Ad Exchanges: Research Issues

S. Muthukrishnan

Google Inc.

Abstract. An emerging way to sell and buy display ads on the Internet is via ad exchanges. RightMedia [1], AdECN [2] and DoubleClick Ad Exchange [3] are examples of such real-time two-sided markets. We describe an abstraction of ad exchanges. Based on that abstraction, we present several research directions and discuss some insights.

1 Introduction

Activities on the Internet can be abstracted as interactions due to three parties. There are

- the *users* who navigate to various web pages,
- the *publishers* who control the web pages and generate the content in them, and
- are *advertisers* who wish to get the attention of the users using the publishers as the channel for placing ads on the pages. We call such ads on webpages as the *display ads*.

Precisely what display ads show when a user accesses a page is a detailed process. A central issue is matching advertisers to publishers, where the number of advertisers and publishers is very large. Direct negotiation between advertisers and publishers might work for large companies, but currently there are *intermediaries* such as ad agencies, ad networks and publisher networks that aggregate several parties. An advertisers might use an ad agency to develop the marketing campaign and use an ad network that pools many advertisers to negotiate with a publisher network that places ads with many publishers. There are several intermediaries and their services overlap eg., networks comprise both advertisers and publishers, or agencies that are also ad networks, etc. Thus there are several ad paths between an advertiser and a user; some ad paths are shown in Figure 1.

An emerging way of selling and buying ads on the Internet is via an *exchange* that brings sellers (publishers) and buyers (advertisers) together to a common marketplace. There are exchanges in the world for trading financial securities, currency, physical goods, virtual credits, and much more. Exchanges serve many purposes from bringing efficiency, to eliciting prices, generating capital, aggregating information etc. Market Microstructure is the area that studies all aspects of such exchanges [4, 5].

Ad exchanges are recent. RightMedia [1], AdECN [2] and DoubleClick [3] are examples. Ad exchanges let ad networks and publishers transact centrally for ads.

– Publishers expect to get the best price from the exchange, better than from any specific ad network; in addition, publishers get liquidity.
– Advertisers get access to a large inventory at the exchange, and in addition, the ability to target more precisely across web pages.
– Finally, the exchange is a clearing house ensuring the flow of money.

In many ways, these ad exchanges are modeled after financial stock exchanges. Since 2005 when RightMedia appeared, ad exchanges have become popular. In Sept 2009, RightMedia averaged 9 billion transactions a day with 100's of thousands of buyers and sellers. Recently, DoubleClick announced their new ad exchange. It seems ad exchanges are likely to become a major platform for trading ads.

We abstract a model for ad exchanges. Based on the model, we present research problems in auction theory, optimization and game theory. The goal is to present a blueprint for research in design, analyses and use of ad exchanges.

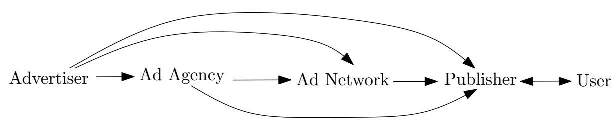

Fig. 1. Ad paths

2 Ad Exchange

We present an abstract *AdX* model to describe ad exchanges. It is defined as a sequence of events.

1. User u visits the webpage w of publisher $P(w)$.
   For now, we assume page w has a single slot for ads.
2. Publisher $P(w)$ contacts the exchange E with $(w, P(u), \rho)$ where ρ is the *minimum price* $P(w)$ is willing to take for the slot in w and $E(u)$ is the information about u that P shares with E.
   We assume that E knows all the contents of w as well as the various specifics of the ad slot in it, including its dimensions and inappropriate ads for that slot as agreed on with $P(w)$. We also assume that the exchange manages information about user u in a manner agreed upon with $P(w)$ via $P(u)$; $P(u)$ may only be partial information about u, possibly even empty. $P(w)$ has moral and contractual relationship with its viewers u while also having the incentive to help advertisers target users suitably; therefore $P(u)$ is determined by P. Also, E can independently crawl contents of w if needed, in many cases, so we might as well assume that in the model.
3. The exchange E contacts ad networks $a_1, \ldots, a_m$ with $(E(w), E(u), E(\rho))$, where $E(w)$ is information about w provided by E, and likewise, $E(u)$ is the

information about u provided by E and $E(\rho)$ is the pricing information E shares with ad networks.

$E(u)$ is the information E provides to the ad networks as agreed upon with $P(w)$; this is possibly different from $P(u)$. When $P(w)$ entrusts E to reveal w, we assume that ad networks know contents of w and also information processed from w, such as keywords, topics of w etc. The ad networks may have the resources to obtain this on their own or resourceful exchanges can do it for them. This is modeled as $E(w)$. There are instances when $P(w)$ does not wish ad networks to know w, in which case, $E(w)$ will be only derived information about w and w's identity will remain unknown to a_i's. Finally, $E(\rho)$ may be different from ρ, for example, by adding a reserve price to ρ.

4. Each ad network a_i returns (b_i, d_i) on behalf of its customers which are the advertisers; b_i is its *bid*, that is, the maximum it is willing to pay for the slot in page w and d_i is the ad it wishes to be shown. The ad networks may also choose not to return a bid.

   It is assumed that $b_i \geq \max\{\rho, E(\rho)\}$, else no bid is returned. It is also assumed that d_i is suitable for the ad slot. Further, it is assumed that a_i targets ads based on its contracts with its customers and negotiates prices for the service with them.
5. Exchange E determines a winner i^* for the ad slot among all (b_i, d_i)'s and its price $\rho \leq c_{i^*} \leq b_{i^*}$ via an *auction* and returns (c_{i^*}, d_{i^*}) to $P(w)$.

   It is assumed that the winning network i^* becomes aware of the outcome including the price, and in some but not necessarily all instances, the losing networks can determine that too. E is responsible that d_i is suitable for w, u according to its contract with $P(w)$. This may be accomplished in a variety of ways from pre-verification to outsourcing the task to the networks. E negotiates pricing for its service with various $P(w)$'s and a_i's. Also, E generates bills and reports, collects payments from a_i's, and makes payments to $P(w)$'s.
6. The publisher $P(w)$ serves webpage w with ad d_{i^*} to user u. This is known as an *impression* of ad d_{i^*}.

   The ad d_{i^*} is rendered by user's browser and the user interacts with the ad in a variety of ways.

We have left out system level details in the model above. Various internet protocols are used for forwarding, referencing, accounting, etc. For example, ad d_i may be directly returned by the ad network via E or passed by reference as supported by common http protocol; then user's browser can fetch the ad while rendering w. For efficiency reasons, certain steps may be optimized out and or made offline. For example, ads and bids may be loaded onto the Exchange *a priori* so networks are accessed only infrequently. Also, due to time constraints, crawling and processing w may be overlooked and $E(w)$ will be minimal or there may be timeouts. Finally, there are considerable details to avoiding inappropriate ads and contextual targeting of ads. For our purposes, the details above suffice. The flow of the model is shown in Figure 2.

The AdX model captures the essence of ad exchanges that the instrument traded is *ad impressions*. So, AdX model has real time auction for each impres-

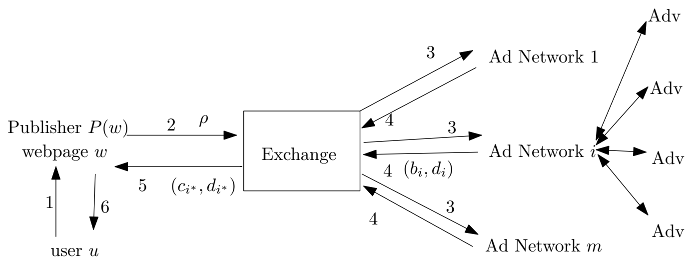

Fig. 2. AdX Model

sion and the bids are CPM (cost per mille or thousand impressions). The entire execution above begins when a user arrives at a webpage and must be completed before the page is rendered on their screen. This sets an upper bound for the execution in 10's or 100's of milliseconds. Thus AdX model provides a spot market for ad impressions. The assumption is that higher order campaign goals such as maximizing number of impressions subject to budget criteria, or reaching as many target users as possible, and other goals will be executed by optimized bidding by the networks, and indeed one can design more sophisticated instruments on top of the spot market.

AdX vs Financial Exchanges. As in financial exchanges, the buyers and sellers come to AdX only when they choose. Further, individual advertisers access AdX through intermediaries like in financial exchanges. However, a significant difference is in the nature of what is traded. Unlike financial securities, ad impressions are heterogeneous, significantly differing from one instance to another in their value depending on their impact on the user which varies widely across millions of users and publishers. Further, impressions are supremely perishable: if the trade does not happen within the browser performance anticipated by the user, the opportunity to place an ad is lost. As a result, AdX has two important tasks. First is informational. AdX embodies the effort by various parties to help advertisers discover users to target, eg via $E(w), E(u)$. Second is economics. With vastly heterogeneous goods, many of the pricing methods face challenges and AdX enables the market to discover prices by making trading automatic and via auctions.

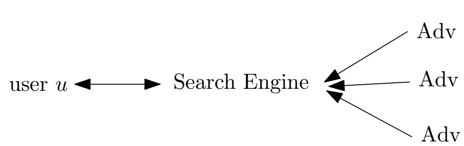

Fig. 3. Sponsored Search

AdX vs Sponsored Search. The AdX model is distinct from *sponsored search*. In sponsored search, a user poses a query at a search engine and gets search results together with ads arranged top to bottom. The assignment of ads to positions is by an auction among all advertisers who placed a cost-per-click (CPC) bid on a keyword that matches the query. If the user clicks on an ad, that advertiser pays the search engine the auction price. The ad path is simple and shown in Figure 3. The dynamics are simpler since there is a single publisher and a one-sided marketplace of buyers. On the other hand, sponsored search aligns the incentives of advertisers and search engines with the quality of ads for the users, and hence, the publisher faces the challenge of monitoring and maintaining quality. The underlying auction for sponsored search is analyzed in [8, 9]. There are many outstanding research issues in sponsored search [6], but they are not the focus here. We will henceforth focus on the AdX model.

3 Research Issues

We describe several research issues using examples. The issues when formalized more precisely will lead to different research problems.

3.1 Basic auction at the exchange

A central question concerns the auction at the exchange. Let us assume that $E(\rho) \geq \rho$. There is a single slot of ad being auctioned, and the bidders know ρ and participate only if their bid is above ρ. Further it is a sealed bid auction since each network bids directly with the exchange. The setting appears to be a standard auction in one-sided market. Revenue equivalence theorem [7] would then indicate that first price, second price and several other auctions will all yield identical expected revenue in Bayesian equilibrium under certain natural assumptions. Still, second price or Vickrey auctions are preferable since they encourage more stable dynamics. So, we will assume second price auction as the benchmark when needed. We now discuss a nuance.

Example. Say network a_1 has 2 advertisers a_{11} and a_{12} with bids 10 and 8 resp., while network a_2 has a single advertiser a_{21} with bid 5. a_1 forwards 10 and a_2 forwards 5 to the exchange. Second price auction at the exchange will declare a_1 the winner with price 5. a_{11} wins the ad slot. Now a_1 may charge a_{11} either 5 (the exchange price) or 8 (the second price in a_1). Instead, if a_1 revealed both bids 10 and 8 to the exchange, second price auction at the exchange would charge 8 to the winner a_{11}. Thus since the exchange does not know the entire book of the ad networks, a network has an excess to disburse and the publisher gets less than the best price from the book. ■

Book value is the second largest value of all the bidders.

Problem 1 *Assuming ρ is exogenous and assuming the advertisers reveal their bids truthfully to the networks, is there a possibly truthful auction at the exchange that will extract a large fraction of the book value?*

The tricky case is when the bid that sets the price of the winner is in the winner's network. A straightforward approach is to generalize the protocol and ask each network i to forward *two* bids (b_1^i, b_2^i) where b_1^i is the largest of the bids in i and b_2^i the second largest. A naive strategy is to declare $i^* = argmax_i b_1^i$ the winner with price $\max\{b_2^{i^*}, \max_{j \neq i^*} b_1^j\}$. In this case, ad networks have no incentive to declare nonzero b_2^i. A different allocation strategy would be to declare largest $b_1^i + b_2^i$ as the winner. This incentivizes networks to bid $((b_1^i + b_2^i)/2, (b_1^i + b_2^i)/2)$ and is not truthful.

3.2 Auction and bidding by ad networks

Consider mechanisms for the ad networks. In general, the exchange can not assume ad networks will follow any particular mechanism. Still, exchange's choice of auction will impact mechanisms ad networks will use, and choice of mechanisms of ad networks will influence auction at the exchange.

Problem 2 *Assuming ρ is exogenous, the exchange runs a second price auction with reserve price $r \geq \rho$, ie., $E(\rho) = \rho + r$, and advertisers are* captive*, that is, remain with their choice of ad network throughout, what is a revenue optimal mechanism for an ad network?*

Seen from an ad network i's point of view, there are several advertisers bidding to be chosen for the impression. The questions are, which ad to choose and what price to charge? This appears to be identical to the standard auction framework where optimal mechanisms are known [11]. Under suitable assumptions, the solution is to run a second price auction with an appropriately chosen reserve price r_i. The key difference here is that the ad network i does not have a guaranteed good to auction, rather it has a *contingent* good. More formally, there is some probability $\alpha(b)$ that it will win the impression if it bid b at the exchange. This probability function α is determined by the bidding strategy of other networks and their advertisers' valuations.

Under standard assumptions that advertisers have values drawn from a known distribution, one can extend the theory of optimal auctions [11] to derive an optimal mechanism for the ad networks [12]. For example, one can show that the revenue-optimal mechanism for a network is to randomize its bids in some range $[l_r, u_r]$, dependent on r. A more detailed understanding of the equilibrium will be of great interest. Even without the complication of what auction to run in the network, the problem of bidding for higher campaign goals is challenging in presence of $\alpha()$. Recently, this was studied in [13].

3.3 Auctioning with heterogenous valuations

While the exchange auctions impressions, there are instances when an impression is valued very differently by the bidders. For example, an impression to a particular user may be far more valuable to one bidder than others. This targeting is enabled by the information $E(u)$ shared by the publisher $P(w)$ with the ad networks.[1]

Example. We have two bidders A and B, A has value $v_A = 100$ and B has value $v_B = 1$ for the ad slot. In equilibrium of the first price auction, A wins by bidding $1 + \epsilon$ for some $\epsilon > 0$ and exchange's revenue is $1 + \epsilon$. This outcome holds for second price auction as well. On the other hand, maximum value to be extracted is $\max v_A, v_B = 100$. What is the maximum revenue a mechanism can extract in equilibrium? ∎

Say bidder i has value v_i, bids b_i, and there is a single item to be sold. Assume $b_1 \geq b_2 \geq \cdots$. Then, the maximum revenue is $R^* = \max_i v_i$. Classical results [11] would provide most revenue assuming the distribution from which v_i's is drawn is known. Here, instead, we focus on *prior-free* case where distributions are not known *a priori*, not even the maximum of individual distributions.[2] We know from [14] that no truthful mechanism exists with (roughly) expected revenue $\Omega(\frac{R^*}{\log R^*})$; on the other hand, they show an auction with expected revenue $\Omega(\frac{R^*}{(\log R^*)^{1+\epsilon}})$ for fixed $\epsilon > 0$. Their auction [14] is as follows. With probability $1 - \delta$, use the second price auction. With probability δ, choose r according to the distribution below and if $b_1 \geq r$, highest bidder wins at price r:

$$f(x) = \frac{\epsilon}{x(\log(x/b_2) + 1)^{1+\epsilon}},\ x \in [b_2, \infty).$$

Problem 3 *Design a non-truthful mechanism for prior-free auction of a single slot with near-optimal revenue, but with good equilibrium properties.*

An approach is to use quasi-proportional allocation in which the good is allocated to one bidder using a distribution where ith bidder is picked with probability $\frac{f(b_i)}{\sum_i f(b_i)}$, for some suitable function f. If $f(x) = x$, we have the well-known proportional allocation. Revenue properties of quasi-proportional mechanisms are not well-studied. In [17], authors show that with $f(x) = \sqrt{x}$, one can extract $(R^*)^\gamma$ revenue for some $\gamma > 0$ under certain conditions. Further, using [18], the equilibrium can be characterized for quasi-proportional auctions. A deeper study of quasi-proportional auctions will be of great interest. In particular, are there functions f for which quasi-proportional allocation generates more revenue in its equilibrium than the lower bound in [14]?

For AdX, we need a generalization of the standard prior-free setting to ad networks as well as the exchange (much as [12] extends [11]). For an economics perspective on market clearing and role of intermediaries in prior-free auctions, see [15].

---

[1] `http://www.adecn.com/faq_4.html`.

[2] In practice, one would argue that in repeated auctions as in the exchange, one can learn the distributions. It is an interesting exercise what can be learned from the data – not necessarily the entire distribution, just enough to extract revenue – and how, as well as its impact in engineering the system, in particular, when learning is likely to be only approximate.

Problem 4 *Design (even non-truthful) mechanisms for prior-free bidding of ad networks in AdX, with good equilibrium properties and (near-)optimal revenue.*

3.4 Callout optimization

In Step 3, E seeks bids from ad network a_i's. This may be accomplished at a system level by (a) having a_i's preload their ad campaigns into E so suitable ad campaigns and their bids that apply to $(E(w), E(u), E(\rho))$ can be retrieved locally, (b) hosting a_i's bidding software in E's machines so the software can manage the network's campaigns and generate suitable (b_i, d_i), or (c) making http calls to a_i's servers, and awaiting (b_i, d_i) to be determined by the network. (a) and (b) are essentially not real-time because networks cannot update bids, bidding logic, key parameters and relevant data from impression to impression. We focus on (c) and discuss an optimization problem.

It is resource-intensive for E to call out to each network for each impression. In practice, ad networks can describe $(E(w), E(u), E(\rho))$'s of cumulative interest to their customers, and therefore for each impression, only a subset $S_{(E(w),E(u),\rho)}$ of ad networks need to be called. Still, we may assume that it will be difficult if not impossible for each network to take all the http calls from E. So, here is an optimization problem E faces.

Problem 5 *Each ad network i has bandwidth budget B_i. Say E has bandwidth budget of B. Design an online algorithm for E that for each incoming call (w_j, u_j, ρ_j), chooses a subset $S_j \subseteq S_{(E(w_j),E(u_j),\rho_j)}$ of networks to call such that no ad network i gets more than B_i calls per second, E make fewer than B calls per second, and optimizes the expected*

- *number of bids, ie, number of nonempty $(b_i(j), d_i(j))$'s received at E, or*
- *efficiency $\sum_j \max_i b_i(j)$, or*
- *sales revenue $\sum_j \max_{i \mid b_i(j) \neq \max_i b_i(j)} b_i(j)$, or*
- *profit for E.*

In the problem above, we need stochastics. The algorithm has to know about likelihood of ad network i making a bid for $(E(w_j), E(u_j), E(\rho_j))$. In particular, $\Pr(b_i(j) \geq E(\rho_j) | E(w_j), E(u_j))$ and $Exp(b_i(j) | b_i(j) \geq E(\rho_j))$ are useful (assuming $E(\rho) \geq \rho$). Hence, E needs to estimate these terms for each j, which is a challenging Machine Learning challenge by itself. Given these parameters, initial approximation results for the problem are in [19].

3.5 Publisher optimization and strategies

Publisher $P(w)$ has many decisions to make when user u visits w.

Accessing AdX. $P(w)$ has several sources for ads including contracts with advertisers previously signed by their own sales teams, contracts with specific ad networks that commit certain slots or traffic, and the ability to reach AdX on any impression real time. Filling many impressions with highest price ads from

AdX will have long term impact on other sources, and possible failure to deliver on contracts; filling contract commitments might lead to loss of revenue from the AdX spot market. The task then is to design an online algorithm to commit impressions to different ad sources in order to honor contracts as well as maximize revenue. Some initial results are in [21] where solutions are based on the publisher virtually bidding on behalf of ad sources. More detailed understanding of strategies will be of great use to individual publishers.

Form of inventory. If publisher $P(w)$ wishes, E can limit information via $E(w)$, not reveal identity of w or $P(w)$, and merely provide information about the nature of the site w eg., sports/baseball, etc. We call this *undisclosed* inventory. [3]; in contrast, when identity of w is shared with the ad networks, we call it *disclosed* inventory. Publishers may choose to keep some of their inventory in AdX undisclosed, in order to avoid conflicts with their other ad sources described earlier. It is assumed that undisclosed inventory fetches less than disclosed inventory. Then, $P(w)$ has an online optimization problem: for each impression, how to choose disclosed or undisclosed option in E to trade off short term vs long term revenue. A suitable model to address this problem will be of interest.

Price. Publisher $P(w)$ has to choose ρ while accessing E: large ρ may not generate a bid, and a small ρ may undervalue the inventory. One way to choose ρ is as the maximum over other sales channels for an ad at the slot, but in general, one needs an online ρ setting algorithm that endogenizes the demand and supply in the system adaptively. Again, a clean model and implementable algorithm is of interest.

Here is the combined problem of inventory management and accessing AdX.

Problem 6 *Given models for impressions inventory (w, u), models for bids (b_{i^*}, d_{i^*}) from E, models of ad sales and prices through other channels, design an algorithm that on each impression (a) decides whether to go to AdX, (b) chooses disclosed or undisclosed inventory at AdX, and (c) selects min price ρ, in order to optimize the expected overall (long term) revenue.*

3.6 Arbitrage bidding and risk analysis

AdX model trades impressions and uses CPM prices. Ad networks can sell other pricing methods, such as pay-per-click with cost-per-click (CPC) prices to their advertisers. This comes with an arbitrage opportunity and associated risk as described below.

Example. Define click-through-rate (CTR) as average fractional number of clicks per impression. Consider an ad with CTR 0.1 and CPC of \$1. Ideally, the network should bid \$100 CPM; then spend is \$100 for 1000 impressions, revenue is \$100, and they are even. However, in any empirical run of auctions,

[3] See blind vs disclosed inventory in `http://www.adecn.com/faq_5.html`, or branded vs anonymous in `http://www.doubleclick.com/products/advertisingexchange/benefits_for_sellers.aspx`.

CTR estimates are not precise. (1) Say CTR was overestimated to be 0.2. Then the network bids \$200 CPM, spend is \$200 for 1000 impressions, and revenue is \$100. This assumes the network will get at least 1000 impressions if they overbid, which is reasonable given large inventory in ad slots. (2) Say CTR was underestimated to be 0.05 and the network bids \$50 CPM. Then spend is $\$50x$ for x fraction of 1000 impressions, and revenue is $\$100x$. The assumption is the network will only get x fraction of 1000 impressions. To summarize, the outcome is $(200 - 100)$ loss if network overbid vs $(100 - 50)x$ profit if network underbid and loss of $(1 - x)$ fraction of impressions. ■

Problem 7 *Consider advertisers who contract with the network for CPC bids, certain reach (number of distinct users reached) and frequency counts (the number of times a user sees an ad). Design an algorithm for ad networks to place CPM bids into AdX that for a given risk level and volatility, maximizes expected revenue and guarantees contract counts. Take into account bidding into the exchange for contingent good as in Section 3.2.*

3.7 AdX integrity

Exchanges in general strive for transparency. The ability of participants to understand the inventory they buy or be convinced of the integrity of underlying process will, in general, induce more to access the exchange. In reality, systems need a balance of overall trust and with legal or technological checks of key parts, if needed. While this balance will be worked out by the marketplace, here is a theoretical problem.

Problem 8 *Design a cryptographically sound real-time auction protocol so that any participating party in AdX can verify that (a) all communication, accounting and computations were performed correctly, and (b) auction was closed envelope, that is, no bidder sees others' bids prior to the auction. This has to work for repeated auction of impressions in AdX where some information is revealed between impressions.*

Secure and collusion free auction design has been explored by the cryptography community, eg using homomorphic encryption scheme [20]. However, one still needs realtime methods for AdX application, methods which will run with round trip times of browsers accessing webpages. Also, one needs a clear model to prove strong cryptographic properties of the various protocols. Some progress is in [22].

3.8 Configuration auctions

Even when there is a single slot in w, different configurations of ads — eg 1 video ad or 2 image ads or 4 text ads — might fit there. While the exchange will return the most efficient configuration in total CPM bids, an interesting question is how to spread the price among the ads in the winning configuration

(Problem 6 in [6]). Further, inherently, the bids from the networks are for entire configurations which makes it more like configuration auctions and associated externalities [23].

We go beyond to look at cases where w has multiple ad slots. How should they be sold? One approach is to sell the slots independently. This is easy on the system but advertisers may wish to appear in only one slot on a page which can not be guaranteed. Another approach is to sell all the slots in one block, but such exclusive packages may not be popular except on premium sites. A realistic approach is to make AdX more sophisticated and consider all slots together in some auction. In reality, advertisers have complex preferences. For example, an advertiser may value a slot on the right highly *only* if they did not win the top slot. That is, advertisers have *conditional values.*

Problem 9 *Devise a suitable bidding language for advertisers and ad networks to express (conditional) preferences for slots, and design suitable auction mechanisms for the exchange as well as for ad networks to allocate substitutable winners if any.*

We propose a simple approach. Publisher $P(w)$ presents a list of slots in w in the order in which will be auctioned. Then ad networks can return one bid for each slot, and specify the maximum number of impressions they wish. The auction proceeds slot by slot; winners are removed from consideration for successive slots if they are fulfilled. We call this *linearized auction* (LA). LA trades off efficiency for expressiveness and lets bidders represent some conditional preferences but not others. Studying suitable LA mechanisms including the impact of publisher's choice of the list will be of interest. Also, comparison of LA mechanisms with richer tree bidding [24] will be of interest.

4 Concluding Remarks

This writeup provides some insight into the research issues in AdX that models exchanges like RightMedia, AdGCN and DoubleClick. These exchanges are emerging as major new platforms to sell display ads on the Internet, but are still nascent. Progress on research issues listed here will likely impact the design and growth of not only these ad exchanges but also the "ecosystem" of bidders, optimizers and quantifiers around them. We list some more issues:

1. *Game theory of adverisers.* Advertisers may go to multiple networks, or choose networks strategically. What are the resulting dynamics? How does their strategic behavior affect competition within their campaigns along multiple ad paths, across advertisers, across ad networks and ultimately, across exchanges?
2. *Ad Quality.* In sponsored search, quality of an ad is correlated with click-through, and so is the pricing and incentives of the advertiser. We need a similar quality metric for impressions and endogenize that in the auction to align advertiser incentives. A proposal is to generate a suitable Markov model for users that will capture even the long term impact of ad impressions.

Acknowledgements I sincerely thank my colleagues Rahul Bafna, Eyal Manor, Yishay Mansour, Noam Nisan and Scott Spencer at Google for weekly AdX lessons. Thanks to Vahab Mirrokni, Mallesh Pai, David Parkes, Al Roth, Hal Varian and Ricky Vohra for pointers, and Sudipto Guha, Boulos Harb, Aranyak Mehta and Mukund Sundarajan for comments. Thanks to Michael Rabin and Moti Yung for an enjoyable collaboration.

References

1. RightMedia. `http://www.rightmedia.com/` More info at `http://www.rightmedia.com/right-media-101/`.
2. AdECN. `http://www.adecn.com/`. Whitepaper at `http://www.adecn.com/resources/ATrueExchangeforOnlineAdvertising.pdf`. Presentation at `http://www.adecn.com/resources/AdECNPresentation.pps`.
3. DoubleClick Ad Exchange. `http://www.doubleclick.com/products/advertisingexchange/index.aspx`
4. M. O'Hara, Maureen. *Market Microstructure Theory*, Blackwell, Oxford, 1995.
5. A. Madhavan. Market Microstructure: A Survey. *Journal of Financial Markets*, 3(3), 2000, 205–258.
6. S. Muthukrishnan. Internet Ad Auctions: Insights and Directions. *Proc. ICALP* (1) 2008, 14-23.
7. P. Klemperer. Auction Theory: A Guide to the Literature. *J of Economic Surveys* 1999, 227-286.
8. H. Varian. Position Auctions. *Intl J of Industrial Organization*, 2006.
9. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, vol. 97(1), 2007, 242-259.
10. W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 1961, 16: 8-37.
11. R. Myerson. Optimal auction design. *Math of Operations Research*, 6:58 73, 1981.
12. J. Feldman, V. Mirrokni, S. Muthukrishnan and M. Pai. Charaterizing Equilibria of Auctions with Mediators. *Manuscript*, 2009.
13. A. Ghosh, B. Rubinstein, S. Vassilvitskii and M. Zinkevich. Adaptive bidding for display advertising *Proc. WWW*, 2009.
14. P. Lu, S. Teng and C. Yu. Truthful Auctions with Optimal Profit. *Proc. WINE* 2006: 27-36.
15. S. Baliga and R. Vohra. Market Research and Market Design. `http://www.kellogg.northwestern.edu/faculty/baliga/htm/mrandmd.pdf`, 2003.
16. G. Aggarwal, G. Goel and A. Mehta. Efficiency of (revenue-)optimal mechanisms. *ACM Conference on Electronic Commerce* (EC) 2009: 235-242.
17. V. Mirrokni, S. Muthukrishnan and U. Nadav. Quasi-Proportional Mechanisms: Prior-free Revenue Maximization. arXivs, 2009. `http://arxiv1.library.cornell.edu/pdf/0909.5365v1`.
18. E. Even Dar, Y. Mansour and U. Nadav. Convergence in Proportional Games. *Proc. STOC*, 2009.
19. T. Chakraborty, E. Even-Dar, S. Guha, Y. Mansour and S. Muthukrishnan. Callout optimization. *Manuscript*, 2009.
20. D. Parkes, M. Rabin,S. Shieber and C. Thorpe. Practical secrecy-preserving, verifiably correct and trustworthy auctions. *Proc. EC*, 2006, 70–81.

21. A. Ghosh, P. McAfee, K. Papineni, and S. Vassilvitskii. Bidding for Representative Allocations for Display Advertising. *Proc of the 4th intl Wkshp on Internet and Network Economics* (WINE), 2009.
22. M Rabin, S. Muthukrishnan and M. Yung. Fast, communication efficient, provably secure protocols for real-time Vickrey auctions. *Manuscript*, 2009.
23. S. Muthukrishnan: Bidding on Configurations in Internet Ad Auctions. *Proc. COCOON* 2009: 1-6.
24. S. Lahaie, D. Parkes and D. Pennock. An Expressive Auction Design for Online Display Advertising. *AAAI* 2008:108-113.