

A Revenue Sharing Mechanism for Federated Search and Advertising

M. Brambilla, S. Ceppi, N. Gatti
Politecnico di Milano, Italy
{mbrambil, ceppi, ngatti}@elet.polimi.it

E. H. Gerding
University of Southampton, UK
eg@ecs.soton.ac.uk

ABSTRACT

Federated search engines combine search results from two or more (general-purpose or domain-specific) content providers. They enable complex searches (e.g., complete vacation planning) or more reliable results by allowing users to receive high quality results from a variety of sources. We propose a new revenue sharing mechanism for federated search engines, considering different actors involved in the search results generation (i.e., content providers, advertising providers, hybrid content+advertising providers, and content integrators). We extend the existing sponsored search auctions by supporting heterogeneous participants and redistribution of monetary values to the different actors, while maintaining flexibility in the payment scheme.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics; H.3.3 [Information Systems]: Information Search and Retrieval

General Terms

Design, Economics, Experimentation

Keywords

Computational Advertising, Federated Search Engine, Revenue Sharing, Auctions, Mechanism Design, Web Search

1. INTRODUCTION

General-purpose search engines are the main way for users to access information. However, there is an increasing demand for more sophisticated queries, higher quality of results, and richer media related to a query, such as images, videos, and pieces of news. *Federated Search Engines (FSEs)*, which integrate a broad set of data sources, are better able to address these user needs. This class of search engines supports the publishing and integration of high-quality data sources for vertical domains (extracted from the deep web or curated data repositories) or for general-purpose information, possibly allowing the user to select the sources and get improved results with respect to traditional search engines. FSEs must also integrate *advertising service providers*, which are crucial to ensure both the economic sustainability and quality of the results. Our focus is on the latter.

Our work is closely related to the auctions for pay-per-click sponsored search and display advertising. However, even though there are many works that study sponsored search auctions, none of these deal with the problem faced by an FSE which needs to aggregate information from different types of sources, and provide them with adequate compensation. Only a few works have addressed the problem of revenue sharing mechanisms in search applications [2, 3], but none of them have addressed the general scenario of federated search, where different classes of actors are involved in the search results generation.

To address the above problem, we extend the work in [2] and we propose a novel *ad-hoc revenue sharing mechanism for FSEs*. Specifically, we use *mechanism design* [4] to introduce a new payment mechanism that redistributes the yields from advertising to produce adequate monetary incentives to all the actors in the system.

2. FSE CHARACTERIZATION

An FSE is essentially a sophisticated content integrator that exploits and integrates existing information sources, including both content providers and advertisement providers. The FSE can implement two main business cases:

1) *multi-domain search approach*, where users can submit complex queries about different topics or items at a time. The multi-domain query is split into a number of single-domain queries, each one addressed by a domain-specific content service provider. Results are then produced by combining the various domain-specific items;

2) *meta-search approach*, where users submit traditional, simple queries but expect results to be more accurate than in a general-purpose search engine. The same query is forwarded to several engines and the results consist of the union and re-ranking of the different result sets.

In both cases, the FSE collects contents and advertisements, and publishes them together in every result page generated as a response to a user query. The choice of how many ads and content items to show on the Web page is determined by the FSE, which also selects the information sources used to generate the results shown.

Each of these sources belongs to one of the following classes:

Content-Only Providers (CPs) (e.g., zillow.com) provide content for the FSE. The content consists of datasets with a high information value for the users because of their precision, completeness, detail and/or coverage. It is the core intellectual property for the CP and is costly to produce.

Integrated Content and Advertising Providers (CAPs) (e.g., Google) provide contents while at the same time serving their own ads (as opposed to using third-party services).

Typically, these websites would allow third parties, such as the FSE, to publish their search results by requiring the FSE to show (some of) their ads too.

Advertising Providers (APs) (e.g., Yahoo!’s *RightMedia*, Google’s *AdSense*) only offer advertisement services. They publish ads from a set of advertisers and their goal is to provide targeted ads to users who visit partner web sites.

In addition to the three classes of providers, the FSE involves two classes of stakeholders interested in using the system: *end users*, who visit the FSE web site, submit their queries, read the results and possibly click on the shown advertisements; and *advertisers*, who provide (and pay) the ads to the APs and CAPs.

3. REVENUE SHARING MECHANISM

The FSE scenario requires an intricate revenue sharing mechanism that provides the appropriate incentives to all the providers. To this end, our approach is based on the following assumptions/requirements: (i) the only *revenue sources* of the system are the advertisers, who pay the APs and CAPs for getting their ads published, with a pay-per-click model; (ii) the FSE can integrate different *heterogeneous providers*, adopting, for each of them, a possibly different payment or compensation mechanism; (iii) to grant *adequate compensation* to every provider that takes part to the FSE’s activity, an appropriate redistribution of the revenue may be necessary; (iv) the approach must grant *flexibility* to the needs of both the FSE and the providers in a variety of business scenarios, in particular including: profit-seeking and non-profit FSEs; different sizes of the eco-system, in terms both of classes and numbers of providers.

The revenue-sharing process starts with the advertisements, which are communicated by the APs and CAPs to the FSE, which in turn allocates (a subset of) them to the available slots. Once an ad has been clicked, the corresponding advertiser pays the provider in which it is registered. For simplicity, we assume that payments are per click and that they are defined according to some auction model (e.g. using GSP, although here we do not consider the specifics of the mechanism used by the APs and CAPs, and our revenue sharing mechanism is independent of this choice). The mechanism’s aim is to produce an allocation that maximizes the total revenue of the system. In order to achieve that, the mechanism needs to elicit private information from the providers about the individual ads (e.g., values). However, providers could misreport their true information, if they gain more by doing so, thereby hindering the mechanism to find the allocation maximizing the total revenue.

Once the total revenue has been established, in the next step, APs and CAPs give a portion of the revenue received by the advertisers, as defined by the *payment function* of the mechanism (as in [2]), to the FSE. The FSE subsequently redistributes a portion of the received revenue, as per the *redistribution function* of the mechanism, to the actors. The redistribution function we propose is an extension of the work presented in [1]. In particular, the novelty of our redistribution function is that it is parametric, a weight being assigned to each providers’ class. Our function is asymptotically optimal, i.e., it is not possible to redistribute more without creating incentives for some providers to misreport. Note that this redistribution includes also the APs and CAPs even if they already possess part of the revenue coming from the advertisers. This redistribution is important to satisfy the *adequate compensation* requirement, which provides the proper




Figure 1: Average expected revenues of a single CP, AP, and CAP as ratio of the total revenue.

incentives to all the actors to take part in the mechanism with all their assets. In particular, the mechanism must ensure that actors that can play a certain role (in particular, CAPs) are not tempted to change it, by dropping either the content or advertising service provisioning aspects.

4. EXPERIMENTAL EVALUATION

We empirically evaluate our revenue sharing mechanism with the aim of comparing the revenues of the actors using different configurations. In particular, we consider the configuration in which the weight assigned to CPs or APs is zero (R1 and R2, respectively), the weights are equal for all the providers’ classes (R3.1), the weight of CAPs is higher than the one of the other classes (R3.2, R3.3). Our empirical evaluation is based on the *Yahoo! Webscope A3* dataset.

We start by considering the FSE’s expected revenue. With no redistribution, the FSE’s expected revenue is about 49.3% of the total revenue, while with redistribution the revenue is always below 10% of the total revenue. This implies lower incomes for the FSE, but in general a higher motivation for the providers, which will be pushed towards producing better results and advertising, so as to be kept in the loop.

We now focus on the average redistribution to each single CP, AP, and CAP, as a ratio of the total revenue. Notice that, while the revenue of CPs is exclusively due to the redistributions, the revenue of APs and CAPs is composed of two parts: what they receive without redistribution and the redistribution itself. In Fig. 1 we visualize this difference by dividing the bars corresponding to APs and CAPs into two parts, where the upper part is due to the redistributions. Since CAPs receive revenue even without redistribution, the overall expected revenue of this class of providers is always larger than that of CPs. Moreover, it is also never smaller than the revenue of an AP. Thus, the risk that CAPs misreport and behave like CPs or APs is avoided.

We also observe that the overall redistribution monotonically increase as the number of providers increases.

Acknowledgments. This research is partly supported by the “Search Computing” (SeCo) project, funded by ERC, under the 2008 Call for “IDEAS Advanced Grants”.

5. REFERENCES

- [1] R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *AAMAS*, pages 882–889. ACM, 2001.
- [2] S. Ceppi, N. Gatti, and E. Gerding. Mechanism design for federated sponsored search auctions. In *AAAI*, 2011.
- [3] Y. Chen, A. Ghosh, R. McAfee, and D. Pennock. Sharing online advertising revenue with consumers. *WINE*, pages 556–565, 2008.
- [4] Y. Narahari, D. Garg, R. Narayanan, and H. Prakash. *Game Theoretic Problems in Network Economics and Mechanism Design Solutions*. Springer, February 2009.