PACIS 2004 Proceedings

Pacific Asia Conference on Information Systems (PACIS)

December 2004

Computational Evaluation of Alternative Paid Placement Mechanisms in Search Engines

Juan Feng
University of Florida

Hemant Bhargava
University of California at Davis

David Pennock
Yahoo Research Labs

Computational Evaluation of Alternative Paid Placement Mechanisms in Search Engines

Juan Feng
University of Florida
jane.feng@cba.ufl.edu

Hemant K. Bhargava
University of California at Davis
bhargava@computer.org

David M. Pennock
Yahoo Research Labs
david.pennock@overture.com

Abstract

The practice of paid placement advertising—where advertisers pay a fee to appear alongside particular Web search results—is now one of the largest and fastest growing source of revenue for Web search engines. This article studies the allocation of placement slots. The payperclick standard for paid placement makes the problem challenging because it is not optimal to simply allocate slots to the highest bidders. We model several allocation strategies, including stylized versions of those used by Overture and Google, the two biggest brokers in paid placement, along with two alternative designs, and compare their performance via computational experiments. All mechanisms perform better when providers' willingness to pay for paid placement is positively correlated with their true relevance. Ranking providers based on the product of clickthrough rate and bid price fares well broadly, while ranking purely by bid price performs well when the content providers' relevance is positively correlated with their bid. Ranking by bid improves greatly with editorial filtering. Search engine's placement revenue decreases when users' attention is significantly lower for lowerranked listings, emphasizing the need to develop better user interfaces and control features. Due to the tradeoff between direct revenue increases and indirect revenue losses (due to consumer defection), the search engine must carefully choose the total number of paid slots. We also study how the rank allocations for each mechanism change over time as the search engine obtains more information about clickthroughs at each rank. We propose a rank revision strategy that weights clicks on lower ranked items more than clicks on higher ranked items. This method is shown to converge to the optimal (maximum revenue) ordering faster and more consistently than other methods.

Keywords: searching engine, ranking, simulation, dynamic, revenue

1 Introduction

Internet search engines index billions of Web pages and employ information retrieval algorithms to display, in response to a user's query, links to a subset of Web pages deemed relevant to the query. These pages might represent commercial firms selling goods or services, information sites, government entities, and so forth. The most recent scientific study (Lawrence and Giles 1999) estimated the publicly indexable Web at 800 million pages, containing 6 terabytes of text data on 2.8 million servers; today, Internet statistics indicate that the publicly indexable Web contains billions of pages, served by over 40 million Web servers.[1] Due to the vast amount of information on the Web, search engines act as an information gateway to many search and decisionmaking tasks. Industry surveys indicate the importance of search engines: more than 50% of Web users visit a search engine every few days, the leading search engine (Google) gets over 250 million search requests each day, over 13% of traffic to commercial sites was generated by search engines, and over 40% of product searches on the Web were initiated via search engines.[2] For the purpose of this paper, the term search engine encompasses various applications of these indexingretrieval technologies, including pure Web search engines (e.g., Google), information portals with search functionality (e.g., Yahoo!), metasearch engines (e.g., Metacrawler), niche search engines (e.g., CiteSeer (Bollacker, Lawrence, and Giles 2000)), and comparison shopping engines (e.g., mySimon, Shopping.com).

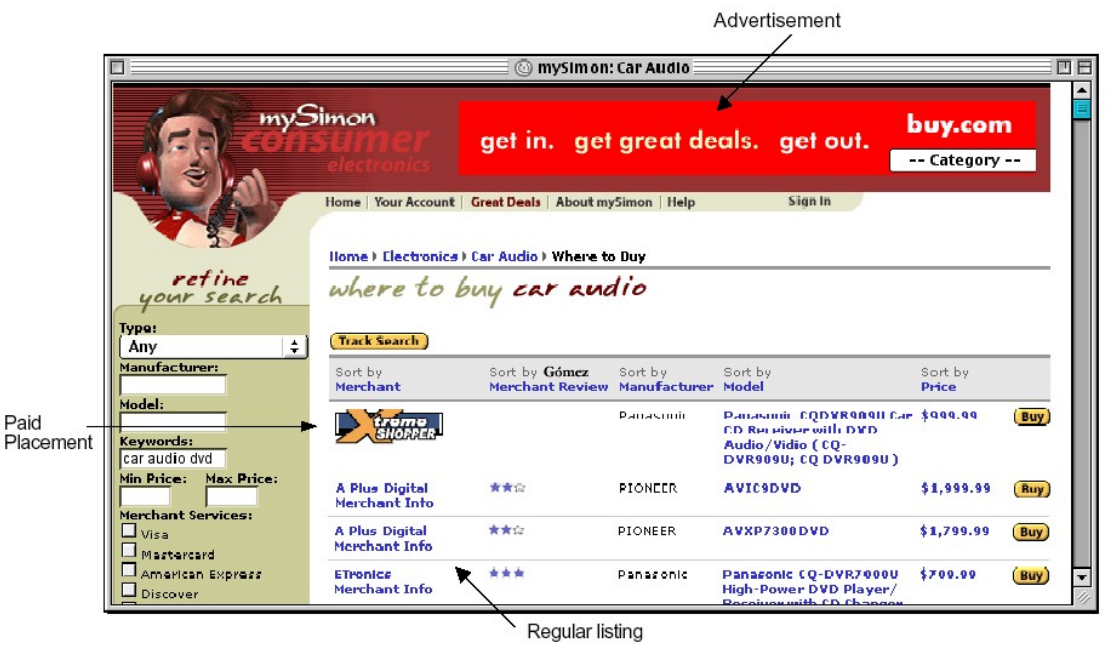

Figure 1: Paid placement advertising, regular algorithmic listings, and banner advertising in a comparison shopping engine. The top listing and graphical icon increase the likelihood that the paid placement listing will be followed up.

[2] See the following sources on the Web:
SearchEngineWatch: http://www.searchenginewatch.com/reports/article.php/2156461
Internet news: http://www.internetnews.com/IAR/article.php/2108921
Overture: http://www.content.overture.com/d/USm/ays/ays keystats.jhtml

Due to the critical influence of search engines on Web users’ actions, many commercial firms have realized the importance of gaining a high position on the search results for specific queries. As noted earlier, search engines initiate a significant amount of traffic to commercial web sites. Research has shown that good placement on a search page leads to high traffic, and eventually an increased financial payoff (Cotlier 2001, Cotriss 2002). In fact, entire niche industries exist touting services to boost a Web page's ranking on the popular search engines, in part by reverse engineering the search engines' information retrieval algorithms. In recent years, search engines have begun practicing what is called paid placement or pay for performance advertising: inserting links from companies that pay a fee for prominent position alongside results for a particular search query. These companies may be content providers, vendors of goods or services, or other advertisers interested in attracting Web users. Paid placement advertising has become an important and fastgrowing revenue source for Internet search engines (Reinhardt 2003, Dolbeck 2003). Usually, paid listings are shown on top of, or to the side of, any standard unpaid search results (also called algorithmic results). In general, paid placement listings are explicitly marked as sponsored results or advertising, and the FTC has advised search engines to appropriately disclose paid links (Tantono, Selby, Bagner, and Sonu 2002). Paid placement is widespread in all types of search engines. Figure 1 displays a screen shot from a comparison shopping site that includes algorithmic placement, paid placement, as well as more traditional banner advertising.

The study of paid placement strategies is crucial in understanding the future design, quality levels, and market structure of Web search engines. Bhargava and Feng (2001, 2002) discuss the practice of paid placement bias in Internet search engines, and analyze the tradeoff between paid placement revenue and the potential loss of market demand because of bias. Hoffman and Novak (2002) discuss the trend in Internet advertising towards CPC (cost per click) pricing rather than the traditional CPM (cost per thousand impressions) model. This move is meaningful because of the uncertainty in the benefit to the advertiser from showing an advertisement; the Web's capability for two way communication makes it possible to measure users' attention towards an ad, hence advertisers want payment to be linked to performance. Bhargava and Sundaresan (2002) have shown that such performancebased pricing is an efficient solution when there is uncertainty about a product's performance. Asdemir, Kumar, and Jacob (2002) also study the pricing choice of CPM and CPC, while Hu (2003) uses contract theory to show that performancebased pricing models can give the publisher proper incentives to improve the effectiveness of advertising campaigns. Weber and Zheng (2003) study the implementation of paid placement strategies, and find that the revenue maximizing search engine design bases rankings on a weighted average of relative quality performance and bid amount. Rolland and Patterson (2003) propose a methodology using expert systems to improve the matching between advertisers and web users. Kumar, Dawande, and Mookerjee (2003) study the optimal advertising schedule to maximize the web site (search engine)'s revenue under a hybrid pricing model (based on both the number of impressions of the ad and the number of clicks on the ad).

This paper addresses the question of how a search engine allocates its (usually very few) paid placement slots. For several reasons, search engines should not simply allocate paid positions to the highest bidders. First, irrelevant or offensive results can turn off the search engine's users, resulting in long term revenue losses. Second, because the industry has largely converged on a payperclick mechanism rather than a payperimpression system, there is an inherent incentive on the part of the search engine to display relevant advertisements that are likely to be clicked on. In this paper, we report the results of computational studies that compare the efficacy of common paid placement mechanisms as well as three of our own, based on a model that we believe captures many of the essential dynamics of paid placement. Our analysis is inspired by the practices of the leading paid placement firms: Google and Overture Services. Overture (formerly GoTo.com, recently acquired by Yahoo!), the firm credited with pioneering paidplacement advertising, is not a consumerfacing brand but rather acts as a broker between content providers (advertisers) and information gatekeepers like Yahoo! and MSN. Overture's success has prompted several other companies to adopt similar business models involving paid placement, most prominently Google, which is the leading Internet search engine today. Although our mechanisms are stylized versions of those run by Overture and Google, there are many variables in both companies' systems that we do not or cannot model, including querymatching algorithms, human editorial intervention, non-searchbased advertising, new pricing models (e.g., pay per conversion), marketing efforts, brand awareness, user interfaces, legal controls, strategic alliances, etc. Therefore our results should not be construed as indicators of either company's business outlook.

The rest of this paper is organized as follows. Section 2 studies the question of how a search engine determines which content providers are included in the paid slots and how they are ranked within the slots. We provide a detailed description of the mechanics of paid placement, develop a model of placement revenues, and describe four alternative mechanisms for choosing the allocation of paid slots to advertisers. In section 3, we report results of computational simulations of the equilibrium performance of these four mechanisms. Ranking by the product of bid price times clickthrough rate weakly dominates other mechanisms in all tested regimes, though the mechanism is statistically equal to ranking by bid price in the expected region of positive correlation between willingness to pay and relevance. Editorial filtering helps the basic rankbybid mechanism significantly. Section 4 provides concluding remarks.

2 Paid placement Ranking Mechanisms

The positive correlation between top placement and increased traffic creates significant demand among businesses for top placement on search engines, especially for popular and commercially-relevant search terms. However, since Web users face negative utility if the search engine becomes impartial, most search engines limit the number of paid placement requests they accept. Thus, the paid placement slots are a scarce resource that need to be allocated carefully. In this section, we describe the mechanics of paid placement, develop a model of the search engine's placement revenues, and describe four alternative mechanisms for choosing the allocation of paid slots to advertisers. We develop a simplified model of the situation that we believe captures several of the important aspects of the problem. Of the four mechanisms described here, two are stylized versions of those used by Google and Overture, leading firms in the paid placement industry.

2.1 How does paid placement work?

The predominant mechanism employed in the industry is to auction off space alongside particular search terms. In practice, hundreds of thousands of advertisers compete for positions alongside several millions of search queries. So, for example, travel vendors like Expedia and Orbitz may compete in an auction for the right to appear alongside the result of a user's search on "las vegas travel". Generally auctions are dynamic, meaning that advertisers can change their bids at any time, and a new auction clears for every search query. In this way, advertisers can adapt to changing environments, for example boosting their bids for "Harry Potter" around the release date of the latest book in the series.

In practice, we observe many different approaches for allocating and pricing paid placement slots. Overture screens listings according to both automated and manual editorial policies to control for relevancy, then ranks all qualified advertisers according to how much they are willing to pay per click. For any given query, paid slots are allocated to the top k bidders in order of their bids, where k is the number of paid listings on the search results page. Listings may be shut down if they do not achieve sufficient clickthrough rates. Google's AdWords Select paid placement program also employs automated and manual editorial filtering, though generally relies less on human intervention than Overture's program. Google ranks qualified listings according to the product of the advertiser's bid times the actual clickthrough rate of the listing. Since advertiser payments are per click, boosting higher clickrate listings may help improve overall revenue. Both companies' deployed services include many practical complications and additional variables that are beyond the scope of this paper, beyond the knowledge of the authors, and/or in constant flux; therefore, associating any of our results to these specific companies is not justified. For example, another differentiator that we do not model are possible inexact querymatching algorithms that match listings to search queries that are close to—but not exactly the same as—the search phrase assigned to the listing. Other paid placement brokers (including FindWhat.com and LookSmart.com) use variations on these strategies for identifying, ranking, distributing, pricing, and displaying their paid listings.

2.2 Revenue Model

Consider the search engine's placement revenue for a single search term. Suppose that s listing companies, interested in advertising for this term, compete for k paid slots on the search engine's results page. Let v_j be advertiser j's willingness to pay (WTP) per click for preferential placement on this term. Let α_j represent the "true" relevance score of listing j, encoding how useful the link is to users. Let $f(\alpha_j, v_j)$ denote the joint density function of the variables. The true relevance is unobservable to the search engine and users, but can be approximated over time by the number of clicks received by the listing. However, since users are more likely to click on higherranked items closer to the top of the page, the clickthrough rate for any listing depends on not only its relevance, but also its rank within the paid placement section. To compute the expected clickthrough rate for an item j at position i, our simulation employs an exponentially decaying attention model with factor $\delta > 1$. Formally, we compute the average clickthrough as α_j / δ^{i-1}. Exponential decay of attention is a fairly standard assumption (Breese, Heckerman, and Kadie 1998) that is borne out in practice. We assume that $0 < v_j, \alpha_j < 1$. Define $r : I \rightarrow J$ to be the search engine's ranking function which allocates position i to listing company j. The set J also contains a fictitious null provider, since the search engine may not fill all slots. Let P_i represent the payment for position i. For the clickthrough based mechanisms, let p_i represent the payment per clickthrough at position i, so that in equilibrium $P_i = \sum p_i$ where $j = r(i)$.

2.3 Allocation Mechanisms

We consider four mechanisms. The v ranking mechanism is inspired by Overture's firstprice auction that ranks by bid price, while the $v\,\alpha$ mechanism is inspired by Google's secondprice auction that combines bid price with the measured clickthrough rate. Table 1 summarizes the notation.

2.3.1 v ranking: Highest payers at the top

This mechanism allocates slots based on the listing company's willingness to pay. For a certain keyword, each listing company j makes a bid Bj for payment per click. The listings associated with the highest k bids are displayed, ranked according to their bids. We assume that companies pay what they bid (first price auction). Due to the per click payment format of paid placement, each listing firm's bid is independent of which slot it might be allocated (which influences its expected clickthroughs), hence we can use standard auction theory to derive the equilibrium bidding strategy. Therefore Bj is the expected highest value among the remaining n-1 bidders given that this value is below vj (Klemperer 1999). Hence, the firm with i^{th} highest value makes the i^{th} highest bid. The highest k bidders win and are assigned positions according to their bids, therefore under this mechanism r(i) is the firm with i^{th} highest valuation. The search engine realizes a priceperclick vfor slot i.

This mechanism is meant as a stylized version of Overture's approach. Note however that our formulation ignores many factors present in the commercially deployed mechanism, including editorial control, bid proxy agents, inexact query matching, and less than perfect correspondence between relevance and clickthrough rate, among other factors.

2.3.2 v α ranking: Relevance and bid price jointly determine rank

This mechanism describes a stylized version of Google's approach to preferential placement. Again, every listing company j bids Bj , its willingness to pay (per click). We compute the product Bj αj (where αj approximates the expected clickthrough rate for the listing). The listings associated with the highest k products are displayed, ranked according to this product. The actual price paid by each winner is a variant of the standard secondprice auction, and computed as follows. Let r(i) be the firm that wins slot i. Then, if Br(i) >Br(i+1) then firm r(i) pays Br(i+1); otherwise, it pays the least amount Bi such that Bi αr(i) ≥ Br(i+1) αr(i+1). In other words, the winning firm pays either the bid just below it, or an amount such that its total predicted payment just exceeds that of firm r(i+1). Even though this is not a standard second price auction any more, we make the simplification that the listing firms bid their true value, which is reasonable since the firm does not pay its own winning bid.

2.3.3 α ranking

This mechanism selects the highest kbids and ranks the bidders by their expected clickthrough rates. The placement slots are assigned according to this ranking, and all winners pay the highest rejected bid. Consequently, every listing company bids their true willingness to pay (per click) vj . Thus this is a generalized version of second price auction. Hence pi = vr(k+1) is the $(k+ 1)^{th}$ highest valuation.

2.3.4 Postedprice mechanism

In this mechanism, the search engine sets a reserve price for each of the positions it has for a certain period. Each listing company submits their bid to be listed for that period. The highest k bids are the potential winners, except that the search engine may fill fewer than k slots if certain of the winners do not meet the reserve price constraints, as explained below. For each i if the ith highest bid is higher than the reserve price for the ith position, the listing company who submits this bid is admitted for the i'th position and pays the reserve price; otherwise, it is compared to the $(i + 1)^{th}$ reserve price. If it is higher than the $(i + 1)^{th}$ reserve price, it is admitted and pays the $(i+ 1)^{th}$ reserve price, but this way the total available positions for paid placement will also be reduced by 1, and so on.

In setting up the auction, the search engine determines the k reserve prices by computing the expected revenue potential for each of the k positions. In our simulation, for every sample correlation (from 1 to 1) we generate s pairs of "α"s and "v"s, calculate the average order statistics for the product αv over 100 runs, and use this as the reserve price. Note that this procedure imposes a quality control for the winner determination rule: those listings with higher relevance score are more likely to have a greater revenue potential (because they are most likely to generate traffic), thus they are the ones which are most able to pay the reserve price.

3 Revenue Comparisons

This section describes our simulation experiments to compare the steadystate performance of alternative mechanisms. Section 3.1 motivates the design of the experiments, the rest of the section provides the results. Appendix A.1 summarizes the design of the simulation code.

3.1 Experiment Design

Intuitively, because of the "pay per click" format of paid placement pricing, the comparative revenue performance of the different mechanisms is influenced by the extent of correlation between the listing firms' willingness to pay for paid placement, and their ability to attract traffic (i.e., their relevance to the search term). While this extent of correlation has an impact on the relative outcomes of the different mechanisms, it is unobservable to the search engine, hence the choice of mechanism requires a good understanding of the relationship between correlation and relative performance. Our computational analysis accounts for this requirement by treating v and α as joint random variables, and by systematically varying the degree of correlation ρ between these variables over multiple simulation runs. We normalize the variables to lie in the interval [0,1]. We model $f(\alpha j ,vj)$ as a truncated bivariate normal distribution between 0 and 1. The means are 0.5 and the covariance can vary between 0.167 and 0.167, making the correlation between these two factors vary from 1 to 1. We expect the region of positive correlation to be the most realistic, because advertisers generally have an incentive to request relevant placement, in order to attract traffic from genuinely interested users. Moreover, we expect that, in the long run, advertisers themselves will withdraw irrelevant or ineffective advertisements, retaining only those that genuinely attract interested consumers and are cost effective.

Similarly, the performance of each mechanism if affected by the intensity of competition for the paid placement slots. Should we expect the effect of more (or less) competition to be identical across all mechanisms? Intuitively, we believe the effect should be different, because of the differences in how each mechanism considers the anticipated clickthrough rate for the bidding firms. We also note that the intensity of competition is affected by two variables, the number of interested advertisers (not controllable by the search engine) and the number of slots made available for paid placement (which is set by the search engine). Hence our analysis covers variations in both the number of available slots k and the number of firms bidding for paid placements.

Finally, we expect the performance of different mechanisms to be impacted differentially by the extent of attention decay over the placement slots, hence we vary the parameter δ across our simulations.

3.2 Impact of Correlation between WTP and Relevance

First we compute the relative sensitivity of each mechanism to the correlation between an advertiser's WTP and relevance. The following result is derived with $k = 5$ available slots, $s = 15$ advertisers who desire paid placement, and an attention decay factor $\delta = 2$. For each sample value of correlation in [1,1], we compute the average revenue over 200 runs. Each run consists of a random draw from $f(\alpha j, vj)$ for each of the s content provider's WTP and relevance. We present the findings for the four different mechanisms discussed above.

Finding 1 For every mechanism, the revenue earned is increasing in the correlation between the

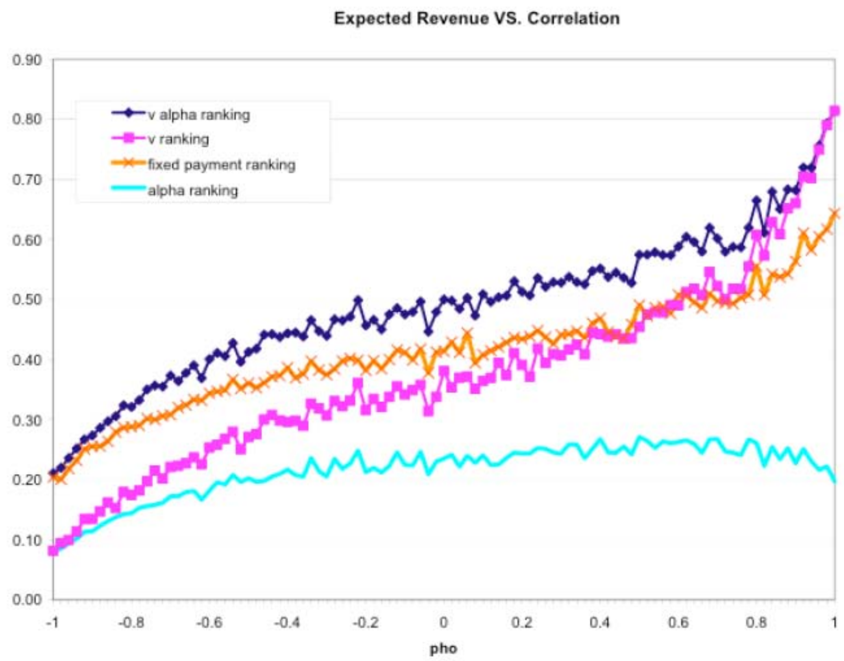

Figure 2: Average revenue versus correlation for the four tested ranking strategies.

content provider's relevance score and its willingness to pay. The effect is most pronounced for the v ranking mechanism.

Finding 1 is intuitive, since a greater correlation results in picking more relevant links, as indicated in Figure 2. The v ranking mechanim, because it ignores relevance, benefits the most from an increase in the correlation. Figure 2 also displays the dominance of the v and v α mechanisms in the region of high correlation. When v and α are negatively correlated (which we expect to be rare in practice), v α ranking performs significantly better than v ranking. In the expected region of positive correlation v α ranking remains marginally better, though the difference is not significant at significance level 0.05.

Finding 2 v α ranking weakly dominates the other three mechanisms, and strongly dominates in the region of negative correlation. When the correlation is positive, the difference between the v and v α mechanisms is not significant.

Note that the v ranking mechanism displays inferior performance in the region of negative correlation. This underscores the need for stronger editorial control for relevance. Interestingly, this is consistent with industry practice: Overture, which ranks by v, does generally expend more resources on human editorial control than Google, which employs a form of v α ranking.

3.3 Impact of Attention Decay Factor

Now we examine how the performance of the different mechanisms is influenced by δ, the difference in the attention that a certain listing item can get in different positions. Figures 3(a) and

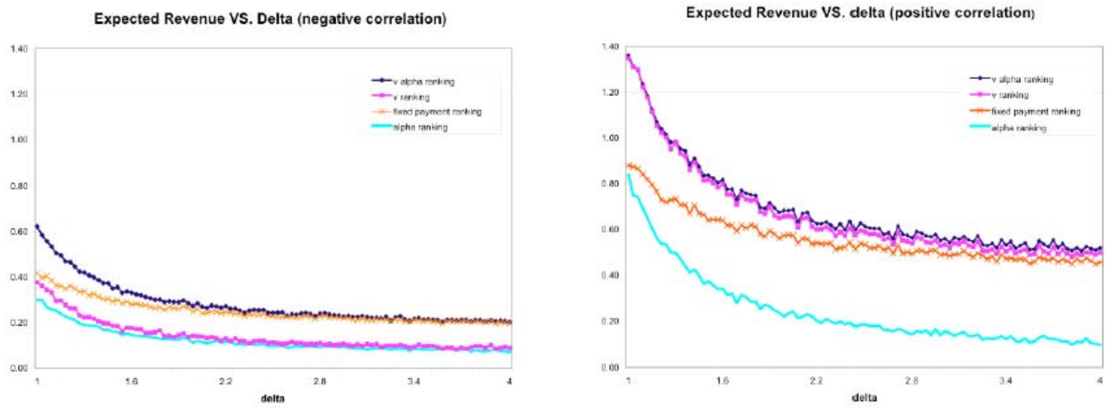

(a) negative correlation (b) positive correlation

Figure 3: Revenue versus the attention decay parameter δ, under negative and positive correlation between relevance and willingness to pay.

3(b) show the performance of each mechanism with respect to δ, when the correlation between paid listing firm's relevance and WTP is strongly negative (left subfigure) and strongly positive (right subfigure), respectively.

Finding 3 The revenue generated from paid placement decreases as the attention decay factor δ increases.

This result is intuitive. As δ increases, the lower rank positions become less attractive thus will generate lower revenues. We note that the search engine can exert some control over the decay factor, for example, by designing a better user interface that maintains attention over a larger subset of paid listings. Still, as the limits of improved user interfaces are reached, a fundamental limitation on human attention is unavoidable. Figure 3 also indicates that the posted price mechanism converges to the performance of the v α mechanism when there is significant attention decay.

3.4 Impact of Demand for Paid Placement

The search engine's potential for placement revenues depends on the overall demand for preferential placement. Intuitively, when more providers compete for paid placement, this will increase the market clearing price and the search engine's revenues. Our analysis extends this intuition by revealing that this relationship is also affected by the correlation between the paid listings' relevance score and willingness to pay. Figure 4 demonstrates that placement revenues increase with the level of demand when WTP and relevance are positively correlated, but that an increase in demand does not yield greater revenues (except in the case of v α ranking, where revenue increases modestly with demand) when the correlation is negative.

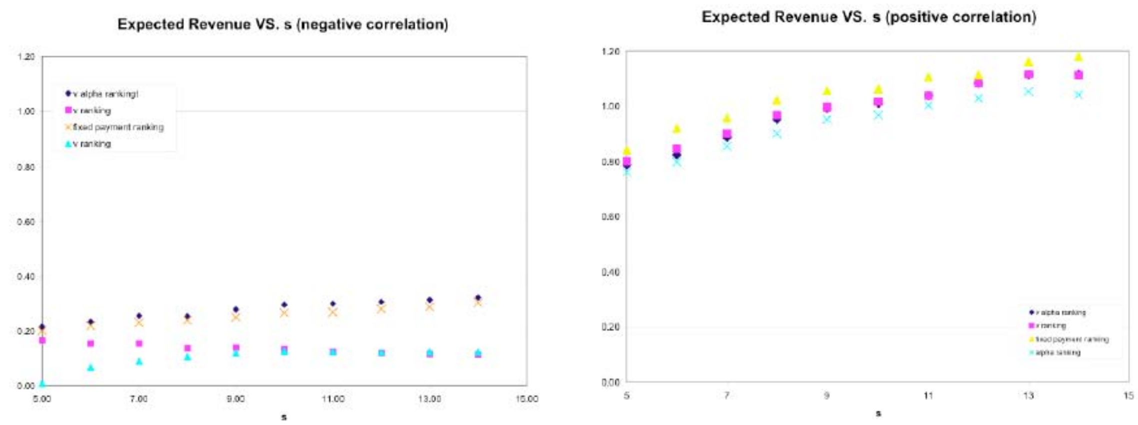

(a) negative correlation (b) positive correlation

Figure 4: Revenue versus the demand for paid listings s, under negative and positive correlation between relevance and willingness to pay.

Finding 4 When the correlation between the paid listings' relevance and WTP is highly positive (cov = 0.15, correlation=0.9 in this case), the increase in the demand of paid placement (number of potential paid links) increases the search engine's revenue. When the advertisers' relevance and WTP are very negatively correlated (correlation = 0.9 in this case), there is no obvious increasing trend when increasing the demand.

It might appear counterintuitive at first sight that an increase in demand for placement slots fails to increase placement revenues. However, when relevance and WTP are negatively correlated, the v ranking and posted price mechanisms systematically favor those advertisers with lower relevance since these are the advertisers with the k highest bids. An increase in the number of advertisers competing for the k slots amplifies this effect. The v αranking mechanism, on the other hand, cancels out the negative correlation to some extent since it incorporates relevance into the selection of advertisers.

4 Conclusions

This paper analyzes the implementation of paid placement strategies for Web search engines. Via computational simulations, we compare alternative mechanisms for allocating paid placement slots, including stylized versions of mechanisms employed by leading services (Overture's v ranking mechanism, and Google's vα ranking mechanism). We find that vα ranking performs best in almost all cases, while v ranking also works well in the expected case of positive correlation between v and α. Editorial filtering can improve the performance of v ranking significantly. Placement revenues decrease when users' attention is significantly lower for lowerranked listings, emphasizing the need to develop better user interfaces and control features.

Paid placement in search engines is a thriving and growing industry. The average price paid per click on Overture's network in 2003 was roughly US $0.40. Paid placement revenues of the two leading firms were between $500 million and $1 billion in 2003. Paid placement is widely credited for the revitalization of the search engine business.[5] This paper has taken early steps in studying the implementation of paid placement in search engines. Much more work remains to be done, with respect to understanding user attitudes towards different forms of paid placement, the impact of various user interfaces on users' willingness to accept and browse placement slots, and on the design of optimal mechanisms for allocating placement slots.

Acknowledgements
We thank Sandip Debnath, Dan Fain, Gary Flake, Lee Giles, and Steve Lawrence.

References

Asdemir, K., N. Kumar, and V. Jacob (2002): "INTERNET ADVERTISING PRICING MODELS," Discussion paper, School of Management, The University of Texas at Dallas.

Bhargava, H. K., and J. Feng (2001): "Paid Placement Strategies for Information Gatekeepers," in Proceedings of the Eleven'th Annual Workshop on Information Technology and Systems, New Orleans, LA. (2002): "Preferential Placement in Internet Search Engines," in Proceedings of the Eleventh World Wide Web Conference, Honolulu, HI.

Bhargava, H. K., and S. Sundaresan (2002): "Optimal Design of Contingency Pricing in IT-Intensive Commerce," in Proceedings of the TwentyThird International Conference on Information Systems, Barcelona, Spain.

Bollacker, K., S. Lawrence, and C. L. Giles (2000): "Discovering Relevant Scientific Literature on the Web," IEEE Intelligent Systems, 15(2), 42–47.

Breese, J. S., D. Heckerman, and C. Kadie (1998): "Empirical analysis of predictive algorithms for collaborative filtering," in Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, pp. 43–52.

Cotlier, M. (2001): "The payoff of paid search listings," Catalog Age, 18.

Cotriss, D. (2002): "Marketers report high ROI with paid listings," B to B, 87.

Dolbeck, A. (2003): "Search engine providers: Dawn of a new tech bubble?," Weekly Corporate Growth Report.

Hoffman, D., and T. Novak (2002): "How to Acquire Customers on the Web," Harvard Business Review, 106(31), 179–188.

Hu, J. Y. (2003): "Performancebased Pricing Models in Online Advertising," Discussion paper, Sloan School of Management, MIT.

Klemperer, P. (1999): "Auction Theory: A Guide to the Literature," Joünal of Economic Surveys, 13(3), 227–286.

Kumar, S., M. Dawande, and V. Mookerjee (2003): "Revenue Maximization in Web Ad-

vertisement," in Proceedings of The Thirteenth Annual Workshop On Information Technologies And Systems (WITS'03), Seattle, WA.
Lawrence, S., and C. L. Giles (1999): "Accessibility of Information on the Web," Nature, 400(6740), 107–109.
Reinhardt, A. (2003): "AND YOU THOUGHT THE WEB AD MARKET WAS DEAD," Business Week.
Rolland, E., and R. A. Patterson (2003): "Classification in Online PayForPerformance
Advertising," in Proceedings of The Thirteenth Annual Workshop On Information Technologies And Systems (WITS'03), Seattle, WA.
Tantono, W., S. Selby, J. Bagner, and C. Sonu (2002): "FTC advises search engines to disclose paid placements," Intellectual Property & Technology Law Journal, 14.
Weber, T. A., and E. Z. Zheng (2003): "A Model of Search Intermediaries and Paid Referrals," Discussion paper, 021201, The Wharton School, University of Pennsylvania.