
Online Advertising: Auctions, Targeting and Tracking

Dr Enrico H. Gerding, Lecturer at the University of Southampton

8 May, 2015

Online display advertising has been rapidly moving from traditional bilateral contracts, where advertisers have agreements with individual websites offering advertising space, to *real-time bidding* where advertisers bid in auctions for each individual user visiting a website. Companies running these auctions are called *ad exchanges* and they are based on the same principles as financial exchanges. The main ones are DoubleClick (owned by Google), Yahoo! Ad Exchange (previously Right-Media), AdECN (owned by Microsoft) and, more recently, Facebook Exchange (FBX). As with financial exchanges, the aim is to match demand and supply. In the case of ad exchanges, the supply is in the form of available space on web sites from so-called *publishers*, whereas the demand is generated by advertisers. Specifically, each user visit to a website is called an *impression*, and each impression is matched with an advertiser using a special type of auction (the process is explained in more detail below).

Advertisers cannot bid directly in these auctions. Rather, this is done through intermediaries, such as *ad networks* (who serve both advertisers and publishers, and also offer an internal matching service but can, at the same time, participate in the large ad exchanges) and *demand-side platforms* (who only serve advertisers and typically participate across multiple ad exchange auctions as well as in the smaller ad networks). These companies run ad campaigns on behalf of multiple clients, i.e. the advertisers, and they need to decide which client to bid for and how much to bid, for each impression (i.e. user). These bids are based on many factors, such as the size and position of the ad on the web site, but also on the user's *behavioural profile* which indicates a user's potential

interests and can be used to target the ads. This profile, in turn, is determined by tracking the user across multiple websites, and using machine learning and data mining algorithms to try and predict a user's *intent*.

In what follows, I will explain the auction process in more detail and describe how tracking works.

Ad Exchange Auctions

As soon as a user visits a publisher's website, the user's impression is auctioned off in an extremely fast ad auction, where multiple intermediaries bid on behalf of their clients, the advertisers. Once the winner is determined, the actual ad content (e.g. a banner or, increasingly, a video) is sent from the ad server and embedded in the website for the user to see. The entire process (from the user arriving to the website to the ad being served) should take less than 100ms to avoid a delay in loading the website. Typically there are multiple ads on a single website, resulting in multiple auctions simultaneously at possibly different ad exchanges. As a result, every part of this complex process has to be extremely efficient.

In terms of the auction itself, most ad exchanges (with the exception of FBX which runs a more advanced auction) use an interesting format called the second-price sealed bid or *Vickrey* auction, named after its inventor, the economist and nobel laureate William Vickrey. In this auction, all bidders bid once (without seeing the bids of others), and the highest bidder wins, unless the highest bid is below a reserve price (a minimum price set by the publisher), in which case there is no winner. The interesting part is the payment: the winners pay the *second-highest*

bid or the reserve price, whichever is highest.

The pricing rule ensures that there is no need for a bidder to speculate about the other bidders in the auction. To understand this better, suppose that a bidder, let's call her Alice, has a maximum willingness to pay or *value* for this impression of £1.00 (we will discuss later how this value is determined). In that case, there is no incentive to bid less than £1.00 since her payment is not affected by her bid, and this only determines the chances of winning. For example, if the highest *other* bid is £0.90, Alice's payment will be £0.90 if she wins, regardless of her bid. At the same time, there is no incentive to bid higher than £1.00. This does increase the chance of winning, but only in situations where the highest other bid is above £1.00, in which case the payment will be more than £1.00. For example, suppose that the highest other bid is £1.05. By bidding £1.00 Alice will lose the auction. With hindsight, she could have won the auction by bidding (say) £1.10. However, her payment would then have been £1.05 which is more than she was willing to pay. Therefore, even with hindsight, it is always optimal to bid the (true) value in the auction, regardless what other bidders are doing. Hence this auction is also called *truthful*.

In contrast, if winners were paying their own bid, they would want to speculate, i.e. try to guess the highest other bid, and bid just above it. This leads to more complex bidding strategies, which is undesirable in ad auctions due to the high speed required, but also leads to fluctuations in prices and inefficient markets (because often guesses turn out to be wrong). Hence the Vickrey auction is an effective auction in this setting. In practice, there are typically multiple ads on a single page (each having their own auctions), and clients have certain budgets and set performance targets (e.g. number of daily impressions) which need to be met, which complicate the bidding strategy, but the auction format still reduces the amount of speculation needed.

Targeting and Tracking

An important component of the auction process is to determine the monetary *value* of an impression, i.e. the value of showing an ad to a specific user visiting a specific page. Often this value is the estimated *return on investment* (ROI) which can be the immediate return, i.e. a user clicking on the ad and completing a purchase, or it can (also) be long-term increase in sales. Typically this value is very small for an individual impression (typically less than £0.01), but

the number of daily impressions for a single advertiser can be very high (typically in the tens of thousands).

To accurately estimate ROI, it is important to have information about the user's interests or *behavioural profile*. For example, a user who has recently searched for gardening tools on Amazon has a high value to other companies selling gardening tools. There are many companies specialising in this so-called *re-targeting*, i.e. showing ads about things that users have shown explicit interest. More sophisticated approaches use techniques such as *recommender systems*, which predict other items a user may be interested in, given their past browsing behaviour.

As we can see, there is a clear benefit for advertisers to keep track of a user's browsing and search behaviour. This is done by data collection services, such as DoubleClick and BlueKai, typically using so-called *third-party cookies*. They, in turn, sell this information to ad exchanges and advertising companies. Cookies are small files on a user's computer which can only be accessed by the server (i.e. the website) placing the cookie. Third parties (e.g. DoubleClick) will embed some web code in numerous publisher's pages, such as Amazon or Daily Mail, allowing them to place unique cookies. When the same user visits another page with the third party code, this code will be able to read the cookie from the user's computer, and identify it as being the same user. This allows the third party to track a user in an *pseudonymous* fashion. Note that other approaches have been developed that can track users without using cookies or similar technologies,¹ but these are typically less reliable.

Conclusion

The ad auction process for display advertising is clearly a complex one with many stakeholders, and it is still very much evolving. Specifically, the *mobile advertising* market is growing rapidly since it has unique targeting opportunities such as geo-location, but is also faced with challenges (not many ads can be shown, and they are less likely to be clicked). Similarly, advertising on social networking allows for more accurate profiling based on someone's friends, photos, posts, etc, and it is a growing area. Still, many of the algorithms currently used for profiling, such as re-targeting, are very simplistic, and there remains significant opportunity to investigate more intelligent algorithms.

¹See e.g. <https://panopticlick.eff.org/>.