Negotiating Multimedia Advertising with Attention Owners

Yue Zhang*
cfactor Works
250 2nd Avenue S.
Saskatoon, SK, S7K 1K9, Canada
yue.zhang@agents.usask.ca

Nadeem Jamali
Agents Lab, University of Saskatchewan
176 Thorvaldson Bldg., 110 Science Place
Saskatoon, SK, S7N 5C9, Canada
jamali@cs.usask.ca

ABSTRACT

Advertising is increasingly an integral part of multimedia delivery over the Internet. Traditionally, brokers – intermediaries between content providers, advertisers, and viewers – have determined the fine balance between the content desired by viewers and the advertising embedded in the content. Parameters of this balance are informed by fields of psychology and marketing, which help target viewer segments identified by their viewing habits. Oddly, mechanisms available to individual viewers to inform this balance are coarse grained: one can change the channel!

We take an *owned attention* view of the problem with an explicit treatment of both attention and its ownership. This approach specializes the CyberOrgs model for encapsulating computations with owned resources available for their execution. Particularly, we treat a multimedia consumer's attention space as a precious resource owned by the viewer. Viewers pay for the content they wish to view in dollars, as well as in terms of their attention. Advertisers pay for viewers' attention by subsidizing the cost of their content.

This paper presents the rationale, design, implementation, and evaluation of our solution, FlexAdSense. Our approach affords finer grained control capability to viewers than what is offered by existing approaches. Pluggable customizable policies specify negotiation positions of different parties, scalably automating typical negotiations. Experimental work demonstrates that the approach scales well, and informs decisions about allocating resources to servers.

Categories and Subject Descriptors

D.2.11 [**Software Engineering**]: Software Architectures—*Domain-specific architectures*; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Performance, Design, Economics, Human Factors

Keywords

Attention, Online video, CyberOrgs, Interactive control

1. INTRODUCTION

Advertising has historically played a key enabling role in the delivery of valuable multimedia content to consumers by subsidizing consumers' access to it. And now, increasingly, advertising revenue is the essential part of the business model for companies providing a wide range of products over the Internet. Online multimedia content lies at the intersection of the former legacy and the latter trend.

Online video advertising spending rose 40.19% from $734 million to $1.029 billion from 2008 to 2009 [3] – and is expected to reach $1.44 billion in 2010, and $5.2 billion by 2014 – just as total Internet ad spending fell 4.6% from $23.4 billion to $22.4 billion because of the economic downturn.

There are fundamentally two models used by multimedia brokers, such as cable operators and multimedia web sites, to pay for content viewed by consumers: one delivers paid high-quality content to viewers ("pay-per-view") without any advertising, and the other provides free content viewing, but at the cost of embedded advertising. Many approaches, including the popular subscription fees model, essentially combine these two. In all cases except strict pay-per-view, the broker makes key decisions about whether or not to display advertisements, how many and which types of ads to display and when to display them. Viewers typically have only limited coarse-grained control. Figure 1 compares these various approaches according to the granularity of control over both multimedia content and advertising that they offer viewers.

A number of advertising models have emerged for supporting delivery of digital content over the Internet; however, their primary focus is on personalizing advertisements for viewer segments, with minimal decision making capacity for individual viewers. The iMEDIA [9] business model applies technologies exploring viewers' interactive data and empowers viewers with control on their personal information. Yoon et al. [11] have proposed a system based on the TV-Anytime standard, which provides both media library services and targeted advertisement services. In the web domain, personalization techniques have been proposed for effective advertising. These can be roughly classified into four categories [8] according to the mathematical techniques adopted: *data mining*, *decision trees*, *linear programming*, and *nearest-*

*This work was carried out as part of Ms. Zhang's M.Sc. thesis at the Agents Lab, University of Saskatchewan.

MM'10, October 25–29, 2010, Firenze, Italy.

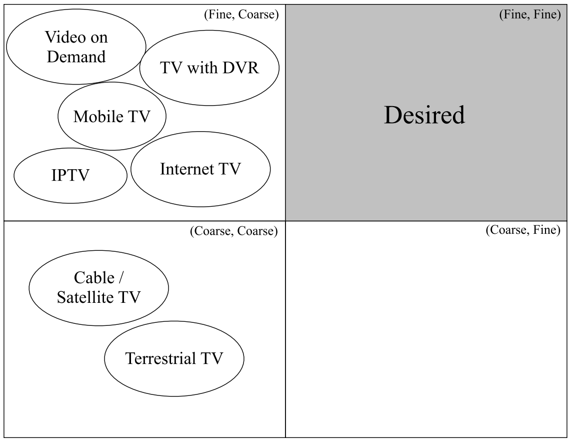

Figure 1: Granularity of control in existing platforms (Content, Advertising).

neighbour collaborative filtering algorithms. Viewers typically have limited control over the advertisements displayed.

2. TRADE IN ATTENTION RESOURCE

The study of *Attention* was pioneered by psychologist William James in 1890 [5]. He identified two characteristics of attention: focalization and concentration. The topic was first introduced to Computer Science by Herbert Simon in 1971 [10], when he highlighted the imbalance between the "wealth of information" and the "poverty of attention." This recognition continues to inspire a variety of personalization techniques and recommender systems [2].

We take a resource ownership view on advertising problem. We view a multimedia consumer's attention space – abstracted as a display screen for an engaged viewer – as a precious resource owned by the viewer. Consumers pay in cash (as well as in terms of their attention) for the content they wish to view. Advertisers may partially pay for a viewer's content, in return for receiving the viewer's attention to their advertising. We build a market of viewers' attention spaces in which advertisers can trade, just as viewers can trade in a content market. Consider these scenarios involving a fictional viewer Jack:

Scenario 1: Jack browses through the 100 or so channels on his TV, and finally settles on his favorite detective show. His viewing is periodically interrupted by advertisements. He wishes he could view his favorite content on his own schedule without the ads even if it required paying extra. Additionally, he wishes there were better mechanisms to filter available choices to his tastes.

Scenario 2: Jack wishes he could set a monthly budget that the system could use to automatically scheduled content and advertisements for him, while accommodating explicit overrides when he chose: relatively more ads when the budget is tight and fewer ads when the budget is freer.

Scenario 3: Jack dislikes pop up ads and wishes that he could select the types of ads displayed on his screen.

Scenario 4: Jack wants to protect his 6-year-old daughter from depiction of aggressive behavior, smoking, etc.

None of the existing mechanisms satisfy the requirements identified in these scenarios. Our approach of enabling fine-grained resource trade – in real time – between owners of the attention resource and the parties interested in acquiring them, seeks to address this need.

2.1 FlexAdSense

Consider the four parties involved in multimedia delivery: content publisher, content viewer (consumer), broker and advertiser. Each party owns certain resources and seeks to obtain certain resources from others. Multimedia publishers own multimedia content, which is a type of information resource. They wish to make profit by selling content resource using brokers' intermediary service. Viewers own cash and attention resource, and want to be entertained by viewing multimedia content. Brokers own delivery resources such as cables, network bandwidth and storage servers, and provide services such as payment mechanisms and customer support modules. Advertisers own cash and are interested in viewers' attention. The resources of the greatest interest to us are *viewer attention* and *multimedia content*.

We treat a display – an abstraction of a viewer's attention space – as a type of resource owned by its viewer, with both a spatial extension and a temporal duration. We specialize the CyberOrgs model [4] for enabling trade in this type of resource. Cyberorgs are distributed resource encapsulations which use *eCash* to buy and sell resources from/to each other. Contracts are negotiated between cyberorgs; these *contracts* stipulate the types, quantities and costs at which resources would be made available to a buyer by a seller. To specialize the model for our purpose, we replace *eCash* with *real cash*, and viewer attention becomes an owned resource. Screen real estate is owned by viewers, and advertisers have to purchase viewers' attention resources by proposing to partially pay for the viewers' content. Consequently, viewers are empowered to decide what is displayed in their attention space by managing access admission.

FlexAdSense allows the different parties to use a number of parameters to specify their objectives in the exchange. For example, viewers can specify desired quality, playback capabilities, content category, etc. for content, or the type (banner, video, etc.), category (apparel, travel, etc.), and the specific format (position on screen, duration, etc.) for advertising. Viewers can further specify their preferred payment mechanism. Similarly, advertisers and contend publishers specify their objectives. Specifications of these objectives from different parties convert an otherwise more explicit negotiation into a matchmaking exercise leading to brokered agreements and *contracts* between the parties.

2.2 Policies

Because fine-grained negotiation can require significant user interactions, leading to unwanted additional demands on viewer attention, we allow parties to specify policies for positions to be taken in automated negotiations in predictable situations. Users may create their own policies, or adopt or customize policies available in a repository. There are default policies in place for users who have not created customized ones. Policies are also composable. We have implemented three types of policies: preference policies, payment policies and privacy policies, which respectively specify policies for content/ad selection, payment and privacy.

Preference Policies. Viewers, advertisers and publishers can all use preference policies. Preference policies reduce explicit user interaction by specifying rules for filtering out unwanted or unrelated ad/content. Viewers' preference policies provide constraints which are used to create choices of display sequences with advertisements embedded in multimedia content streams. Recall scenarios 1 and 3 involving our viewer Jack: he can now specify the multimedia content category of his interests, as well as the ad types he dislikes. Jack can specify constraints about price, category, date, type, language, etc., which are used for selecting ads. Similarly, for his content preference policy, Jack can specify the category, language, video quality, and price constraints, in much the same way as he did for ad preference policies.

Advertisers and publishers can specify ad/content attributes, by which they can target their audience and ease the ad/content sequence creation process. For example, a lingerie advertiser not interested in wasting advertising money on men, can set the preferred viewer gender to female (subject to privacy settings allowing access to such information). At the same time, these policies can be used to aid viewers' preferences, such as to support Jack's Scenario 4 preference to exclude adult–only content/ads. Advertisers (and publishers) can also choose to target viewers in specific geographic locations.

Payment Policies. An advertiser can specify payment policies by specifying *pricing models* or setting monthly or daily spending *budgets*. Pricing models can be static or dynamic: static pricing is independent of the context in which ads appear; dynamic pricing depend on attributes defining the context. For example, an advertiser may want to vary the price to be paid to a viewer depending on the points of insertion of ads in the content. Consider a model which gives preference to ads shown close to the middle of the content being viewed. Price paid by the advertiser would be $P_{max} - \left|\frac{(T_{insert} - D_{con}/2)*(P_{max} - P_{min})}{D_{con}}\right|$, where P_{max} and P_{min} are the maximum and minimum prices for the ad, T_{insert} is the point of insertion, and D_{con} is the duration of the content.

Viewers can choose the balance between viewing preferences and the price paid, by specifying payment policies. Recall Scenario 2 in which Jack wishes to view uninterrupted content within constraints of a monthly budget. This could be specified using the policy $P_i = \frac{B_i}{f*(30-i)}$, where P_i is the budget for the i^{th} day of the month, B_i is the remaining balance in the monthly budget as of the i^{th} day, and f is the frequency (between 0 and 1, representing percentage of days) that the viewer watches content on television. This policy tries to distribute the budget evenly over the month.

Privacy Policies. Privacy issues have been thoroughly addressed in the literature [7]. In our approach, viewers specify privacy policies determining how their data can be used. Particularly, viewers' personal information is treated as a type of resource with commercial value, and is traded in a market at a fine grain.

Policies are adopted in stages before a contract is finalized. Preference policies are adopted first, leading to generation of content and ad lists. Next, depending on whether or not the viewer wants to make the final selection of ads and content, either the viewer makes these selections, or they are made for the viewer. Finally, payment policies are adopted, following which a contract is deemed to be agreed.

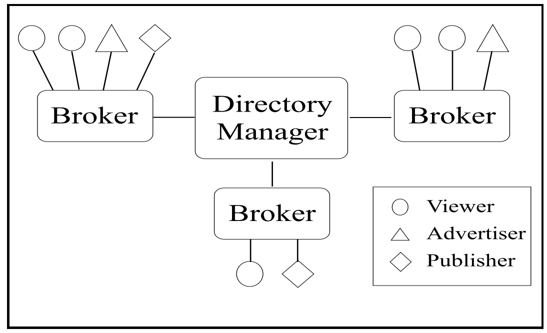

Figure 2: Distributed Structure of FlexAdSense

3. IMPLEMENTATION

Our prototype is built using a distributed multi-agent system called Actor Architecture (AA) [6] based on Actors [1]. The structure of our system is shown in Figure 2. Each party is implemented as an agent; agents are designed according to their roles. Advertisers, publishers and viewers are represented by client agents responsible for interacting with and acting on behalf of the users they represent. Broker agents serve the clients using copies of a shared database of users' profiles, policies and information about multimedia content and ads. A special agent called *Directory Manager* (DM) offers a *Yellow Pages* service. Several example policies are implemented and user interfaces are designed implementing key mechanisms required for a functioning system.

Advertiser agents represent advertisers in managing their ads, policies and accounts, as well as investigate viewers' information when required. Publisher agents represent publishers in managing their multimedia content resources and accounts. Viewer agents represent viewers in managing their attention space by searching for and selecting content resources, selling attention resources at a fine grain, as well as managing their profiles, policies and accounts. Brokers maintain a shared database which stores information about all users, including users' account information, viewers' profiles, various customized policies, as well as advertisement and content information. The Directory Manager (DM) maintains a directory of all brokers. New brokers register with the DM and replicate the database before going online. When a new client agent is created, the DM arranges a rendezvous with a broker, which subsequently takes care of requests from that client. The DM attempts to balance brokers' loads: the broker with the lightest load is selected to serve a new client agent. Additionally, the DM coordinates updates to brokers' copies of the database to keep all copies consistent. Note that although we chose to implement a shared database because it simplified our design, it would not be necessary in a deployed system.

There are two types of communication between client agents and their brokers: query requests and modification requests. Query requests are requests between users and their designated brokers for searching for or subscribing to content/ad, or retrieving users' personal information. These requests are made by sending messages to the relevant brokers. Modification requests are requests which change the database. These less frequent communications request registration of a new ad/content, registration of a new user or an update to a user's preferences. These modification requests trigger a global update to the database.

4. EVALUATION

This work can be evaluated along multiple dimensions. Key among these is the level of flexible fine-grained control afforded all involved parties. A simple comparison with existing systems' control granularity shows that FlexAdSense does indeed offer greater flexibility and a finer grain of control. The question which then arises is: at what cost? Two interesting metrics of evaluation of this cost are server scalability and the added demand on viewer attention resulting from interactions with the system. We examined the latter in a preliminary study counting the number of mouse clicks required from a viewer to carry out various tasks: tasks identified in Jack's scenarios take between 2 and 6 clicks. The remainder of this section addresses scalability.

Experiments were carried out using six Mac OS X Servers each running an actor platform of AA. Servers had 2×2.8 GHz Quad-Core Intel Xeon CPUs with 8 GB memory each; they were connected using a Gigabit network switch.

Our experiments used simulated load, with a number of viewers concurrently sending information requests at preset rates. These are the more frequent types of requests in the system; requests to modify schedules would be orders of magnitude less frequent, and thus have relatively insignificant impact on scalability. The frequency of view requests (10 per sec) is also orders of magnitude higher than what would happen in practice: not all viewers would be actively searching at the same time, nor would they search at as high a rate. In fact, the cumulative rate of generation of requests, not the number of viewers, turn out to be the significant determiner of performance. Consequently, our results about highly demanding viewers are equally applicable to orders of magnitude larger numbers (we expect 10^6) of actual viewers with typical demands.

A set of experiments measured the total amount of wall-clock time required to complete serving requests as the number of viewers (distributed over three machines) grows from 1 to 3000. The requests were generated by viewers at the rate of 10 per second; the time required to process a request was set to 10 ms. All brokers were located on the same (multi-core) server. As Figure 3 shows, for the one broker case, as the number of viewers grows, the execution time increases significantly before becoming linear. However, there are orders of magnitude improvements when the load is divided between 2 or 3 brokers (note the logarithmic scale on y-axis). This shows that relatively few brokers executing in parallel can sufficiently improve performance.

Additional experimental work studied broker capacity, and impact of the number of servers and brokers per server on execution time. All of these showed high scalability; interestingly, the experiments on number of servers showed that the only a small number of distributed servers are needed for maximum performance; alternatively, simply adding additional servers does not significantly improve performance. Details are excluded here because of space limitations.

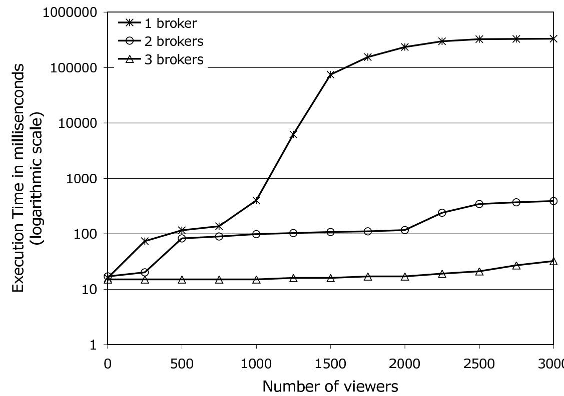

Figure 3: Execution time vs. numbers of viewers.

5. CONCLUSION AND FUTURE WORK

We presented an owned attention resource based approach to making decisions about advertising embedded in multimedia content, where viewers can sell their attention resource to advertisers. Different parties can install pluggable policies which enable their participation in negotiations without explicit interaction requiring attention resource. A distributed prototype has been implemented and evaluated, which shows that the finer granularity of control can be enabled scalably. An interesting direction for future exploration is to extend the idea to dynamically evolving groups of viewers, where viewers' agents would negotiate with each other before the group's preferences are negotiated with advertisers. Finally, some broader questions require further thought. For instance, advertisers may be interested in targeting new market segments; matchmaking based on stated preferences alone may preclude such matches.

Acknowledgement: Support from NSERC, CFI, and the Province of Saskatchewan is gratefully acknowledged.

6. REFERENCES

[1] G. Agha. *Actors: A Model of Concurrent Computation in Distributed Systems.* MIT Press, Cambridge, MA, USA, 1986.

[2] R. P. Bagozzi. Marketing as Exchange. *Journal of Marketing*, 39(4):32–39, 1975.

[3] eMarketer Newsletter. US Online Ad Spend Turns the Corner. Article on findings in "US Ad Spending" by eMarketer's D. Hallerman, December 2009.

[4] N. Jamali. *CyberOrgs: A Model for Resource Bounded Complex Agents.* PhD thesis, UIUC, 2004.

[5] W. James. *The Principles of Psychology.* Dover, 1950.

[6] M.-W. Jang. The Actor Architecture Manual, 2004.

[7] A. Juels. Targeted Advertising ... And Privacy Too. In *Proc. of Conference on Topics in Cryptology*, pages 408–424, London, UK, 2001. Springer-Verlag.

[8] T. Oshiba, Y. Koike, M. Tabuchi, and T. Kamba. Personalized Advertisement-Duration Control for Streaming Delivery. In *Proc. of ACM Multimedia*, pages 21–28, 2002.

[9] K. Pramataris, D. A. Papakyriakopoulos, G. Lekakos, and N. A. Mylonopoulos. Personalized Interactive TV Advertising: The iMEDIA Business Model. *J. of Electronic Markets*, 11(1):17–25, 2001.

[10] H. A. Simon. *Designing Organizations for an Information-Rich World*. Johns Hopkins Press, 1971.

[11] J. H. Yoon, H. K. Lee, J. M. Seok, H. Y. Kim, and H.-K. Lee. TV-Anytime Based Personalized Bi-Directional Metadata Service System. In *Symposium on Consumer Electronics*, pages 1–6, 2006.