

# COMPETITIVE MARKET-BASED ALLOCATION OF CONSUMER ATTENTION SPACE

---

This re-print corresponds to the paper “*Competitive Market-based Allocation of Consumer Attention Space*”, by Sander M. Bohte, Enrico H. Gerding, and Han La Poutré, and appeared in the Proceedings of the 3rd ACM Conference on Electronic Commerce (EC-01), editor M. Wellman, Tampa, USA, The ACM Press, pp 202-206, 2001.

# Competitive Market-based Allocation of Consumer Attention Space \*

[Extended Abstract]

Sander M. Bohte<sup>†</sup>

Enrico Gerding<sup>†</sup>

Han La Poutré<sup>†‡</sup>

<sup>†</sup>CWI, Centre for Mathematics and Computer Science  
Amsterdam, The Netherlands

<sup>‡</sup>TU Eindhoven  
Eindhoven, The Netherlands

sbohte,egerding,hlp@cw.nl

## ABSTRACT

The amount of attention space available for recommending suppliers to consumers on e-commerce sites is typically limited. We present a competitive distributed recommendation mechanism based on adaptive software agents for efficiently allocating the “consumer attention space”, or banners. In our approach, each agent bids in an auction for the momentary attention of each consumer. Successive auctions allow agents to rapidly adapt their bidding strategy to focus on consumers interested in their offerings. We demonstrate the feasibility of our system by an evolutionary simulation, and reflect on the advantages of this distributed market-based approach.

## 1. INTRODUCTION

With the advent of electronic marketplaces, scale limitations as encountered in the brick-and-mortar world no longer apply. At the same time, novel problems are encountered, like how consumers can find their way in large marketplaces.

A preferred solution is to have a central party propose relevant suppliers and products to a consumer. This central filtering mechanism uses knowledge of both users and shops and the product domain to determine recommendations; such recommender systems work well within limited domains, e.g. a book or music store. However, to maintain accuracy in a large marketplace with many consumers and suppliers, the latter will need to reveal detailed and perhaps sensitive business information to the central party. Central filtering mechanisms may thus suffer from objections from the suppliers. Furthermore, the amount of information to be processed and maintained by the central party can become unmanageable. Other approaches are needed to complement central or personal filtering systems.

\*Extended Abstract, research carried out within the framework “Autonomous Systems of Trade Agents in E-Commerce”, funded by the Telematics Institute.

In this paper, we present a framework for a distributed Competitive Attention-space System, *CASy*, to allocate the scarce resource that is consumer attention, via the techniques of dynamic market-based control [2] and adaptive software agents [6]. In the example of an electronic shopping mall, the task of matching a specific online consumer to a set of suitable shops is delegated to the individual shops, each of which evaluates the information that is available about the current consumer: the consumer’s current interests and other information which the consumer is or has been willing to provide; e.g. keywords, product queries, and available parts of a personal profile. Based on this information and on their own domain knowledge, shops can make a monetary bid in an auction where a limited amount of consumer attention space, or banners, for the particular momentary consumer is sold. To facilitate *CASy*, the system is designed as a multi-agent system where each shop is represented by a software agent that executes the task of bidding for the attention of each individual consumer. The use of learning software agents allows shops to rapidly adapt their bidding strategy such that they only bid for consumers that are likely to be interested in their offerings. Furthermore, the shops can incorporate up-to-date and detailed domain knowledge into their agents, which allows for accurate targeting of a shop’s audience, taking into account sensitive information and issues like service, quality, price, product diversity, sales, etc.

For various basic and simple models for on-line consumers, shops, and profiles, we demonstrate the feasibility of our system, i.e., that proper matchings of consumers with shops are achieved, and that shops can learn their niche in the market, even in the case of interdependencies resulting from competition among similar stores. We use evolutionary algorithms [3] to simulate the adaptive market behavior. We also reflect on the merits of the system, and assess the issues that need further attention, from both the technological and the economical point of view. We note that the mechanism we describe is not limited to the example of the electronic shopping mall, but can easily be extended to other domains where (pre) selection of possibilities has to be guided, like banners on more general websites, or other types of marketplaces.

In this extended abstract, we present the design of *CASy* in Section 2. The simulation model and validation results are described in Sections 3 and 4, and Section 5 concludes. For further details, we refer to the full paper [1].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.


EC’01, October 14-17, 2001, Tampa, Florida, USA.

Copyright 2001 ACM 1-58113-387-1/01/0010 ...\$5.00.

## 2. THE DESIGN OF CASY

Here, we present the framework of CASy (Competitive Attention-space System) for matching online consumers with relevant suppliers in the case of an electronic shopping-mall.

Within an electronic shopping-mall, CASy operates by taking the expressed momentary interest of a consumer, say a product and a business sector, and then presenting a suitable short-list of shops. CASy assembles the short-list via a competitive market-based mechanism. This mechanism works as follows: the information about the consumer’s interest, possibly augmented by additional knowledge, is passed on to potential suppliers. These suppliers subsequently compete against each other in an auction, by each placing bids to “purchase” one of a limited number of entries of attention space for this specific consumer. The consumer is shown the list of winning suppliers, for example as depicted in figure 1.



**Figure 1: Banner-advertisements tailored towards a consumer’s characteristics.**

### 2.1 Agents in a marketplace model

Software agents are used to manage the fine grain of interaction, bidding, and selection in CASy. The system consists of supplier agents and a central manager agent (CMA). The supplier agents purchase attention space by bidding on interesting consumers, whereas the CMA executes the auction process. The system-setup is depicted in figure 2A. We discuss the agents in more detail below.

**Consumers** The consumer communicates his interest and preferences to the CMA, e.g. via a web page. Preferences may be the product that is being searched and various attributes. The CMA can also maintain information on a consumer’s profile, which consists of more generic information, e.g. general interests, previous acquisitions, zip code, and price and quality-level. The profile information can be based on historical data or direct queries, where the consumer can restrict dissemination of this information.

**Central Manager Agent** The CMA enables the selection of a set of suppliers for each arriving consumer through an auction: it collects the bids of the supplier agents, selects the winners, charges the selected suppliers, and enables their display. Furthermore, the CMA provides information from the consumer (see above) to the supplier agents.

**Suppliers** Each supplier “owns” an agent that is equipped with supplier-specific knowledge and an (adaptive) bidding strategy in order to satisfy the supplier’s goal(s). Examples of such goals are maximizing immediate profits or increasing market share. Typical knowledge could include the products carried, intended audience, and current stock or service level. Obtaining an effective bidding strategy can be complicated: the variety of consumers can be great, and the competitive environment can change rapidly. The supplier’s conception of the targeted audience may also deviate

from its actual audience, which is also affected by the competitive environment.

**Auctions** Several types of auction protocols can be applied. In this extended abstract, we use the communication-efficient single-bid sealed auction. With this procedure, each supplier submits a single sealed bid for a particular consumer. The CMA allocates the top banner space to the highest bidder, the second banner space to the next highest bidder, and so on. The winners all pay the  $N+1$  price,  $N$  being the total number of banner spaces available (if less than  $N+1$  suppliers participate, the winners obtain the banner space for free). This protocol is an instance of the Generalized Vickrey Auction (GVA) [5], where buyers are restricted to purchasing at most one good (i.e. banner space). The GVA is a prominent and widely-used auction type, which has shown to be (dominantly) efficient for independent valuations of the items.

### 2.2 Effectiveness and Feasibility

The supplier must find and improve its actual niche in the market, especially in the fine-grained advertisement mechanism CASy. Although the typical business information for the supplier agent can contain many variables that relate to those in a consumer profile, these cannot be matched directly. Similar observations hold even more for the valuation of a consumer.

The need for accurate targeting is especially important when consumers are significantly contested by competing suppliers. We illustrate this by the case of a very expensive department store: consumers arriving in a fancy car are *a priori* as likely to buy at the store as consumers arriving in a middle-class car. However, when a cheaper department store exists across the street, this competition changes the behavior of the latter consumers much more than that of the former. Similarly, in CASy the valuation of an advertisement space can be influenced by which other shops are displayed concurrently. This is a form of interdependent valuation due to allocation.

An  $N+1$  auction mechanism is theoretically efficient in the case of fully rational agents and independent valuations. However, with interdependent values and/or bounded rationality, it is unclear whether an efficient allocation of the attention space will emerge, i.e., a correct match between consumers and suppliers with the largest appearing interests. In practice, this task is even more difficult as the software agents have imperfect knowledge of their environment.

## 3. MARKETPLACE SIMULATION

In this section, we design a simulation to assess the feasibility of the market mechanism of CASy. To this end, we use evolutionary algorithms for simulations, as in the field of agent-based computational economics (ACE, [4]), to study the emergence of efficient banner-space allocation in CASy. Here, we model the behavior of the *consumers* and *supplier agents* for our simulations.

**Consumers** Consumers are modeled through their profile and purchasing behavior. In the simulations in this extended abstract, we use an abstract profile, viz. a two-dimensional vector of real values in a  $[0, 1]^2$  range, the consumer profile space.

**Consumer Behavior** The purchasing behavior of a single consumer for one isolated supplier is modeled as follows. After consumer  $c$  observes the advertisement of a sup-




Figure 2: A) Agent model of the marketplace. B,C) Consumer behavior models 2 and 3.

plier  $i$ , a purchase is made with probability  $P_i(c)$ . The *expected payoff* for the supplier is then given by  $\mu_i P_i(c)$ , where  $\mu_i$  is a constant value related to the supplier's marginal profits. In the simulation,  $P_i(c)$  is a Gaussian surface with width  $\sigma_i$ . The  $c$ -value corresponding to the highest point of the surface is called the supplier's *center of attraction*. In the reported experiments in this extended abstract, we use  $\mu_i = 1$ ,  $\sigma = 0.2$ , 8 suppliers, 3 banners, and maximal purchase probability scaled to unity.

The purchasing behavior of a consumer faced with multiple suppliers is modeled in three types of consumer models:

1. *Independent visits with several purchases.* In this model, the consumer visits all displayed suppliers, and can buy products at several suppliers (e.g. CDs). The expected payoff for each shown supplier  $i$  is then the same as with one isolated supplier:  $\mu_i P_i(c)$ . The consumer purchase behavior for the above settings is depicted in figure 3.




Figure 3: Expected purchases for consumers with a 2-dimensional consumer profile (the (x,y)-plane) at 8 different suppliers.

2. *Independent visits with one total expected purchase.* In this model, a consumer visits all displayed suppliers and then buys on average one product in total (e.g. a computer). A supplier's payoff now depends on other suppliers in the list and equals  $\mu_i P_i(c) / \sum_{j \in \text{list}} P_j(c)$ . The model is depicted in figure 2B.
3. *Search-till-found behavior.* In this model, the consumer visits the suppliers in sequential order from top to bottom, until he finds a supplier with the proper product, which he buys (e.g. a raisin bread). Expected payoff depends both on position and on suppliers higher in the list. For instance, if suppliers  $i$  and  $j$  are placed at positions 1 and 2 respectively, supplier  $j$  obtains on average  $\mu_j P_j(c)(1 - P_i(c))$ , e.g. figure 2C.

**Supplier Agents** A supplier in the simulation is characterized by its *center of attraction* within the consumer profile space. The goal of the supplier agent in the simulation is to maximize immediate profits. The suppliers have no

initial knowledge of their own actual niche, nor of the (type of) consumer behavior. A supplier agent can adapt to the appropriate niche(s) or market segment(s) by learning an effective bidding strategy from actual consumer purchases. The agent thereby indirectly learns the consumer behavior and the competitive environment determined by other supplier agents.

The bidding strategies are represented by piece-wise linear functions and are learned using an evolutionary algorithm (EA, [3]). EA's (which include genetic algorithms) are inspired by biological evolution theory and are based on selection of the fittest solutions to participate in the forming of new solutions. For details on the implementation and settings, see the full report [1].

## 4. RESULTS

We discuss simulation results for the settings outlined above; additional results are reported in the full paper [1].

First, the accuracy of the match between suppliers and consumers is considered. Consider the  $n$  suppliers with the largest expected payoffs for a given consumer. We measure the fraction of these  $n$  suppliers that are present in the actual list of 3 banners displayed to the consumer. For instance, when  $n = 2$ , matching accuracy is 1 if the supplier with the highest and the supplier with the next-highest expected payoff are displayed in two of the three banners, 0.5 if only one of them is displayed, and 0 otherwise.

Table 1 shows the results of the evolutionary simulation (with EA's, section 3) for different values of  $n$  after 500 iterations (generations). Results are averaged over 100 consumer visits per iteration, and 10 runs of the same experiment are executed to obtain the standard deviations. The 100 consumers are equi-distributed over consumer profile space.


Model	$n = 3$	$n = 2$	$n = 1$
1	$0.90 \pm 0.02$	$0.98 \pm 0.01$	$0.99 \pm 0.01$
2	$0.92 \pm 0.01$	$0.98 \pm 0.01$	$1.00 \pm 0.00$
3	$0.75 \pm 0.02$	$0.90 \pm 0.02$	$0.97 \pm 0.01$

Table 1: Matching results for consumer models 1-3.

The results for consumer behavior model 1 and 2 show very efficient allocation of the banners, particularly considering that stochastic effects are always present with evolving bidding strategies. In case of consumer models 2 and 3, obtaining a stable, efficient allocation is more difficult, as the expected pay-off is dependent on the allocation of the other banners: cases where the  $N + 1$ -price sealed bid auction is no longer known to be necessarily theoretically efficient. For model 2, the results match those of model 1. In model 3, the actual ranking on the banner list and that of the other occupants determines the payoff. Nevertheless, the results

nearly match those of the other models for  $n = 2$  and  $n = 1$ , and provide a reasonably accurate match for  $n = 3$ .

Note that model 3 actually entails extreme ‘search-till-found’ behavior, amongst others due to the maximal purchase probability scaled to 1 (see Section 3). More moderate settings yield results in between those for models 1 and 3; we refer to [1].



**Figure 4: Profit per supplier as obtained from 100 consumer visits, during learning in a run for model 2.**


In figure 4, we plotted the obtained profits (payoff) as a function of learning by the supplier agents. For each iteration (generation), the profit obtained from the 100 consumers is plotted. Based on this result, new bidding-strategies are evolved that are tested in the next iteration. As can be seen, after only a few generations of evolutionary learning, all suppliers make a profit that quickly stabilizes in a steady state. The actual profit obtained differs per supplier as each supplier encounters a different competitive environment (see Fig. 5 and Section 4.1).

#### 4.1 Specialization and niches in the market

The simulations show that supplier agents learn to properly evaluate their environment and locate their niche in the market in the case of competition. Figure 5 shows the contour of a supplier’s bidding strategy for bid values 0.5, for two different consumer models, viz. 1 and 2. In the latter model, the expected payoff of a supplier depends on the competition, where in the former it does not. We find that this gives suppliers an incentive to locate their own niches in the market in model 2, and bid more in places where less competition is present. In Figure 5, the depicted supplier clearly expands its market to the upper right in model 2, and reduces its bids in the lower left region, where competition is relatively greater. Furthermore, for the profits obtained by the suppliers, a decreasing payoff with increasing competition can be observed (compare figures 5 and 4). This is in accordance with economics theory: with intense competition in an efficient market, the net profit of competing suppliers tends to zero. We refer to the full paper [1].

## 5. DISCUSSION AND CONCLUSION

We can identify a number of commercial and technological advantages of CASy. In CASy, proper matching does not have to be performed or enabled by a third party. This significantly reduces the combinatorial complexity as compared to centrally processing all product ontologies and information about consumers and shops. Furthermore, shops have substantial autonomy and can thus incorporate local domain knowledge and momentary business considerations



**Figure 5: Contours of the 0.5 bid level from the average evolved strategy of supplier 1 for consumer models 1 (left) and 2 (right). Crosses indicate the suppliers’ centers of attraction.**

in their bidding strategies and thus in the ultimate matching process. Especially, they can take more sales aspects relevant for the consumer into account: not only product pricing, but also service level, quality, product diversity, or customization of products. Also, they do not have to reveal sensitive business information to a third party. The system also enables them to quickly adapt to market dynamics or their own internal situation (out-of-stock, discount periods, promotion). Note that the relevance of the shop for the consumer is expressed and transmitted via the monetary bidding procedure. Finally, we remark that the mechanism is also a form of dynamic pricing of attention space. For further discussion, see [1].

**Conclusion** In this paper, we have presented a competitive distributed system, CASy, for allocating consumer attention space. By evolutionary simulation as in agent-based computational economics (ACE), we have shown the conceptual feasibility of the system. I.e., we modeled the various parts in the system in a basic and simple way suitable for analysis, visualization, and comparison, and showed that proper matchings emerged when suppliers learn their niche in the market. Furthermore, we have designed adaptive software agents for suppliers for the above models and these performed successfully (not presented here). Our system CASy is complementary to existing recommendation systems. It is important to know in what way these together could be used as part of a broader system, and which application areas are more suited for either one of these.

## 6. REFERENCES

- [1] S.M. Bohte, E.H. Gerding, and H. La Poutré. Competitive market-based allocation of consumer attention space. Technical Report forthcoming, CWI, Amsterdam, 2001.
- [2] S.H. Clearwater, editor. *Market based Control of Distributed Systems*. World Scientific Press, Singapore., 1995.
- [3] M. Mitchell. *An introduction to genetic algorithms*. The MIT press, Cambridge, MA, 1996.
- [4] L. Tesfatsion. Introduction to the special issue on agent-based computational economics. *Journal of Economic Dynamics and Control*, 25(3/4):281–293, 2001.
- [5] H.R. Varian. Economic mechanism design for computerized agents. In *First USENIX Workshop on Electronic Commerce*, pages 13–21, New York, 1995.
- [6] G. Weiss, editor. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. The MIT Press, Cambridge, MA, 1999.