Decentralizing Sponsored Web Advertising

Rohit Khare, Ph.D.
Benjamin C. W. Sittler

CommerceNet Labs Technical Report 05-04
April 2005

Decentralizing Sponsored Web Advertising

Rohit Khare
CommerceNet Labs
490 South California Avenue
Palo Alto, CA
+1 (650) 289-4040
rkhare@commerce.net

Benjamin C.W. Sittler
CommerceNet Labs
490 South California Avenue
Palo Alto, CA
+1 (650) 289-4040
bsittler@commerce.net

ABSTRACT

CommerceNet has recently established a research lab focused on the impact of decentralization on electronic commerce. Specifically, we are interested in developing new software architectural styles that go beyond the limitations of traditional client/server systems.

Today, sponsored search auctions are purely centralized. A single, trusted agency controls who can advertise; acceptable advertising messages; the mapping of ads to pages; placement within pages; and, of course, the entire bidding process.

We are interested in learning more about the challenges of building a peer-to-peer alternative; and to discuss our experiments with an alternative mechanism, Ross Mayfield's "cost-per-influence."

Categories and Subject Descriptors

K.4.4 [**Computers and Society**]: Electronic Commerce – *Distributed Commercial Transactions*. D.2.11 [**Software Engineering**]: Software Architecture – *Domain-specific Architectures*.

General Terms

Design, Economics, Experimentation

Keywords

Decentralization, Advertising, Auctions

1. INTRODUCTION

CommerceNet is interested in decentralizing electronic commerce. One of the hottest e-commerce applications today is the renaissance of Web advertising with real-time, pay-per-click keyword matching [5].

Today, centralized 'ad networks' such as Google AdSense and Yahoo! (neé Overture) control:

- *Who* can advertise;
- *What* can be advertised (and *how*);
- *Who* can publish ads;
- *Which* ads run on each page; and
- *How* prices are set and revenue is shared.

Some other ad networks compete by increasing freedom along one or more of these axes. So-called Publisher-Driven Advertising (PDA, [2]) services such as BlogAds, AdBrite, and FederatedMedia assist publishers in launching their own ad networks by positing kinder, gentler hubs to match buyers and sellers.

However, a genuinely peer-to-peer solution that offers freedom along all of these axes remains an open research question in electronic commerce [34]. There are strong reasons to believe that decentralized mechanisms can approach the efficiency of centralized ones with global knowledge [70], but such an ad market has not been demonstrated yet, nor shown to be economically efficient enough.

Furthermore, keyword auctions are by no means the last word in advertising business models. There is a backlash to the democratizing impact of automation, because blind placement of ads according to keywords and click-through rates rewards publishers roughly in proportion to their traffic, not their social influence.

2. OUR BACKGROUND

CommerceNet is a nonprofit organization that combines research and entrepreneurship to accelerate Internet innovation in Silicon Valley. Over the past decade, some of our pioneering accomplishments include: the first secured Web credit-card transaction; the first consumer Internet adoption survey; and the first XML business vocabularies.

Rohit Khare joined CommerceNet in 2004 to establish a new research effort, the zLab Center for Decentralization. The initial research theme was based on his doctoral dissertation work in the field of software architecture [6]; earlier, he spun out a successful venture-backed startup from the same effort [9].

Ben Sittler is an associate member of technical staff who is implementing our prototype event-notification network and our experimental advertising server.

3. OUR APPROACH

Financial markets inspire our emerging theory for the software architecture of such decentralized systems. While there are many examples of centralized or distributed markets in the real world (such as the NYSE and NASDAQ stock markets, respectively), the most robust ones (such as the foreign currency trading markets) must tolerate disagreement rather than relying upon a single, globally-correct value. In those markets, the power to establish an equilibrium clearing price becomes decentralized over all the traders.

In the context of sponsored search auctions, we are proposing that there may not be a single cost-per-click in the future. If subsets of total market supply and demand meet at multiple points, and within hard real-time deadlines prevent 'backhauling' all of the data to a single point, local price equilibria may diverge significantly from agency to agency.

In addition, decentralized markets are regulated differently by society. Ideally, they are self-policing; in this case, that suggests that at a minimum, the trusted-third-party roles of 1) auctioneer; 2) ad server; and 3) accountant must be separated apart and partitioned to limit exposure to rogue agencies.

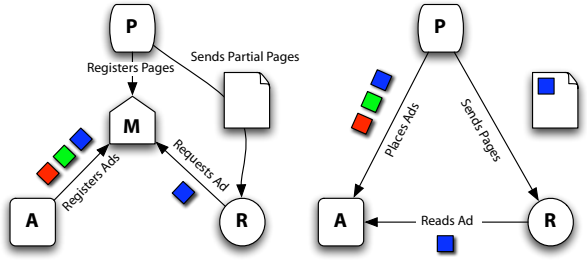

Figure 1: Contrasting intermediated and peer-to-peer SSAs

3.1 The ARRESTED Architectural Style

Whereas traditional client/server software architectures depend on consensus — as emphasized by the so-called ACID properties of transactions (*Atomicity, Consistency, Isolation, and Durability*) — eliminating critical dependency on an Earth-bound mission control center requires new architectural styles that ensure what we term the 'BASE' properties: using only *Best-effort* networking, substituting *Approximate* estimates for exact values, managing *Self-centered* trust relationships, and ensuring *Efficient* use of the network by discarding outdated information.

A first step towards such an architectural style has an extension to the Representational State Transfer (REST) style that describes how the World Wide Web works. We have incrementally modified REST to support *Asynchronous* notification of dynamically changing data; Routing event notifications to interested & trusted parties; Estimating current values from cached data; and assessing multiple agencies' opinions when making *Decisions* (ARRESTED). [6]

3.2 An Idealized Model

Figure 1 illustrated the idealized relationships in an advertising relationship: advertisers A offer coupons to publishers P to present ads to readers R; the coupons pay off if R "redeems" the ad to A's satisfaction. In practice, a mediator M is required to operate the market today, for legal and technical reasons.

These two factors correspond to the two basic forces of decentralization: latency and agency. Because it takes time and energy to transmit information, it is impossible for all As to reach all Ps in time; M instead substitutes a low-latency interconnect (their data center) and precomputes possible matches using keyword-indexes of pages and ads.

The second factor refers to political independences of sovereign agencies: in the ideal case, P has the right to accept or reject each advertiser and advertisement, but in practice must cede judgment to M in order to rapidly compute the utility of competing offers. Even more troubling, all As and Ps must trust M to operate a sealed-bid second-price auction fairly.

We admit that we do not yet have a rich enough model to capture the nuances of current practice, much less neatly delineate the transformations required to decentralize it.

Nonetheless, we speculate about techniques that may be applicable. For example, to reduce ad-serving fraud, all parties could participate in a content distribution network, with randomized responsibilities. Similarly, judging a sealed-bid auction might be more tractable if randomly-chosen observers could only compare pairs of bids. Auditing could rely on notarizing logfiles against a Merkle hash tree [4]. Payment could be handled by Chaum's blinded e-cash [3]. And the arduous task of classifying ads as spam, adult, sports-related and so on might be tackled by tagging and folksonomies [7].

4. EXPERIMENTS

Our view of decentralization encompasses the freedom to adopt new business models, not just adding freedom-of-association and freedom-of-speech to current sponsored search auctions. One such proposal is Ross Mayfield's "Sell-Side Advertising" or "Cost-per-Influence" model: syndicated ads selected by bloggers that pay off for prior sources [1, 8]. Its intent is to reward authors proportional to their "influ-

ence" over other authors (as indicated by choice to run the ad on their own sites), not the number of readers.

We have developed a prototype server in Python using the Twisted framework. Users can create ads, browse ads, or choose to syndicate a copy of an existing ad to insert on their own website. Ad copy is arbitrary XHTML that will be loaded into an IFRAME on the publisher's pages. An ad also specifies a total budget and deadline. The current "settlement" algorithm divides the pot equally for each ad served, then collects a 'tax' back up the syndication tree using the HTTP Referer header.

5. FUTURE WORK

Our informal motto is to develop "software that works the way society works" — so while our experimental work is not, strictly speaking, an auction (nor does it use keywords), we believe it would be valuable to consider it as part of the space of future advertising models a toolkit for decentralized web advertising ought to support.

6. ACKNOWLEDGMENTS

We would like to thank our colleagues for many productive discussions to help clarify our thinking so far. Not least is Ross Mayfield, for presenting his model in detail and suggesting experiments to pursue.

7. REFERENCES

More at
http://labs.commerce.net/wiki/index.php/SellSideAds

[1] Battelle, J. *Sell Side Advertising: A New Model?* in *Searchblog*, 26 August, 2004.

[2] Battelle, J. *A New Idea for Publishing in Technology Review*, January, 2005.

[3] Chaum, D., Fiat, A. and Naor, M. *Untraceable Electronic Cash, in Advances in Cryptology*, (1988), pp. 319-327.

[4] Haber, S. and Stornetta, W. S. *How to Time-Stamp a Digital Document in LNCS #537*, 1991. p. 437.

[5] Khare, R. *An Introduction to zLab*. CommerceNet, 2004.

[6] Khare, R. and Taylor, R. N. *Extending the Representational State Transfer Architectural Style for Decentralized Systems (in preparation) in TOSEM*, 2005.

[7] Mathes, A. *Folksonomies - Cooperative Classification and Communication Through Shared Metadata*. University of Illinois Urbana-Champaign, 2004.

[8] Mayfield, R. *Cost-Per-Influence* in *Many2Many, a group weblog on social software*, 12 July, 2004.

[9] Udell, J., Gillmor, S. and Eds. 2002 *Technology of the Year Award: Publish/Subscribe Technology: KnowNow 1.5* in *Infoworld*, January 24, 2003.