Internet Ad Auctions: Insights and Directions

S. Muthukrishnan

Google Inc., 76 9th Av, 4th Fl., New York, NY, 10011
muthu@google.com

Abstract. On the Internet, there are advertisements (ads) of different kinds: image, text, video and other specially marked objects that are distinct from the underlying content of the page. There is an industry behind the management of such ads, and they face a number of algorithmic challenges. This note will present a small selection of such problems, some insights and open research directions.

1 Introduction

Everyday we interact with the Internet in several ways. For example, we read news from an online source, go to a portal to check email or start at a search engine to navigate the web or for discovery. We belong to some explicit social network and interact with "friends" online. We plan and execute projects or events. Through these activities, we express our interests, intent, an implicit openness to discover new things, not only individually, but also as members of different groups. These are the *signals* we generate on the Internet.

This Internet world is valuable to businesses. They seek to benefit from this online world not only by transacting their business but also by marketing themselves. Marketing is predominantly done via advertisements (ads), using the many signals. The ads could be in different forms from text snippets to images and even videos. These ads provide some information and seek to get the attention of the users, and ultimately induce some action, direct or indirect.

There is now an industry of companies that enable the process above. This includes companies that bring users to their site, procure ads and manage the ad presentation and billing processes, as well as companies, that are mediators, merely placing ads where they are slotted and needed in others' properties. For most part, the ad placement is determined by *auction*s, so there is focus on the strategizing and gaming aspects of the behavior of advertisers and ad placement companies. In addition, there are companies that enable advertisers as well as content providers to strategize and optimize their goals across multiple parties.

All of this generates many algorithmic challenges. In this paper, we present an overview Internet ad auctions typically used. We will present some insights and directions for research. In many cases, the directions are only abstractions that are aimed at spurring some theoretical research thought and are not direct business problems. We believe that Internet ad auctions is an area of research where novel ideas will have great impact in practice because large scale Internet ad systems exist, are successful, and can be modified with reasonable effort. This

L. Aceto et al. (Eds.): ICALP 2008, Part I, LNCS 5125, pp. 14–23, 2008.

is an evolving area, and this write-up covers a small vantage point only; we will maintain an updated version of the write-up over time.

2 Basics

In this section, we describe the type of ads, the process of ad campaign development and an overview of the auction mechanism that is involved.

2.1 Types of Ads

There are different types of Internet ads, depending on the nature of signals used as well as the nature of ads. We give a few examples.

Sponsored search ads. When a user poses a query at a search engine, the search engine returns search results together with advertisements that are placed into positions, usually arranged linearly down the page, top to bottom. On most major search engines, the assignment of ads to positions is determined by an auction among all advertisers who placed a bid on a keyword that matches the query. The user might click on one or more of the ads, in which case (in the pay-per-click model) the advertiser receiving the click pays the search engine a price determined by the auction. This is known as sponsored search ads.

Content ads. Users go to several sites for their intrinsic content. For example, this includes content providers such as established news sources or more individualistic blog sites and other publishing sites. Hence, content at such sites matches users' interests or intent. Content providers use these signals derived from their content to target ads to users. This is known as content ads.

Display ads. Users go to portals and other pivotal sites as starting points of their interaction with the Internet. Advertisers seek to get such users' attention by display of ads which may not be directly determined by the content of such sites. Banners and pop-ups are examples of such display ads.

Social networking ads. Users belong to one or more social networks, and interact with friends and contacts. The signal of such friends, friends' signals and so on, may be used to present ads to a user. Such ads are social networking ads. ■

2.2 Life of an Ad

Planning and execution of an ad campaign involves at least three mains stages.

- *Targeting.* Advertisers determine the potential target for ad campaigns. Targeting may take the form of demo or psycho-graphics and rely on users' signals.
- *Ad placement and optimization.* Advertisers strategize with budgets and prices for favored placement of their ads and optimize the overall impact of their budgets.

- *Ad effectiveness.* Any ad campaign needs to evaluate its over all impact. In offline ad media, effectiveness of an ad may have to be measured via surveys or coupons etc. In Internet ads, there are other measures of ad effectiveness including click-through as well as change in traffic levels due to an ad.

Life of an ad starts with its targeting, proceeds to production, optimization and execution of the ad campaign, and finally, evaluation of the effectiveness of the campaign. Algorithmic problems arise in each of these stages.

2.3 Ordered Ad Auctions

We describe the popular Internet ad auctions formally, and call them *ordered ad auctions.* An ordered ad auction is defined by a tuple (N, K, v, α, β). The set $N = \{1, \ldots n\}$ is the set of bidders (advertisers) and the set $K = \{1, \ldots, n\}$ is the set of positions, ordered top to bottom. Each bidder $i \in N$ is associated with two values, v_i which is her valuation for a click and α_i which is her *click-through rate* (ctr). Each position $\ell \in K$ is associated with a click-through multiplier β_ℓ. We have $\beta_1 = 1$ and $\beta_\ell > \beta_{\ell+1}$ thereafter, ie., the position multiplier goes down top to bottom. The standard assumption is that the actual α of bidder i in position j is *separable* [7], that is, it is the product of the bidder's ctr α_i and the position multiplier β_j: if bidder i is placed at position j then she receives a click with probability $\alpha_{i,j} = \alpha_i \beta_j$. The value of v_i is known only to bidder i while all the other parameters are publicly known.

The ad placement is determined by an *auction.* Advertiser i specifies a bid b_i which is the maximum they wish to pay. The rules of the auction determine the *ordering* of the k chosen ads in these positions, as well as their *pricing.* We describe two well known auctions.

GSP. The most natural ordering is to sort by decreasing bid, but that does not take into account the quality of ads and their suitability to users. The Generalized Second Price (GSP) mechanism ranks the bidders by $b_i\alpha_i$. Wlog assume that bidder i is assigned to position i. The price that the bidder at position i pays per click is

$$P_j^{GSP} = \frac{b_{i+1}\alpha_{i+1}}{\alpha_i}.$$

This is the ordering and pricing currently in use by search engines like Yahoo! and Google. ■

VCG. There is implementation of the well known Vickrey-Clarke-Groves (VCG) mechanism [6,1,3]. This mechanism ranks the bidders by $b_i\alpha_i$, which can be thought of as the expected advertiser value if $b_i = v_i$. Wlog assume that bidder i is assigned to position i. The VCG allocation maximizes the *social welfare*, which is the sum of the bidders' expected value, i.e., $\sum_{i \in N} v_i \alpha_i \beta_i$. The VCG mechanism charges each bidder the total value lost to other bidders caused by her presence in the auction. The VCG mechanism has the property that the bidders' dominant strategy is to bid their true value, i.e., $b_i = v_i$.[1] Formally, the VCG price for position j is [5,4]

$$P_j^{VCG} = \sum_{i>j} \frac{b_i \alpha_i (\beta_{i-1} - \beta_i)}{\alpha_j \beta_j}.$$

Note that $\beta_j = \sum_{i>j} (\beta_{i-1} - \beta_i)$ and therefore $P_j^{VCG} \leq b_{j+1}\alpha_{j+1}/\alpha_j$. ■

Ordered ad auctions described above occur in sponsored search where the keywords that match the search query entered by the user determines the advertisers in the auction. Such ordered ads are also used in content ads where keywords that suitably match the content on a web page determine the advertisers. More generally, various properties of the users as well as the content can jointly determine the pool of advertisers in the auction.

As described above, advertisers are charged only for clicks. In alternative models, advertisers may pay for just appearing (impressions) or only if they *acquire* users according to a suitable definition of "acquire" (users buy products or spend time at the advertiser's website familiarizing themselves with the products, etc).

There is an overview of sponsored search auctions in [20], and a nice introduction to the area of mechanism design [21] which applies to Internet ad auctions. Several workshops, conferences and meetings address Internet ad auction problems. Recent plenary talk [22] discusses some of the algorithmic challenges in content ads.

3 Some Directions

This paper does not describe the issues involved in any detail. Instead, we prefer to present certain directions for research thought. Most of these problems have to be formalized more precisely and depending on the creative approaches taken, will lead to many different research problems.

3.1 Game Theory of GSP in Practice

Since GSP is the most widely used mechanism in practice for ordered ad auctions, it is worthwhile to understand its strengths. This is typically done using game theory.

A commonly accepted *utility* of bidder i at position j is:

$$u_i(j) = \alpha_i \beta_j (v_i - p_j),$$

where p_j is the price per click and is a function of all the bids. For advertisers with this utility, GSP mechanism is not truthful [7], ie., they can gain some by not revealing their true values as bids. In contrast, VCG mechanism is truthful. Still,

[1] A *dominant strategy* is a strategy that a bidder always prefers regardless of the other players' strategies. A mechanism is said to be *truthful* if revealing the true valuation is a dominant strategy for every bidder.

recent papers [4,6,7] show that there exists an *equilibrium* of GSP that is identical in prices and positions to the truthful VCG equilibrium. This characterization gives a nice way to understand the properties of GSP.

Problem 1. In practice, the GSP mechanism is implemented with certain tweaks, for example, each advertiser has an advertiser-specific *minimum price, or certain positions have a set reserve price, or more generally, each advertiser and position pair may have a minimum price. What are natural variations of GSP with these tweaks and what are the equilibrium properties of the resulting auctions?*

Recently, the tweak of just adding advertiser-specific minimum price was studied, and the authors show that GSP has a nice equilibrium even with this tweak [14]. The proof structure in [4,6,7] does not work and a significantly new proof approach was developed in [14]. In presence of the other tweaks stated in the problem above, even modifying GSP suitably may present challenges.

3.2 Multiparty Modeling

Let us examine the commodities being sold: these are clicks at various positions. To determine the best allocation, we need to understand the value of these commodities. The probability of a click on an ad depends on the ad being shown, the user seeing the ad provided it is shown, and then conditioned on that, the user clicking the ad. Of these, the last two are user-dependent. Therefore, the game we study has in fact three parties: users, advertisers and ad providers. So far, we assumed that the probability of user clicking on ad i at position j is $\alpha_{i,j} = \alpha_i \beta_j$, *independent* of other ads. In general, this is unrealistic [2] and it exogenizes the role of the user. A more principled approach would be to assume that the user is not strategic, but model the behavior of the user, and conditioned on this model, study the game between the advertisers and the ad provider.

More precisely, consider the following *markov model* for user behavior.

Problem 2. User scans the positions from the top. Consider position i with ad j. User chooses to click on the ad with probability dependent on the ad j, say $p(j)$, and chooses to scan down the list with probability $q(i, j)$ dependent on the ad seen as well as current position. Assume this markov user model. Determine an allocation of ads to slot with maximum total expected value.

One such an allocation is found, we can use pricing as determined by the VCG mechanism to get a truthful auction. See [9] for a development of this approach, for a restricted model when $q(i, j)$ is only a function of ad j. The rationale behind the model is that the suitability of the ad to their task as well as the fatigue of exploring many of them determines the likelihood of an user continuing to scan the ad list.

A more generally scenario emerges in content ads as described below.

Problem 3. We have the content provider *(say a blog writer) who has some flexibility in choosing the story they wish to tell as well as the words to describe their story (for example, draw analogies, use quotations, etc). The content provider* *will (a) choose the story that is likely to generate large number of* readers*, (b) choose the words which are likely to be used by the* ad provider *to target ads (c) indirectly influence the choice of quality ads shown to the readers, and thereby (c) generate most revenue for themselves from the* advertisers*. Thus, the content provider is a strategic player. Formalize and study mechanism design in this world of four players, ie., content providers, readers, advertisers, and ad providers.*

3.3 Optimal Mechanism

Consider the problem of designing an "optimal" mechanism for ordered ad auctions. We will consider the Bayesian case.

Problem 4. Say each of the n *bidders has their value in* $[0, 1]$ *drawn independently randomly from a distribution* F*. There are* k *ordered positions with decreasing position-dependent click-through rates and assume separability of ctr's. What is the optimal mechanism, that is, a mechanism that maximizes the expected profit?*

If $k = 1$, the optimal Bayesian mechanism was shown by Meyerson [15], and it is a VCG auction that involves setting a reserve price for the position, and selling only if the instance of the values has one that exceeds the reserve price. In the problem above, we are interested in extending it to k position as suitable for ordered ad auctions. See [16] for a discussion of Bayesian and worst case optimal mechanism design.

3.4 Target Size Estimation

Early in the life of an ad is the step of targeting ads. In order to do that, often, advertisers need to estimate the size of their target group, say based on demographics. This is computed typically from data accumulated via surveys and other means which are expensive and do not have complete coverage. To address this, we abstract the following.

Problem 5. We have d *attributes which are say* hierarchical[2]*. We are given several* $d + 1$ *tuples for preprocessing, where the first* d *tuples are the attributes and the* $d + 1$*th attribute is the size. Each query specifies a* d *tuple of attributes and the output is an estimate of its size based on some model of data and size distribution. Devise a theoretical way to address the estimate quality vs complexity of estimation in this problem.*

A natural model is to assume that the sizes are uniformly distributed within the region specified by the d attributes in each tuple. There are heuristic ways to project this model to get size estimate for query region, but a theoretically sound approach that quantifies the information complexity vs accuracy tradeoff will have impact.

[2] For example, *time* is hierarchical; it is represented as a tree with root being the year, its children being the months, their children being partial weeks, down to days, hours and so on. The attribute *salary* may not be hierarchical, since $[50k, 100k]$ and $[75k, 150k]$ may be valid ranges.

3.5 Mechanisms for Heterogeneous Ads

Not all ads are equal. Here is a basic problem.

Problem 6. We have ads of Types I and T. We can either put one ad of Type I or two ordered ads of type T in the space provided. Design a GSP like mechanism for running this auction. In particular, can you formalize and satisfy the property that an advertiser pays the smallest amount needed to obtain their allotment?

3.6 Mechanisms for Heterogeneous Utilities

There has been a lot of research that assumes some silo or the other of advertisers. In practice, we get a mixture of advertisers. Some care about impressions, others care about clicks. Some care about profit maximization and others about appearing in some position, no matter the cost. Classical game theory provides insights into mechanisms for specific classes of utility functions, but in practice, one gets a mixture of utility functions.

Problem 7. Consider a collection of advertisers with mixed utilities in terms of impressions/clicks/acquisition, profit/position or other suitable parameters. Design a truthful, intuitive mechanism for such a collection and characterize the equilibrium behavior and advertiser dynamics.

The authors in [18], propose a general model and design a stable matching based mechanism which is truthful for a rich set of advertisers. They also propose an efficient algorithm to implement the mechanism. Alternative models and mechanisms, with richer mix of advertisers, will be of great interest.

3.7 Ad Effectiveness Estimation

An important task is to estimate the effectiveness of ads. While clicks are a natural measure, in some cases such as in branding or display or offline media ads, we need indirect ways to measure effectiveness, say in terms of increased traffic at sites of advertisers, or call backs or innovative methods such as 2d barcodes [23]. Motivated by this, we abstract the following.

Problem 8. We are given two disjoint sets of target groups C and A. Target groups in A are shown ads during time period $[0, T]$ and target groups in C are not shown ads. We have time series measurement of activity that pertains to the ad of each target group in C and A (eg, the measurement is the number of visits to the advertiser's website for each time instant) measured over time $(\ldots, 0, \ldots, T, \ldots)$. The task is to design a concise measure of the change in measurements in A after the ad was shown wrt prior time when the ad was not shown and also wrt other target groups in C during the entire time (since they were not shown ads). Assume the richest model of dependence among the different timeseries of measurements that makes the problem computationally feasible.

3.8 Inferring Profiles: Graph Learning

Many inference problems for ad targeting and elsewhere can be abstracted as follows. There are entities with associated metadata and links. Some of the entities are labeled, and we need to infer labels for the others. We can infer the label for an entity using its metadata. But, instead, we wish to infer the label based on just the link structure between the entities and the known labeling. For example, say each entity is a blog profile, and the label of our interest is the *age* of the blogger. Some of the blog profiles have age value entered (truthful or not) and others have missing age value. It will be beneficial to infer some age value for the unlabeled bloggers. This will be useful for targeting ads, for example, if the advertiser wants their ad shown only to those above a certain age. The problem we state below models this (and other cases with labels like say gender or interest in rock music, etc.) informally.

Problem 9. The input is a graph $G = (V, E)$ with some nodes $v \in V$ with a label $L(v)$. Output is $L(u)$ for all $u \in V$.

See [12] for a random walk approach and see [10] for the Machine Learning approach to this problem. More remains to be done, in particular, in novel approaches and re-formulations of this problem. We propose the following.

Problem 10. Initially we are given a graph H with a labeling $L : V(H) \to [k]$. Then, for each edge e of H, and independently, we remove e with probability $p_{i,j}$, where i and j are the labels of the end-points of e. Now, we present the resulting graph G, sans the labels, to the algorithm. Design a decoding algorithm which will produce a maximum likelihood labeling of the nodes of G, with high probability.

If H were the complete graph on n nodes, results in [11] apply and such decoding is possible. For general H, preliminary results on this problem are in [19]. Results for special cases of H or improved approximate decoding results are of great interest; also, the labeling could be a probability distribution over the label set, and algorithms for such cases are of interest.

3.9 Reservation Auction

We have so far considered spot auctions only, that is, the auction takes place at the instant when the commodity is available. In certain cases, typically with display ads, advertisers need to be able *reserve* ad spots ahead of time. Motivated by this, we abstract the following problem.

Problem 11. There are ad positions available at each time instant. Bids arrive online; each bid is for a subset of ad positions at some specific instant in the future. Bids must be accepted or rejected when they arrive. Assume reservations once accepted may be rejected later, for a bump fee. Design a mechanism for making these decisions online and still result in an allocation which is comparable to offline mechanisms in revenue and welfare, together with desirable game-theoretic properties.

For the case when each bidder requires a single slot only, some positive results appear in [26].

3.10 Budget Optimization and Bidding

Sometimes it is required to view ordered ad auctions from the viewpoint of advertisers. They need tools to strategize and develop a bidding method with multiple keywords and targeting groups.

Problem 12. Consider a particular advertiser. We are given budget U, a keyword-query graph G and distribution of how many clicks we get at what cost for each bid value on each query in the GSP auction, assuming all other parameters remains fixed. The goal is to output a bid for each keyword so that the expected number of clicks obtained by the advertiser is maximized, subject to the condition that budget is not exceeded. Also, characterize the equilibrium if all advertisers follow the click-maximizing strategy in GSP auction.

A bidding strategy (bidding uniformly on all keywords) and its variants are explored in [24]. Certain basic dynamics have been understood [25], but more remains to be done.

4 Concluding Remarks

This writeup provide some insight into Internet ad auctions and directions for further research. Internet ad auctions face algorithmic challenges, and will remain a rich area of research for some time. It is also likely to be useful research area because large scale auction systems exist and are successful; new research ideas can modify them or rework them completely with reasonable effort, and have the potential for great impact.

Acknowledgements

I sincerely thank my colleagues at Google.

References

1. Clarke, E.: Multipart pricing of public goods. Public Choice 11, 17–33 (1971)
2. Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: WSDM 2008: Proceedings of the international conference on Web search and web data mining, pp. 87–94. ACM Press, New York (2008)
3. Groves, T.: Incentives in teams. Econometrica 41(4), 617–631 (1973)
4. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In: Second workshop on sponsored search auctions (2006)

5. Varian, H.: Position auctions. International Journal of Industrial Organization 25(6), 1163–1178 (2007)
6. Vickrey, W.: Counterspeculation, auctions and competitive-sealed tenders. Finance 16(1), 8–37 (1961)
7. Aggarwal, G., Goel, A., Motwani, R.: Truthful auctions for pricing search keywords. In: ACM Conference on Electronic Commerce (EC) (2006)
8. Feldman, J., Muthukrishnan, S.: Algorithmic methods for sponsored search. In: Liu, Z., Xia, C. (eds.) Performance Evaluation and Modelling. Springer, Heidelberg (to appear, 2008)
9. Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M.: Sponsored search auctions with Markovian users (submitted, 2008)
10. Graph Labelling Workshop - ECML/PKDD (2008), `http://graphlab.lip6.fr/pmwiki.php`
11. McSherry, F.: Spectral partitioning of random graphs. In: Proceedings of the 42th Annual IEEE Symposium on Foundations of Computer Science (2001)
12. Baluja, S., Seth, R., Sivakumar, D., Jing, Y., Yagnik, J., Kumar, S., Ravichandran, D., Aly, M.: Video Suggestion and Discovery for YouTube: Taking Random Walks Through the View Graph WWW (2008)
13. Aggarwal, G., Ailon, N., Constantin, F., Even-Dar, E., Feldman, J., Frahling, G., Henzinger, M., Muthukrishnan, S., Nisan, N., Pal, M., Sandler, M., Siridopoulos, A.: Theory research at Google. In: SIGACT News (to appear, 2008)
14. Even-Dar, E., Feldman, J., Mansour, Y., Muthukrishnan, S.: Sponsored search with bidder-specific minimum prices. In: [13] (2008)
15. Meyerson, R.: Optimal auction design. Mathematics of Operations Research 6, 58–73 (1981)
16. Hartline, J.: Lectures on optimal mechanism design. In: [17], `http://www.ece.northwestern.edu/~hartline/omd.pdf`
17. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.): Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
18. Aggarwal, G., Muthukrishnan, S., Pal, D., Pal, M.: General Auction Mechanism for Search Advertising. (manuscript, 2007), `http://www.cs.uwaterloo.ca/ dpal/papers/auction/stable-matching-auctions.pdf`
19. Feldman, J., Muthukrishnan, S., Sidiropoulos, A.: Graph labeling in the planted graph model. (manuscript, 2008) See [13]
20. Lahaie, S., Pennock, D., Saberi, A., Vohra Sponsored, R.: search auctions. In: [17] (2007)
21. Nisan, N.: Introduction to Mechanism Design (for Computer Scientists). In: [17] (2007)
22. Broder, A.: Computational advertising. In: Proc. ACM-SIAM SODA, p. 992 (2008)
23. Print ads 2D barcodes, `http://www.google.com/adwords/printads/ads/barcode/`
24. Feldman, J., Muthukrishnan, S., Pal, M., Stein, C.: Budget optimization in search-based advertising auctions. In: Proc. ACM EC (2007)
25. Borgs, C., Chayes, J., Immorlica, N., Jain, K., Etesami, O., Mahdian, M.: Dynamics of bid optimization in online advertisement auctions. In: Proc. ACM EC (2007)
26. Constantin, F., Feldman, J., Muthukrishnan, S., Pal, M.: Online Ad Slotting With Cancellations. (manuscript, 2008) See [13]