenario where sellers (i.e., publishers), who are willing to dedicate space on their Website for to ads, and buyers (i.e., advertisers), are brought to a common, automatic marketplace. Exchanges in this scenario are advantageous in that they efficiently aggregate information and result in efficient ad allocation and pricing [1]. Examples of such ad exchanges include RightMedia [2], adBrite [3], openX [4] and Google's DoubleClick [5].

In most situations in the real world of ad exchange market applications, the choice of advertisements displayed on a certain private Website is made without taking into consideration the ranking preferences of the publishers. In particular, the tool currently available from Google [6] for publishers to address the above concern (verified on May 1st, 2014), is the ability to specify a list of forbidden keywords and a list of forbidden URLs; nonetheless, there is no option to specify ranked preferences regarding the ad content, despite the fact that it may be important for it to suit the Website owner's preferences, as demonstrated in the bakery example below.

However, when considering private sellers of ad spaces, the Website owners (the sellers) themselves may have preferences considering the ads showed in their ad spaces. Consider, for example, a bakery Website. The bakery will have positive a incentive to show ads on complementary products, such as spreads; it will be neutral concerning ads about bakery accessories, but it will be completely unsatisfied when ads of another bakery appear on its site. The keyword bread, however, may appear in the three types of ads. Nevertheless, the keyword bakery may appear in the two last types of adds. Consequently, an algorithm that matches ads to Websites according to matching keywords may disregard the preferences of the Website owners. To summarize, a solution is required for ad exchange markets where the sellers of ads have detailed preferences considering the ads shown in their ad spaces.

Even when considering a non-profit organization, may be easily convinced to put on its Website ads or links to sites related to its field of activities but that do not clash with its purposes and goals. For example, a site for home education would not want to place ads of traditional schools on its sites, even though the keywords may match.

Most of the existing research [7,8,9] on Web advertising auctions directly focusses on sponsored search and does not directly consider the issue of the self-interested Website advertisement market. Other research was done on private sellers of webspaces [10] without considering their private preferences about the allocated ads.

The problem of suiting advertiser incentives was already addressed by [1] as an open question, and in our previous study [11,12] which dealt with the case of a single Website owner who has to choose which ads to show on its Website, and in which slots. There we proposed three mechanisms for ad placement on publishers' websites: (i) the laddered auction, (2) the greedy-laddered auction, and (3) the Hungarian [13,14] + VCG protocol [15].

However, the solutions proposed in [11,12] are not appropriate for the case of multiple and possibly manipulative ad sellers (Website owners). This is due to the fact that the Laddered and Greedy-Laddered schemes used in [11,12] are inadequate when there are manipulations on the part of the publisher side. For example, the publishers themselves may deviate from their true scoring/ranking function in order to achieve higher gains.

In addition, the Hungarian+VCG protocol is not budget balanced, since it pays the publishers more money than collected from the advertisers, as shown in our simulations

978-1-4673-9618-9/15 $31.00 © 2015 IEEE
DOI 10.1109/WI-IAT.2015.35

in Section IV. Consequently, it cannot be considered a practical solution for the multiple Website problem.

To summarize, new solutions are required to handle the case of multiple publishers and multiple advertisers, in order to guarantee all the required properties that need to be considered in the ad exchange market: incentive compatibility, efficient allocation, computational efficiency, alignment with the publisher's preferences and budget balancing.

In this paper we consider the above problem in which we propose protocols that are truthful for both the publishers and the ads. In particular, we propose two heuristic solutions for the ad-exchange auction mechanisms, in addition to the **HVCG** un-balanced protocol.

1. **SEA:** Simultaneous English Auction, where each publisher publicizes a scoring function and then simultaneous English auctions are run by the different publishers, where each advertiser can decide in which auction to participate and what price to bid.
2. **DRP:** Distributed Reallocation Protocol where the mapping of ads to publishers is randomly initialized, then iteratively, pairs of publishers are chosen randomly, and they exchange ads between themselves in a way that is effective for both, if possible.

For both alternatives we were able to maintain a critical economic property of being truthful/incentive compatible. Simulations show that the efficiency of the allocation by SEA and DRP are close to optimal. Furthermore, the publishers' total share with SEA and DRP are significantly higher than with HVCG.

To conclude, we propose two near optimal mechanisms that achieve high efficiency (incentivize participants to take part for the long term), low complexity, incentive compatibility and a balanced budget for the ad exchange marketplace. Furthermore, these mechanisms are in favor of the publishers, utility-wise and therefore incentivize the publishers to use them in practice.

II. THE MULTI PUBLISHERS MULTI ADVERTISERS MODEL

We consider a model where M publishers (sellers), each denoted Pub_j, offer space on their websites for advertisements. On the other hand N Advertisers (buyers), denoted A_i, compete for these spaces. Moreover, for simplicity, we assume each publisher dedicates K ad slots on his website.

The ad exchange environment is characterized by:

- **Utility per click awarding publisher *Pubj* by being located in slot k** – denoted $u_{i,j,k}$. The utility a certain ad awards different publishers is not the same, and depends on: (i) the matching that exists between the ad's properties and keywords and the Publisher's preferred or unwanted list of keywords, (ii) the quality of the ad, (iii) the aim and type of the Websites the ad links to (i.e., competitive or friendly Websites), and (iv) other business and cultural factors. Note that the value of $u_{i,j,k}$ may be positive or negative. This in turn means that a winning ad that is beneficial to a certain publisher is expressed by a positive utility (for being a complementary product or service). As a result the publisher is willing to compensate the advertiser for this as we discuss later. On the other hand, a winning ad that may harm the publisher if it appears on its Website (e.g., a competitive product or service) will result in a negative utility $u_{i,j,k}$ for the publisher which in turn will be expressed by an extra payment for each click on such an ad. Namely, a negative value of $u_{i,j,k}$ means that Pub_j prefers not to show ad A_i. We assume these utilities are known to the auction center (auctioneer).
- **Click Through Rate (CTR)** - each advertiser A_i is characterized by the $CTR_{i,k,j}$ indicating the click through rate it will attract by being placed in slot k of publisher j. Specifically, the click through rate denotes the expected number of times the ad of advertiser i will be clicked when placed in slot number k out of 100 times that the ad is shown in this particular slot location [16,17] (thus, $0 \le CTR_{i,k,j} \le 100$ for each advertiser i and slot j).
- **Bid:** Each advertiser submits a positive bid $B_{i,j}$ indicating the price it is willing to pay per click on its ad. The bid of the advertiser depends on the advertiser's private knowledge, such as the available budget, keywords' similarity to the Website's content, and the expected utility that will be gained due to future deals performed by visitors who clicked on the ad.
- **Private Value** –defines the value per click advertiser A_i is expecting to gain by being placed on the Webpage of Pub_j, indicated by $V_{i,j}$.

To this end, we can define the total expected utility of pub_j, from showing the set of wining ads $\{A_{i_1},...A_{i_k}\}$ located in slots $\{S_1,...S_k\}$, correspondingly.

$$EU_{pub_j}\left(A_{i_1},...A_{i_k}\right)=\sum_{k=1}^{K}\left(p_{i_k,k,j}+U_{i_k,k,j}\right)\cdot CTR_{i_k,k,j}$$

Where $p_{i_k,k,j}$ is the price that A_{i_k} will be charged for being located in slot S_k of Pub_j.

On the other hand, the utility of advertiser A_i from being displayed in slot S_k of of Pub_j is:

$$EU_{A_i}=\left(V_{i,j}-p_{i,k,j}\right)\cdot CTR_{i,k,j}$$

A. The Auction Design Goals

We propose multiple auction mechanisms to perform the matching between the advertisers and the publishers and to determine the payment rule by which the advertisers will be charged. Our goal is to satisfy the following good economic properties that current auction mechanisms barely even partially satisfy: **Incentive compatible** (to ensure truthful bids); achieve an **efficient allocation**, **computationally efficient** allocation **aligned with the publisher's preferences** and **budget balanced**.

III. Auction Mechanisms for Mediating Publishers and Advertisers

In this section we describe three auction mechanisms for the task of matching advertisers to publishers. For each mechanism, we first briefly describe the auction mechanism and then provide full details. After presenting all the protocols we discuss and compare the advantages and disadvantages of each proposed solution.

HVCG: The Hungarian + VCG Protocol:

The first proposed mechanism, termed the **HVCG**, is basically centralized and is assumed to be run by a center which requires active intervention of the auctioneer. According to this protocol, each publisher reveals its scoring function (by which the auctioneer can calculate the utility $U_{i,k,j}$ of each advertiser A_i located in slot S_k on the Webpage of publisher Pub_j). On the other hand, each advertiser reveals its reservation price associated with each of the publishers $V_{i,j}$. Then, the auctioneer finds the optimal allocation of ads to the publisher using the Hungarian mapping algorithm. Finally, the payments are determined by the auctioneer using the VCG protocol, in order to ensure truth telling of both the Publishers (considering their scoring functions) and the advertisers (considering their reservation prices).

Specifically, the **HVCG** works as follows:

1) Each participant sends its preferences to the auctioneer. Specifically, each publisher Pub_j , provides the auctioneer with $u_{i,j,k}$ for each advertiser-slot combination. In addition each advertiser provides the auctioneer with its private value $V_{i,j}$.

2) The auctioneer finds the optimal mapping of ads to publishers using the Hungarian algorithm.

3) The auctioneer determines the payment of each ad and the income of each publisher by means of the VCG protocol.

The advantage of the VCGH protocol is its stability and the optimality of allocation. Lemma 1 proves the fact that each participant (advertiser and Publisher) will be motivated to report its true preferences when using the VCG protocol. For space reasons proofs are omitted and we refer the reader to [18] for the full proofs.

Lemma 1: The HVCG protocol is truthful.

Though HVCG has the highly desired economic properties of being truthful (strategy proofness) and achieves the optimal efficient allocation, it has two drawbacks: (i) it has a relatively high complexity (determined in Lemma 3), (ii) it is not budget balanced (again inherited by the VCG) as we demonstrate in section 5.

Lemma 2: The allocation made by the HVCG is optimal.

Lemma 3: The time complexity of the HVCG protocol is $O((N+M)\max(N,NK)^3)$.

However, the disadvantage of the HVCG is the fact that it is not budget balanced, as demonstrated in Section IV.

SEA: Simultaneous English Auctions

According to our proposed simultaneous English auctions, the centrality level dramatically decreases relatively to HVCG's centrality, where the center is only responsible for the verification that the published scoring functions are actually being used by the publishers.

Specifically, the **simultaneous English auction** proceeds as follows:

1) Each publisher initiates an English auction, and publicizes its scoring function.

2) Each advertiser (ad owner) sees the list of publishers each running an English auction simultaneously and needs to decide in which auction to place a bid and at what price.

In particular, at each time unit t=1..T, each advertiser can decide to do one of the following:

- Choose the auction to participate in and the price to bid.
- Move from one auction to another.
- Increase its bid in its active auction.

3) In a predefined closing time, the auction is over. The winners of the available slots of each publisher are the highest ranked ads, and the price each ad pays is equal to its bid.

Under the simultaneous English auction the bidders have a best response strategy that frees the advertisers from strategizing about their proposed bids. Their best response bid is to do the following. Once the initial reserved prices are announced by all the publishers who run a separate English auction each, the bidder should place a bid in the most beneficial publisher auction, which will result in the highest difference between the utility announced by the publisher and the required amount of the bid. With a tie the random selection rule is used. If someone has raised their bid in a particular auction, the bidder will again choose the current most beneficial auction and place a bid with the minimum increment required. The process iterates until no other bidder overbids his bid and he will win. Alternatively, if the lowest proposed bid among all open auctions is higher than the bidder's private value associated with each publisher, then he will leave the auction.

Lemma 4: The SEA auction protocol is incentive compatible from the advertiser perspective, and it is also computational incentive compatible from the publishers' perspective.

To conclude, the above protocol is incentive compatible for both sides.

Lemma 5: The time complexity of the simultaneous English protocol is:

$\left(MK \cdot \max_{i=1..N; j=1..M} \{V_{i,j}\}\right) / (m.increment)$

where M is number of publishers, K is the number of slots of each publisher, $V_{i,j}$ is the private value of advertiser A_i for publisher j, and the *m.increment* is the minimum required in order to place a new bid (to beat the current bid).

To summarize, the SEA has the following properties:

- **Distributed** auction mechanism.
- **Individually rational**: no advertiser will pay more than its private value, no publisher will obtain less than its reserved costs published.
- **Incentive compatible** as discussed above and stated in Lemma 4.
- The **convergence** is guaranteed as discussed in Lemma 5

Next we describe the third mechanism we propose.

DRP: Distributed Reallocation Protocol

The **Distributed Reallocation Protocol** works as follows:

1. Each publisher announces his utilities $U_{i,k,j}$ to all advertisers (which can be defined by a scoring function).
2. According to a random assignment, ads are allocated to publishers.
3. In a random and in an iterative manner, at each step, two publishers are chosen; next, they first search in the unallocated ads to find a possible exchange between their allocated ads such that the utility of at least one of them will increase without affecting any of the advertisers. A possible improvement is to perform this step simultaneously where all publishers are matched to pairs (in the case of an odd number one is left out) and this matching can be iterated.
4. Closing rule can be the minimum of:
   a. Hard deadline or;
   b. No significant improvement is possible (i.e., no changes occur)
5. At the closing time, the allocation is performed according to the final agreed allocation. The advertisement payments are determined according to the original Laddered auction [16]:

$$p_{i,k,j} = \sum_{w=k}^{K-1} \left(\frac{CTR_{i,w,j} - CTR_{i,w+1,j}}{CTR_{i,k,j}} \right) \cdot b_{ad\_allocated\_slot(w+1)}$$

In other words, we calculate the laddered price for each slot $p_{i,j,k}$, based on the second best bid of the winner located in the slot below, $b_{ad\_allocated\_slot(w+1)}$, where w runs from the current slot down to the lowest (1 is the first position).

Lemma 6: The **Distributed Reallocation Protocol** is **strategy proof** from the publishers and the advertisers' viewpoints.

Lemma 7: The time complexity of the distributed reallocation is:

$$O(T(2(N - MK) + \binom{K}{2K} \cdot K \log K)$$

where, T indicates the number of allowed iterations.

Table I summarizes the attributes of the different protocols we developed.

TABLE I. THE PROPERTIES OF THE DIFFERENT PROTOCOLS

Protocol	Centralized/ Distributed	Truthful	Balanced	Optimal
HVCG: Hungarian allocation + VCG payment mechanism	Centralized	**V**	**-**	**V**
SEA: Simultaneous English Auction	Distributed	**IC for the advertiser, computational IC for the publisher**	**V**	**-**
DRP: Distributed Reallocation Protocol.	Distributed	**V**	**V**	**-**

IV. Simulation Details and Results

In order to check and compare the performance of the three suggested protocols, we implemented the following simulation. In each simulation, we considered an environment with N ads, M Web publishers and K slots for each Web publisher. A CTR table was randomly generated, for each ad i and slot k, and two scoring functions were created. One scoring function gives the utility of the ad owner from a click on each ad of each Web publisher, and the second scoring function gives the utility of each Web publisher for each click on each type of ad.

We then, ran the three algorithms to check their performance with different sets of parameters. For each performance we checked, 100 random generated environments were created and tested. The results were compared as summarized below.

Our first test checked the difference in the total utility between the three algorithms. Namely we compared the total efficiency of the ad assignments. Figure 1 presents the total utility of the publishers and the advertisers for the three protocols, where each point is based on 100 runs, with 5 Web-publishers, each offering 4 slots for ads (the minimum increment was 0.2).

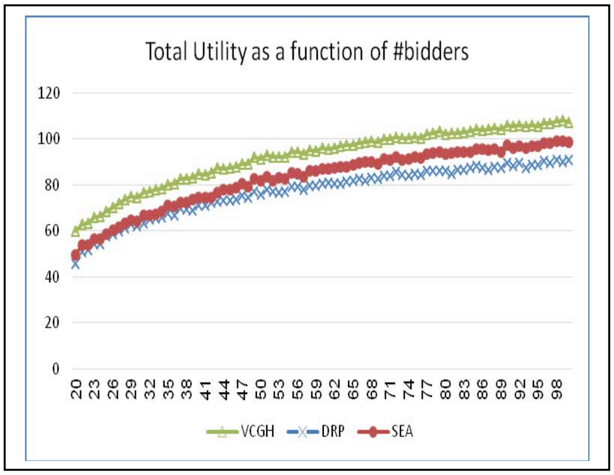

Fig 1. Total utility as a function of the number of advertisers

It is easy to see that the total utility increases as the number of advertisers increases. Moreover, the SEA protocol (Simultaneous English Auction) obtains a slightly higher total utility than the DRP (Distributed Reallocation Protocol). In particular, the SEA obtained 89% utility of the optimal HVCG results, while the DRP protocol obtained 86% out of the HVCG results, for the experiment shown in figure 1.

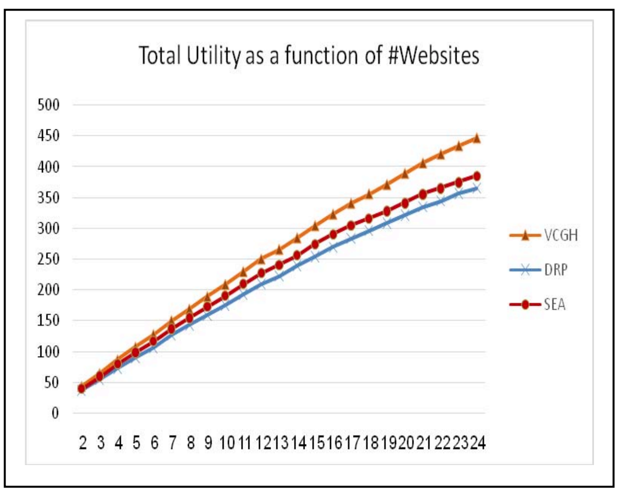

Fig. 2. Total utility as a function of the number of Web publishers

We proceeded by checking the total utility as a function of the number of publishers, where the number of advertisers was fixed at 100 and each publisher offered 4 slots. As depicted in Figure 2, the total utility increases linearly by the number of Publishers, which can be explained by the fact that as the number of publishers is multiplied by two, the number of shown ads is also multiplied by two, and as a result, the total utility is multiplied. In this test, however, the SEA achieved 92% of the VCGH results, while the DRP achieved 85% of the optimal VCGH results. As demonstrated in Figure 1, this can be explained by the fact that, as the number of advertisers increases, the results of the SEA become closer to the optimal VCGH results.

Next, as depicted in Figure 3 we tested the behavior of the different algorithms as a function of the number of slots. It is interesting to see that as the number of slots was 6 or more, the results of SEA were less efficient than the results of DRP.

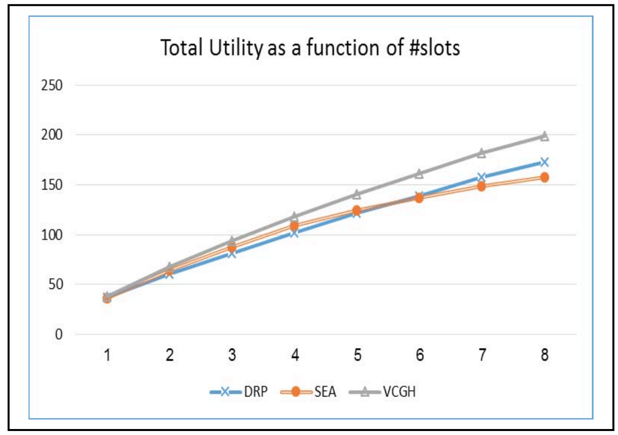

Fig 3. Total utility as a function of the number of slots

As illustrated in Figure 4 we checked the part of the publishers' utility from the total utility obtained by the publishers and bidders. Our findings show that their part is higher with VCGH or DRP than with SEA. The explanation lies in the fact that in SEA the competition between publishers is stronger since advertisers can change publishers during the simultaneous English auction process.

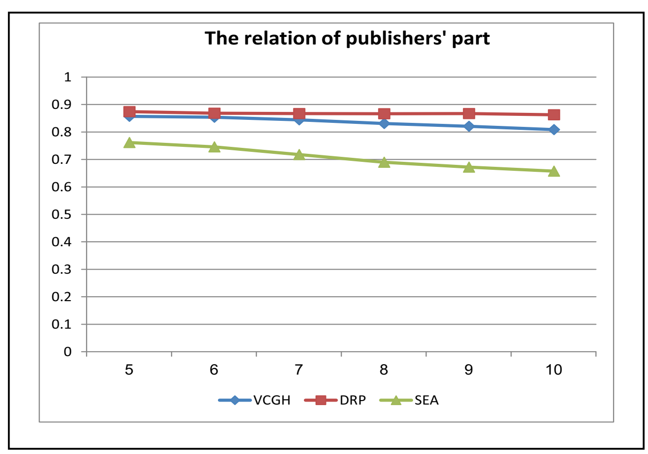

Fig 4. The relation of the publishers' part as a function of the number of publishers

Finally, as displayed in Figure 5 we checked the deficit (loss of balance) with VCGH, since such deficit is its largest drawback. The deficit is measured by the distance between the total utility gained by the advertisers and the publishers, and the total social welfare from the allocation. The existence of this deficit means that the protocol has higher costs than its benefits. Namely the protocol is not budget balanced.

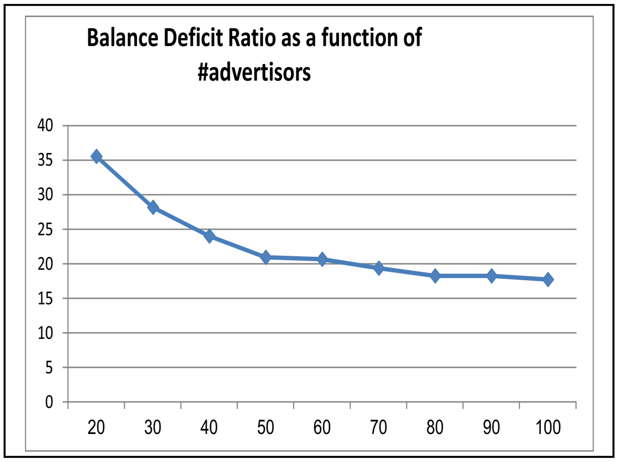

Fig 5. The relation of publishers as part of the function of the number of advertisers

V. Conclusions

In this paper, we consider an ad exchange marketplace where multiple publishers exist, and the publishers may have preferences considering the ads showed in their ad spaces. We suggest three mechanisms for the ad allocation: VCGH (Hungarian algorithm for ad mapping + VCG for payment distribution), SEA (each publisher runs an English auction), and DRP (the allocation is randomly initialized, and then pairs of publishers are matched in order to replace their ads if such an exchange is beneficial for both).

We compared the different protocols according to their attributes and we ran simulations to check their performance for different sets of environments. We found that SEA's results are closer to the results of the optimal VCGH in most of the situations. However, the benefit of the less efficient DRP algorithm lies in the fact that it ensures truthful reports of utilities on the part of both the advertisers and the publishers, while the SEA is incentive compatible from the advertiser's perspective, and computational incentive compatible from the publisher's pespective. The VCGH, however, which promises optimal solutions as well as truthful reports of utilities, suffers from the fact that it is not budget balanced, thus it cannot be implemented in real world situations without meaningful readjustments. Another benefit of the DRP and SEA protocol is the fact that they are distributed and thus they can be implemented without a centralized manager.

In future work, we intend to continue this research in order to find a budget balanced variation to the VCG protocol, while looking for more efficient solution for the distributed heuristic auction protocols for the ad exchange marketplace with publishers with preferences concerning the ad they publish. In addition, we intend to consider dynamic environments where ad auctions are repeatedly performed, and the click through rates vary over time.

Bibliography

[1] S Muthukrishnan, Ad Exchange: Research Issues, Internet and network Economics Lecture Notes in the Computer Science Volume 5929, pp 1-12, 2009.

[2] http://www.rightmesdia.com/rightmedia.com/right-media-101/. In http://www.rightmedia.com/

[3] adBrite http://www.adbrite.com/

[4] OpenX http://www.openx.co/

[5] www.google.com/adexchange/AdExchangeOverview.pdf, product information at http://www.doubleclick.com/products/advertisingexchange/index.aspx. In http://www.doubleclick.com.

[6] Google Adsense, www.google.com/adsense

[7] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. The American Economic Review, 97(1):242–259, (2007).

[8] H. Varian. Position auctions. International Journal of Industrial Organization, 25 (6):1163–1178, (2007).

[9] A. Goel and K. Munagala. Hybrid keyword search auctions. In WWW, pages 221–230, (2009).

[10] Y. Mansour, S. Muthukrishnan and N. Nisan, Double Click Ad Exchange Auction. Computer Science and Game Theory, arXiv:1204.0535v1 [cs.GT], 2012.

[11] R. Azoulay and E. David, Protocols and Strategies for Multi-Slot Publisher Oriented Advertising, IAT (2012).

[12] R. Azoulay and E. David, Truthful and efficient mechanisms for Website dependent advertising auctions. Multiagent and Grid Systems 10(2): 67-94 (2014).

[13] J. Munkres. Algorithms for the assignment and transportation problems. Journal of the Society for Industrial and Applied Mathematics, 5, 1:32–38, 1957.

[14] H. W. Kuhn. The hungarian method for the assignment problem. Naval Research Logistics Quarterly, 2:83–97, 1955.

[15] T. Groves. Incentives in Teams. Econometrica, 41:617–631, 1973.

[16] G. Aggarwal., A. Goel., and R. Motwani, Truthful auctions for pricing search keywords. In: ACM EC, 2006

[17] G. Aggarwal, S. Muthukrishnan, D. Pa'l, and M. Pa'l. General auction mechanism for search advertising. In WWW, pages 241–250, 2009.

[18] www.jct.ac.il/~azrina/IAT15-proofs.pdf