

Market-based Recommendation: Agents that Compete for Consumer Attention

SANDER M. BOHTE

ENRICO GERDING

The Netherlands Center for Mathematics and Computer Science (CWI)

and

HAN LA POUTRÉ

The Netherlands Center for Mathematics and Computer Science (CWI), and

Eindhoven University of Technology

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 the example of an electronic shopping mall, the task is delegated to the individual shops, each of which evaluates the information that is available about the consumer and his or her interests (e.g. keywords, product queries, and available parts of a profile). Shops make a monetary bid in an auction where a limited amount of “consumer attention space” for the arriving consumer is sold. Each shop is represented by a software agent that bids for each consumer. This allows shops to rapidly adapt their bidding strategy to focus on consumers interested in their offerings.

For various basic and simple models for on-line consumers, shops, and profiles, we demonstrate the feasibility of our system by evolutionary simulations as in the field of agent-based computational economics (ACE). We also develop adaptive software agents that learn bidding-strategies, based on neural networks and strategy exploration heuristics. Furthermore, we address the commercial and technological advantages of this distributed market-based approach. The mechanism we describe is not limited to the example of the electronic shopping mall, but can easily be extended to other domains.

Categories and Subject Descriptors: C.2.4 [Computer Systems Organization]: Distributed Systems—*Distributed Applications*; I.2.6 [Computing Methodologies]: Learning—*Connectionism and neural nets*; I.2.11 [Computing Methodologies]: Distributed Artificial Intelligence—*Languages and structures*; J.4 [Computer Applications]: Social and behavioral sciences—*Economics*; K.4.4 [Computing Milieux]: Electronic Commerce—*Distributed commercial transactions*

General Terms: Algorithms, Economics

Authors’ address: Sander Bohte / Enrico Gerding / Han La Poutré, CWI, department SEN4, Kruislaan 413, NL-1098 SJ Amsterdam, The Netherlands.

This research has been performed within the framework of the project “Autonomous Systems of Trade Agents in E-Commerce,” funded by the Telematics Institute in the Netherlands. An earlier version has been published as CWI Technical Report SEN-R0131 [Bohte et al. 2001], of which a short version has been published in the Proceedings of the ACM Conference on Electronic Commerce, 14-17 October 2001, pp 202-206, The ACM Press, Tampa, Florida, USA [Bohte et al. 2001].

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0164-0925/20YY/0500-100001 \$5.00

Additional Key Words and Phrases: Electronic Markets, Market-based programming, competitive multi-agent systems, learning agents, ACE, agent-based computational economics, recommendation systems.

1. INTRODUCTION

With the advent of electronic marketplaces, scale limitations as encountered in the brick-and-mortar world no longer apply: the supply side of the market is no longer restricted by geographical considerations or lack of physical (shelf) space. At the same time, novel problems are encountered, like how consumers can find their way in a large marketplace where very many suppliers offer their products.

To this end, a mechanism provided by a central party is desired to propose relevant shops and products to a consumer in e.g. a virtual shopping mall. A central filtering scheme implemented by a trusted third party is a feasible solution for several different business areas. It uses knowledge of both the user and of the shops, as well as knowledge on the product domain to centrally determine the resulting matches and present them to the user. This approach is used in recommender systems like *Amazon* and *eBay* [Schafer et al. 1999] to recommend goods on specific domains such as books and CD's, and in shopbots or pricebots [Greenwald and Kephart 1999], as for instance Bargain-Finder [Krulwich 1996]. A prototype called MATE [Owen et al. 1999] also takes shops into account: merchant agents receive the profile of the consumer, and each suggests one or more products. A personal consumer agent filters the appropriate products and ranks the remaining products according to the consumer's preferences. Also, in this approach, significant knowledge on a product domain should be incorporated in the personal consumer agent, being a task of a central party to provide. Keyword profiling is also a popular method for ranking online sites in search engines. This amounts to contracts for charging monetary amounts for increased visibility, given specific keyword entries, e.g. [<http://www.google.com> ; <http://www.overture.com> ; <http://www.worldmall.com>]. These technologies currently have a static and limited nature.

A central or personal filtering system works well in the case of suitable and well-demarcated domains, as for instance for a book and music store. However, for a large heterogeneous marketplace with many participating shops and consumers, several complexity difficulties arise. This is due to the amount of relevant information that has to be tracked and processed by the filtering mechanism in the form of relevant up-to-date knowledge of e.g.: the consumer's interest in different product domains and shop categories; the shops' products, ways of doing business, and business interest; and ontologies and domain knowledge for various product categories. Also, the weighting of multiple issues like service, quality, price, and product diversity (add-ons and customization of products) can be important.

Besides the computational complexity problems for information processing, this requires the transfer of business information of shops towards the central system as a trusted third party. Such a practice encounters many objections in the daily business practice, even if only product catalogues are concerned [Vermeer 2001; Moukas et al. 1998]. In addition, a central mechanism still needs to make decisions

about what to display in which order to a consumer, in a way that is reasonable to all parties: all the suppliers and consumers. As is commonly agreed on in economics, a fair and general weighing of interests (utilities) of different market parties is not possible, and concepts like Pareto-efficiency are used instead.

Thus, central filtering mechanisms may suffer from increasing (computational) complexity as well as serious objections and obstructions from commercial parties in various sorts of business areas.

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 [Clearwater 1995; Cheng and Wellman 1998; Gibney et al. 1999] and adaptive software agents [Weiss 1999; Guttman et al. 1998; Kephart et al. 2000]. In the example of an electronic shopping mall, CASy recommends shops to a consumer: the task of matching a consumer to a set of suitable shops is delegated to the individual shops, each of which evaluates the information that is available about the consumer and his or her interests (the consumer's interests and other information which the consumer is willing to provide; e.g. keywords, product queries, and available parts of a profile). Based on this information and on their domain knowledge, shops can make a monetary bid in an auction where a limited amount of consumer attention space, or banners, for the particular 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, efficient bidding for each customer is only feasible when automated: hence the use of software agents. These agents allow a shop to process a large number of small transactions, and enable them to make a deliberated bid for every customer entering the shopping mall.

In CASy, shops react to consumer behavior and to behaviors of other shops, yielding various interdependencies in the commercial effects related to being displayed together with competitors (Section 2.7). 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 such interdependencies. Especially, to validate the economical concept of the market mechanism underlying CASy, we develop an evolutionary agent system for bidding supplier agents. In this approach, the agent system is investigated like an economic market, as in agent-based computational economics (ACE). [Tesaftion 2001; Vriend 1995; van Bragt et al. 2001; Gerding et al. 2000; van Bragt and La Poutré 2001; Kirman and Vriend 2001]. To the best of our knowledge this is a novel approach within the field of market-based programming. Furthermore, we also develop adaptive software agents that learn bidding-strategies, based on neural networks and strategy exploration heuristics. Finally, we reflect on the merits of the system, and assess the advantages and 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, attention spaces on mobile devices, or other types of marketplaces.

We believe that the system as presented is the natural evolution of auction-based allocation systems like those currently employed by internet companies like Google (for sponsored keywords, [<http://www.google.com>]) and Overture (for banner targeting, [<http://www.overture.com>]). Whereas these pre-cursor systems rely on the human factor to set essentially static prices for particular goods, the use of software agents in our system in principle allows a market-party to assess the value of each individual prospect, if desired at a very detailed level, as well as take into account real-time business-related domain knowledge and strategies. The implementation of adaptivity into the software agents allows the “market” for consumer attention to function more efficiently, where the targeting of potential prospects can be more precise, and changing buyer behavior can be tracked and followed. As such, agent-assisted recommendation in competitive markets represents the next logical step for current auction-based allocation systems.

This paper is organized as follows. In Section 2, the design of CASy is presented. The simulation model is given in Section 3, whereas Section 4 contains the results. Section 5 concerns the design of software agents using neural networks. Section 6 reflects on practical implementation issues such as privacy and the communication overhead of the mechanism. Finally, Section 7 concludes.

2. THE DESIGN OF CASY

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

2.1 The General Design

We describe the process of the mechanisms in CASy. When a consumer enters a shopping mall, he¹ expresses his interest for certain products and selects the business sector of his interest. The information about his interest, possibly augmented by additional knowledge, is passed on to potential suppliers in the sector. The suppliers subsequently compete against each other in an auction by placing bids to “purchase” one of a limited number of entries of attention space for this specific consumer. Finally, the consumer is shown the list of winning suppliers, using for instance banner advertisements. An example is depicted in figure 1.

2.2 Software Agents

Software agents are used to facilitate the fine grain of interaction, bidding, and selection in CASy. A *software agent* [Nwana and Ndumu 1999; Weiss 1999] is a program with an owner, whom it represents. The software agent can autonomously perform tasks on behalf of its owner in order to satisfy its goals. To this end, it also has the ability to interact with other agents or external sources.

For our mechanism, we have software agents for the suppliers and for the enabling intermediary: the central manager.

¹“he” stands for “he or she.”




Fig. 1. Advertisements are shown in the form of banners. The banner list is tailored towards a consumer's characteristics.

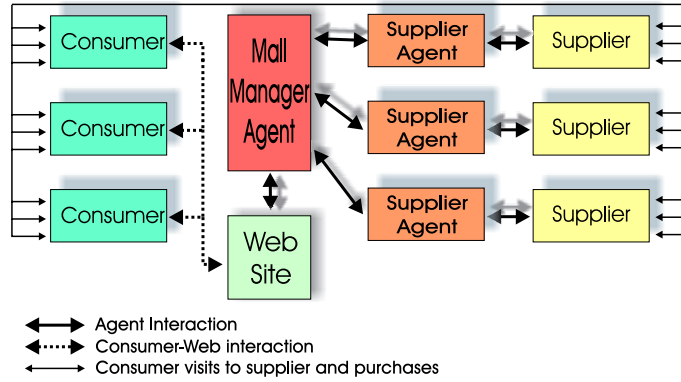


Fig. 2. Components of the shopping mall and their interactions.

The goal of an agent representing a supplier is to effectively purchase attention space. The agent will do this by bidding on attention spaces that are to be displayed to consumers it deems interesting, thus maximizing the supplier's profits. The agent can learn this targeting by for instance using the push-back information from individual customers, e.g. the knowledge whether or not its advertisement was selected by the customer (click-through), subsequent buying actions, or, to be provided by the central manager, (selected) click-stream information (e.g. time spent on pages, mouse actions). Additionally, the agent can use supplier-specific knowledge and (adaptive) rules for accurate targeting.

A central manager agent (CMA) facilitates these bidding and information dissemination processes by providing the auctions and additional customer profiling services to the suppliers.

The model of the electronic shopping mall is depicted in figure 2, showing both the software agents and the actual economic players in the shopping mall: the consumers and the suppliers. Note that although the consumer directly communicates with the mall manager agent (CMA), the assistance of a personal software agent for the consumer is conceivable.

Instead of addressing to the case of "shops" only, we henceforth mainly use the

more general term “supplier” to refer to the suppliers of goods or services. The participants within the shopping mall and their roles are discussed in more detail in the sections that follow.

2.3 Consumers

In the model of figure 2, the consumer directly communicates its interest and preferences to the CMA, e.g. via a web page. Preferences include the product that is being searched after and various values for the attributes of the product. The CMA can also consider information on a consumer’s profile. The consumer profile consists of more generic information on the consumer. This could include regular personal information like general interests, previous acquisitions, as well as age or zip code; but also general sales-related information like style or the interest in issues as price, quality, and service. The consumer can either be queried directly for this information, or the CMA can derive the information from previous interactions. The consumer can restrict or disable the dissemination of his profile information. E.g., distribution of such information can be limited to for specific or anonymized parts, or to general sales-related information that is derived from the private profile.

2.4 Central Manager

The Central Manager Agent (CMA) acts as an intermediary between consumers and supplier agents. The task of the CMA is to enable the selection of a set of suppliers for each arriving consumer. The CMA furthermore provides information from the consumer to the supplier agents. Given privacy concerns, the consumer profile will not automatically be communicated in full to the suppliers, as e.g. described in Subsection 2.3. Information on the consumers could be stored within the CMA for revisiting consumers, leaving open consumers who wish to remain anonymous. The CMA applies the auction: it collects the bids of the supplier agents, selects the winners, charges the selected suppliers, and enables their display. In Subsection 2.6, we address the auctions in more detail.

2.5 Suppliers

Each supplier “owns” an agent that acts on the supplier’s behalf. These agents are equipped with knowledge and a strategy on behalf of the supplier. Such knowledge can contain amongst others relevant business information on the supplier that is needed for the matching process. This information should determine the supplier’s conception of its “niche” in the market, and hence the type of preferred consumer. Typical business information could be the products carried and the intended audience. Furthermore, the goals and limitations of the supplier can be taken into account, such as the current quantity of a certain product in stock or the service level.

The main task of a supplier agent is to bid on arriving consumers. To this end, it has to evaluate (information about) consumers. Namely, the valuation of a consumer by a supplier agent is closely linked to its bidding strategy: the bid should not outweigh the expected profit (if the supplier is to break even) or percentage thereof. This task can be complicated: the variety of consumers can be great, and the competitive environment can change rapidly. Also, the supplier’s conception of the targeted audience may deviate from its actual audience.

2.6 Auctions

The CMA executes the auction protocol, the payment procedure, and the supplier selection mechanism.

Auction protocol. The actual choice of the auction protocol can depend on many factors. In this paper, we focus on the single-bid sealed auction, being a communication-efficient auction. With this procedure, each supplier submits a single sealed bid for a particular consumer. The CMA allocates the available positions to the highest bidders, where the first position is allocated to the highest bidder, the second position to the second highest bidder, and so on. In some environments the ranking is not important, whereas in other cases the profits for the supplier depend on the position obtained. For this reason, the choice payment of payment scheme matters, and is discussed below. Note that, since the CMA executes the auction for each arriving consumer, suppliers losing an auction could increase their bid in the next auction for a similar consumer.

Payment procedure. A payment procedure specifies *what* should be charged and *when*. Several different payment schemes are possible for various auction procedures. In the Vickrey auction, the winner pays the price of the second-highest bid. This is a prominent and widely-used auction type, which has been shown to be efficient for independent valuations of the item [Vickrey 1961; Dasgupta and Maskin 2000; Varian 1995]. The auction is also robust, since revealing ones true preferences is the dominant strategy in case of independent valuations.

In this paper, we focus on an extension of the Vickrey auction where winners pay the $(N+1)$ price, where N is the number of items (here banners). This is an instance of the generalized Vickrey auction [Varian 1995], which has the same auction characteristics as above.

Note that in such a setup, the same price is charged to the winners of a banner placement. The auction is only theoretically guaranteed to work well if the sold goods (the attention spaces) are assumed to be identical, an assumption that is dependent on the way a customer chooses from a list of alternative offerings. In the simulations, we investigate models of customer behavior where this assumption is valid, as well as a model where it does not hold. In the latter case, we also investigate another payment scheme, the so-called *next-price* auction. Here, each winner pays the price of the next-highest bidder. Such more complicated auctions are notoriously hard to theoretically demonstrate optimal behavior for, and we use the ACE methods (as discussed later) to show that in the simulations this auction does work efficiently in the case where the valuation of an attention space depends on the position it has on the list and when the highest position is the most valuable (and the second-highest position is the next most valuable etc..).

2.7 Effectiveness and Feasibility

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. Rather, the supplier must find and improve its actual niche in the market, especially in the fine-grained advertisement mechanism CASy. Similar observations hold even more for the valuation of a consumer.

The need for accurate valuation and targeting is especially pronounced 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 of the former. Similarly, in CASy the valuation of an advertisement space depends on the selection of and competition between suppliers.

An N+1 auction mechanism is theoretically efficient in case of fully rational agents and independent valuation of the items. However, if consumer purchases are like consumer models 2 and 3 (see also Section 3.5), the valuation of advertisement space also depends on the selection and competition between various suppliers. It is then 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 for being displayed together.

In the following, we will show via evolutionary simulation as in the field of agent-based computational economics (ACE, [Tsfatsion 2001]) and by implementations of software agents, that the market mechanism is indeed effective and results in an efficient allocation. Furthermore, supplier agents learn to properly evaluate their environment and thereby locate their niche in the market.

3. EVOLUTIONARY SIMULATION MODEL OF CASY

In this section, we model the electronic shopping mall for an evolutionary simulation as in ACE, based on Section 2. The goal of the simulation is to assess the feasibility of the market mechanism of CASy (see Section 2.7). To this end, we will make some additional assumptions and simplifications, which enables us to study, measure, and visualize the emerging behavior of CASy (results are given in Section 4).

3.1 Central Manager Agent

The CMA has 3 banner advertisements to dispatch (see also figure 1), and executes the auction as described before.

3.2 Consumer Models

We abstract away from any interpretation of the profiles. Profiles are represented by a vector of real values. In the simulations, the consumers are classified by a one or two dimensional vector with entries in a $[0 \dots 1]$ range. The profile can reflect a consumer's interests such as price segment, taste, or quality, or any combination of characteristics projected on 1 or 2 dimensions. We thus model a class of consumers for some given category of products. In the simulation of CASy, several consumers with different profiles arrive and are contested by the suppliers in CASy.

3.3 Supplier Models

We will denote by *gross profit* the profit that a supplier earns on a product, before the cost of advertisement is taken into account (but after accounting for all other costs), and by *net profit* the profit after deduction of all costs, including advertisement cost.

The goal of a supplier is to maximize net profits, and therefore a supplier tries to sell as many items as possible at the lowest possible advertising costs. The net profit of a supplier is also referred to as the supplier's *payoff*. The suppliers in the simulation have no initial knowledge of their own actual niche or payoff function in the market (see Section 2.7).

A bidding strategy specifies the monetary bid for each possible consumer profile. Given the feedback in the form of actual payoff for visiting consumers, a supplier agent adapts its bidding strategy and thereby indirectly learns the consumer behavior and its competitive environment determined by other supplier agents. Note that these two factors are interrelated (see also Subsection 3.5).

Evolutionary Simulation. We use evolutionary simulation like in the field of agent-based computational economics (ACE) [Tesauro 2001; Vriend 1995; Gerdling et al. 2000; van Bragt et al. 2001; van Bragt and La Poutré 2001; Alkemade et al. 2003], where suppliers that interact and compete in a market are evolved, in order to investigate their emerging behavior and the equilibrium situation. Recall that a supplier's goal is to maximize payoff.

We proceed as follows. Each supplier agent is replaced by a population of strategies. These strategies are evaluated and evolved according to the amount of profit they earn in single CASy simulation. In such a CASy simulation, a number of consumers arrive, supplier strategies bid for each of these, and the winners get the expected payoffs as described in Subsections 3.4 and 3.5. The strategies that are evolved after repeating this process many times, show the emerging behavior of the suppliers. Hence, the process of evolution finds effective strategies for a CASy simulation.

An evolutionary algorithm (EA) is used to adapt the strategies of the supplier agents. EAs are strongly inspired by the genetic evolution theory in biology, as developed by Darwin. EAs typically work as follows. First, for each supplier a population of randomly initialized strategies is generated. The populations are subsequently changed and improved in a number of iterations ("generations") by means of selection and mutation. Selection chooses the better strategies (with higher accumulated payoff) which survive in the next generation. This corresponds to the concept of "survival of the fittest" in nature. The selected strategies are subsequently changed slightly in a random way ("mutation"), to enable diversity in the population.

Our implementation is based on "evolution strategies" (ES), a branch of evolutionary algorithms that traditionally focuses on real-coded problems [Bäck 1996].² We use standard parameter settings for EAs (see Appendix A).

3.4 Buying Behavior Model for One Consumer and one Supplier

We model the purchasing behavior of a single consumer for one isolated supplier. For each supplier i , the *expected gross monopolistic profits* $E\langle\pi_i(c)\rangle$ is its average gross profits for a possible purchase following the observation of a consumer of its

²The widely-used genetic algorithms (GAs) are more tailored toward binary-coded search spaces [Holland 1975; Mitchell 1996; Goldberg 1989].

advertisement, while no other supplier is shown. We take

$$E\langle\pi_i(c)\rangle = \mu_i P_i(c),$$

where $P_i(c)$ denotes the monopolistic purchase probability for consumer profile c and μ_i is a constant value related to the supplier's average profit when a purchase is made. Note that both μ_i and $P_i(c)$ are taken as an externally imposed model for interaction and are initially not known or available to the supplier.

In the simulation each supplier is given a *center of attraction* a_i , where $P_i(c)$ is maximized. We used two types of purchase probability functions P_i in the experiments: (1) linear functions, where the P_i is proportional to the Euclidean distance $d(c, a_i)$ in the following way:

$$P_i(c) = 1 - \delta d(c, a_i),$$

and (2) Gaussian functions with the highest point corresponding to the center of attraction. The width of the Gaussian curve is then set by parameter σ_i . For simplicity the maximal monopolistic purchase probability is set constant to 1. This value can be chosen lower, but is chosen for maximal discrimination between various advanced behavior models (see Subsection 3.5).

3.5 Buying Behavior Model for Several Displayed Suppliers

As the consumer is presented with a selection of winning "Consumer Attention Spaces", we assume that with some probability p he or she will buy a product. In effect, this stochastic behavior can be modeled as meaning that a single presentation of banners results in an amount p of products being sold: how much and at which recommended supplier (the buying behavior) is formalized in the Customer Buying Behavior Models. Here, we present several Customer Buying Behavior Models, as the behavior of consumers shopping for a specific product may be different for different product areas or different consumer populations.

We modeled three classes of consumer behavior:

- (1) *Independent visits with several purchases.* In this model (see figure 3), the consumer visits all displayed suppliers, and can buy products at several suppliers (e.g. CDs).
- (2) *Independent visits with one expected purchase.* In this model (see figure 4), a consumer visits all displayed suppliers and then buys on average one product in total (e.g. a computer).
- (3) *Search-till-found behavior.* In this model (see figure 5), 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).

The consumer behavior in these models is stochastic: whether a product is purchased by consumer c at a certain supplier j depends on a probability value $Q_j(c)$. The monopolistic purchase probabilities $P_i(c)$ are the basic parameters, determining these probability values $Q_j(c)$ as shown in figures 3 to 5. The expected gross profits $E\langle\rho_j(c)\rangle$ for supplier j is then given by

$$E\langle\rho_j(c)\rangle = \mu_j Q_j(c).$$




Fig. 3. Consumer model of independent visits with several purchases, where $P_i = P_i(c)$.




Fig. 4. Consumer model of independent visits with up to one purchase, where $P_i = P_i(c)$.




Fig. 5. Consumer model with search-till-found behavior, where $P_i = P_i(c)$.

Notice that in the models of figure 4 and 5, the probability that an item is sold at one supplier depends on the monopolistic purchase probabilities of its competitors within the list. Importantly, for the third model, the actual position of a supplier on the list influences the expected average proceeds, meaning that the individual banners are no longer identical. We will address this issue, and a solution, in detail in Section 4.5.

Parameter	Value
Number of suppliers	8
Number of banner spaces (N)	3
Maximum bid value	1.5
Consumer behavior model	1 / 2 / 3
Expected gross monopolistic profit ($E\langle\pi\rangle$)	set1 / set2/ set3
Profile dimensionality	1 or 2
Number of defining points	8 (1 dimension), 16 (2 dimensions)
Number of consumers	50 (1 dimension), 100 (2 dimensions)

Table I. Default settings of the simulations.

$E\langle\pi\rangle$ function name	Type	μ_i	δ	σ
Set1	Linear	1.0	2.0	-
Set2	Gaussian	1.0	-	0.2
Set3	Linear	variable	variable	-

Table II. Consumer purchase functions and their general settings.

3.6 Measure for Proper Selection of Suppliers in an Auction

The selection procedure in an auction should ultimately lead to an appropriate selection of suppliers for consumers. We start from the economic point of view of optimizing the revenue of the collection of shops in the shopping mall as a whole.

Consider the n suppliers with the largest expected payoffs for a given consumer. We measure the *proportion* of properly selected n suppliers as the fraction of these n suppliers that are present in the actual list of 3 displays shown to the consumer.

From the consumer point of view, we can interpret the expenditures of a consumer at a supplier as a measure for his interest in the supplier. In case that the ratio between expenditures and payoff within a certain business sector is similar for the suppliers in that sector, the above measure is related to both the consumer interests as well as the supplier interests.

4. RESULTS

We performed a number of experiments in the e-shopping-mall simulation outlined in Section 3. The results are given and discussed in the following subsections.

4.1 Simulation Settings

Table I shows the parameters and their values which are varied for different simulation runs. The parameters refer to Section 3. Two of the parameters are further explained below.

—*Expected gross monopolistic profit functions ($E\langle\pi\rangle$)*. The $E\langle\pi\rangle$ -functions are explained in Section 3.4. The applied settings are specified in table II. Figure 6 shows the functions “set2” for 8 different suppliers and a one-dimensional consumer profile. The functions defined in “set3” have different μ_i and δ combinations for each supplier; μ_i varies between 0.5 and 1.0, and δ between 1.0 and 2.0.




Fig. 6. Expected gross monopolistic profits at the different stores for “set2” function settings.

—*Number of defining points.* A supplier has to obtain a bidding function on the space of consumer profiles. The function that is learned is an interpolation function, based on a number of defining points. For the one-dimensional case, this results in a piecewise linear function; for the two-dimensional case, we obtain the function values by triangularisation of the profile surface.

4.2 Single Advertisement Model

In this subsection, we illustrate the use and evolution of the bidding function for a supplier for a very simple setting, where the optimal bidding strategy is known from auction theory.

The setting contains a single store competing against a random opponent for the case of one banner. The random player bids any random value between 0 and 1.5. Since a Vickrey (second-price) auction is used, it is a well-known dominant strategy for the supplier to bid its true valuation (i.e. the expected gross profit) [Vickrey 1961]; any lower bid risks a missed profit-opportunity, whereas a higher bid might result in direct loss. The dominant strategy maximizes the supplier’s net profit, regardless of the opponent’s behavior. Thus, the store should learn the profit function as the bidding function. The results for experiments on this setting show that this happens indeed. Typical, good results are shown in figure 7, where $E\langle\pi\rangle$ is a Gaussian (recall that piecewise linear functions are used).

4.3 Consumer Model 1: Independent Visits with Several Purchases

This consumer model assumes that expected purchases at each supplier can be modeled by the same function as in the single banner case (see Subsection 3.5).

The results are shown in figure 8. Matching accuracy is measured in several ways. We display the proportion of properly selected n suppliers for 3 banners and $n = 3, 2, 1$ (see Subsection 3.6). The reason for including $n = 2, 1$ as well is that the evolutionary system has some degree of stochasticity, and thus small errors occurring frequently can have larger influence on individual outcomes (although relatively little impact on the payoff obtained). Results using these two measures




Fig. 7. Example of a bidding strategy as employed by the supplier after co-evolution no longer increased the profits obtained. Results are shown for a single supplier competing against random supplier. Also shown is the dominant bidding strategy.

Consumer model	$E\langle\pi\rangle$	$n = 3$	$n = 2$	$n = 1$
Regular auction settings				
1	set1	0.95 ± 0.01	0.99 ± 0.00	0.99 ± 0.00
	set2	0.96 ± 0.00	0.99 ± 0.00	1.00 ± 0.00
	set3	0.92 ± 0.01	0.98 ± 0.00	0.99 ± 0.00
2	set1	0.94 ± 0.01	0.99 ± 0.00	0.99 ± 0.00
	set2	0.95 ± 0.00	0.99 ± 0.00	1.00 ± 0.00
	set3	0.90 ± 0.01	0.97 ± 0.01	0.99 ± 0.00
3	set1	0.73 ± 0.03	0.76 ± 0.07	0.79 ± 0.09
	set2	0.83 ± 0.05	0.89 ± 0.06	0.92 ± 0.05
	set3	0.75 ± 0.02	0.89 ± 0.02	0.97 ± 0.01
Next-price auction				
3	set1	0.79 ± 0.03	0.92 ± 0.03	0.97 ± 0.02
	set2	0.75 ± 0.03	0.92 ± 0.02	0.98 ± 0.01
	set3	0.83 ± 0.02	0.95 ± 0.02	0.99 ± 0.00

Table III. Matching results for consumer models 1 through 3. Results denote proportions of properly selected n suppliers for 3 banners and $n = 3, 2, 1$. Averages over 10 runs of the simulation are shown with the standard deviations.

show an almost perfect match. The results after 500 generations of the EA are summarized in table III.

4.4 Consumer Model 2: One Expected Purchase

It is more difficult to get a stable system in this situation, since the expected amount purchased at a supplier (and therefore the valuation of a banner space) depends on which other stores are selected as well. Nevertheless, the simulation does stabilize, and the results are comparable to the previous consumer model. See table III.

4.5 Consumer Model 3: Search-Till-Found

In this model, it is not only important for the stores to be in the list, but also to take into account the position on the list (and the other stores above him). Table III shows that it is indeed more difficult for the stores to find a good matching, in




Fig. 8. Matching results for consumers with independent purchases and $E(\pi)$ is set to “set2”.

particular when using “set1”. This occurs since all relevant suppliers prefer the very top advertisement space and are willing to bid above their valuation (because of the $N + 1$ -price auction their payment remains relatively low). As a result, the bids reach their limit value (even when this is set to 2.5).

Therefore, we have applied another auction payment procedure as well: each of the winning stores pays the price offered by the next following highest bidder, the so-called next-price auction. This procedure appears to improve the matching, giving comparable results to other consumer models (see table III). Note that a store who obtains the first banner position now pays more than the other stores. This is also reasonable, since the first position is actually more valuable.

We want to remark that we have chosen the maximal purchase probability to 1 (see Subsection 3.4) to have maximum difference between this consumer model and the previous ones. When this value is lower, results will become more comparable to the other models also for the regular auction setting.

4.6 Two-Dimensional Profile

We now consider the two-dimensional case, where each consumer profile corresponds to a position within a square. The types of profit functions are similar to the previous case, extended for two dimensions. An example is shown in figure 9.

The matching results are comparable, but slightly less accurate than for one dimension, see table IV. These can be explained through the more difficult learning problem (more defining points are needed for the search function), and thus the settings of the evolutionary algorithms could be further optimized for more accurate learning results in this case (which is beyond the scope of the current paper).

Specialization. Interestingly, the suppliers indeed find a niche in the market in case of competition. This becomes clear in figure 10, which shows the intersection of a supplier’s bidding strategy for two different consumer models, viz. 1 and 2. For consumer model 1, a supplier’s payoff is independent of the other suppliers displayed. In the second consumer model, however, the payoff is shared amongst the displayed suppliers. In the latter model the payoff thus depends on the competition.

Consumer model	$E\langle\pi\rangle$	$n = 3$	$n = 2$	$n = 1$
1	set1	0.95 ± 0.01	0.99 ± 0.00	1.00 ± 0.00
	set2	0.90 ± 0.02	0.97 ± 0.01	0.99 ± 0.01
	set3	0.93 ± 0.01	0.98 ± 0.00	0.99 ± 0.00
2	set1	0.94 ± 0.01	0.98 ± 0.00	0.99 ± 0.00
	set2	0.92 ± 0.01	0.98 ± 0.00	1.00 ± 0.00
	set3	0.93 ± 0.01	0.98 ± 0.00	0.99 ± 0.00
3	set1	0.85 ± 0.01	0.93 ± 0.01	0.97 ± 0.01
	set2	0.75 ± 0.02	0.89 ± 0.02	0.97 ± 0.01
	set3	0.82 ± 0.02	0.91 ± 0.02	0.94 ± 0.02

Table IV. Matching results for consumers with two-dimensional profiles. See also table III for comparison.




Fig. 9. Expected gross monopolistic profits $E\langle\pi\rangle$ for “set2” function settings and a 2-dimensional consumer profile.




Fig. 10. Contours of the average evolved strategy at level 0.5 of a supplier 1 at generation 500 for consumer models 1 (left) and 2 (right) using “set2”. The points indicate the centers of attraction of the suppliers’ Gaussian curves.

We find that this gives supplier an incentive to locate niches in the market, and bid more in places where less competition is present. In figure 10, the depicted supplier clearly expands its market to the upper right, and reduces its bids in the lower left region, where competition is relatively greater.

Supplier Payoff. The above results mainly focus on the proportion of proper selection. We now briefly discuss the supplier payoffs, i.e. the net profits (see




Fig. 11. The average accumulated payoff for each supplier using “set2” and consumer model 2.

Section 3.3). Firstly, we find that in all experiments suppliers obtain positive accumulative payoff in the long run. The strategies emerged are thus individually rational. Secondly, a supplier’s payoff depends both on its function settings $E\langle\pi\rangle$ and on the amount of competition. The latter is shown in figure 11, which displays the accumulated payoff of the suppliers for consumer model 2 and “set2”. The more isolated suppliers, in particular suppliers 4, 6, and 7, obtain a larger payoff than those with much competition (see also figure 10). This is due to the difference in advertisement costs. Note that this is in accordance with economics theory: in case of large competition, the net profit of competing suppliers is close to zero.

4.7 Conclusion

The experiments show that a proper selection of suppliers emerges with very good to perfect matches. In case consumer model 3 is applicable, a next-price auction mechanism further improves the results. Furthermore, we find that all experiments show positive supplier payoffs. Finally, we observe that shops find their customers and their niche in the market via CASy.

5. SOFTWARE AGENTS WITH NEURAL NETWORKS FOR ONLINE AUCTION LEARNING

5.1 Online-learning with Software Agents

In this section, we briefly describe the development of adaptive software agents that can perform online learning from a repeated general Vickrey auction, and we show some of the results we obtained with this adaptive approach based on neural networks.

First we remark that for online-learning, we deal with a variant of reinforcement learning, e.g. [Sutton and Barto 1998]. In the $N + 1$ -price sealed-bid auction, learning signals constitute of the (average) payoff generated by a winning bid for a particular consumer profile, as well as the information that a losing bid is below the winning bid, or “going-price”, in the market. Note that receiving payoff is in principle a stochastic process. Since payoff is normalized in the simulation, this

stochastic nature can be expressed by taking the instant payoff as a discrete value $\in \{0, 1\}$, which averages to the payoff we defined in Section 3.

For a given consumer profile, our agent generates two estimates: the expected payoff (the “value”) and the expected going-price; in addition, uncertainties associated with these expectations are calculated. These values are then combined into a resulting bid according to a heuristic algorithm that balances the exploitation of accumulated information versus exploration aimed at reducing uncertainty in the estimates. When exploiting, the agent bids the estimated payoff, as bidding the actual payoff is a dominant strategy in the $N + 1$ -price auctions considered in Section 5.2.

The algorithm was implemented with two ensembles (sets) of neural networks (multi-layer perceptrons, although alternative architectures could also be used, e.g. [Bohte et al. 2002]) in each software agent, where each ensemble of networks acts as a function-approximator that learns respectively the expected payoff and the going-price, both as a function of the consumer profile. By using ensembles of neural networks, we can use existing techniques for estimating the uncertainty in the respective function-approximations by the neural network ensembles [Heskes 1997]. The uncertainty in the estimates constitutes an important ingredient in our heuristic for learning from losing bids.

5.2 Results with Adaptive Neural Software Agents

For the consumer models 1 and 2 (Section 3.5), it is easy to see that bidding the actual payoff by a shop is a dominant strategy. Within the shopping-mall simulation outlined in Section 3 and these consumer models, we performed a number of experiments to test whether the shop-agents endowed with neural networks are capable of learning the correct valuations from the second-price auctioning of consumer-profiles.

For all examples tested, we found that the agents accurately learned the payoff profiles, both for one and for three banners, a stochastic payoff or averaged payoff, and various numbers of competitors. We observed that the exploration-expenditure stabilized to a small fraction of the revenues after the initial learning phase of typically 50 consumers. After this time, all shops become (accumulated) profitable and generate accurately targeted bids. The time needed for learning was very short: on average it took less than 50 consumers to visit the mall for the shops to learn which consumers are profitable; this held for all simulations we performed, with up to 8 competing shops. An example of online learning for bidding on three available banners is shown in figures 12, 13. The results shown are for the case where for every winning bid the associated average payoff was returned (the case of stochastic payoff took somewhat longer to converge).

6. EVALUATION AND FURTHER RESEARCH

6.1 Reflections

In the previous sections, we have presented CASy and showed its feasibility.

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 cen-




Fig. 12. The consumer valuation as learned by the shop-agents (solid lines) after bidding for 200 consumer profiles, and the actual market valuation (dotted lines), for consumer-behavior model 2, as in Section 3.5.




Fig. 13. Shop-selection resulting from the submitted bids. Plotted is the proportion of properly selected n suppliers (Section 3.6) for the 200 sequential consumers for three banners and $n = 1, n = 2$ and $n = 3$. Regularly one out of three matching shops is “ousted”, but given the low payoff for third place in our experiments (third consumer-valuation or dotted line in fig. 12), the third-highest bid is easily exceeded by even minimal explorations by other shops.

trally 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 in their bidding strategies and thus in the ultimate matching process. Especially, they do not have to reveal sensitive business information to a third party, and can take more sales aspects into account: not only product pricing, but also service level, quality, product diversity, or customization of products. 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 still expressed via the monetary bidding procedure. The mechanism is also a form of dynamic pricing of attention space.

There is much debate about whether or not advances in Information Technology (IT) will increasingly make intermediaries within markets redundant (disintermediation) [Gellmann 1996], or whether such advanced IT will help reestablish intermediaries because of new value-added services that become possible (reintermediation) (cf Chircu and Kauffman, 1999 [Chircu and RJ. 1999]. The results in this paper can be taken either way: on the one hand, we can conceive the basic auction functionality performed by the CMA to be part of the customer agent, replacing the matching function previously performed by central filtering mechanisms. Alternatively, we noted that there are many possible value-added services regarding user-profile enhancement that could be performed by a central shopping mall intermediary. This conclusion is in line with recent arguments regarding the effects of current agent technology on the disintermediation/reintermediation debate [Nissen 2000].

The proceeds the electronic shopping mall can derive from the matching mechanism (through the auctions) can be used to facilitate additional intermediation services to both customers and shops (e.g. micropayments, 24x7 intermediation). Offering an effective matching mechanism adds considerable value to the customer experience, and can thus be expected to be an important selling point for the electronic shopping mall, and entice suppliers to participate in the mechanism. It will be interesting to investigate the exact economic conditions – such as at which price the suppliers are no longer prepared to follow the customers – for this to be relevant, but we leave that for future research here.

Some points need attention when further implementing CASy. In CASy, information about a consumer is (partially) communicated to suppliers. At the same time however, the consumer's privacy requirements must be respected. We will not extensively address this here, but just mention some approaches: having the consumer decide what information he allows to be communicated, restricting the types of communicated information in general, or conversion of personal information to more sales-related properties. The latter could include restricting the profile to attributes of the desired product (instead of the customer), like “expensive vs. cheap”, “ultra trendy vs. conservative” etc.... Such attributes could in principle even be queried from the customer. As argued in [Kobsa 2001], no uniform solution for privacy demands exist, rather “privacy will have to be dynamically tailored to each individual user's needs” and requirements.

There remains the issue whether a central entity like the shopping mall would be willing to convey individual user related profile information. Google for instance

currently considers its clickstream information a business secret. In the setup we introduced in the paper however, the proceeds that the intermediary obtains from the ongoing auctions, and possibly for additional advanced IT services, will be a strong incentive for the intermediary to consider what parts of the profile information are allowed to be disseminated by its clients (here, the suppliers). Note that when the intermediary charges a (fixed) price for customer profile information services, such information would constitute a sunk cost for each seller, and reduce the available funds for placing advertisements, resulting in lower bids. Since such cost will reduce *all* bids from *all* agents, the relative ordering of the bids remains intact and the market-based selection mechanism itself is not affected by such additional cost.

We remark that once individual shops receive customer (related) profiles, they have the tools developed in information intensive personalized marketing research at their disposal for determining how interested they are in each individual customer: i.e. interactive marketing, database marketing, micromarketing and one-to-one marketing [Montgomery and Srinivasan 2002; Montgomery 2001; Blattberg and Deighton 1991; Haeckel 1998; Peppers et al. 1999]; in [Shaffer and Zhang 2002] these slightly different approaches are considered in more detail. The information filtering mechanism we describe in this paper is then the gate controlling the flood of finely targeted business interests.

Another point concerns the communication between suppliers and shopping mall, which is increased because of the bidding process and the communication of consumer profiles to the suppliers. However, the communication in the mall is linear in the number of customers, and also in the number of participating shops, and the size of the consumer profile. The latter is also typically very small, e.g. up to 100 bytes. In a prototype implementation on a single PC, a single market comprising of 100 *learning* shop-agents was easily able to sustain 100 customers per second, and still continuously update the internal state of the agents (the learning mechanism) [’t Hoen et al. 2002].

To scale to even larger settings, the market can be divided into a number of segments, with each market handled by different agents. The profile then only needs to be transmitted to agents within a particular market segments, reducing the overall communication. We pursued this approach in a distributed prototype of the electronic shopping mall ([’t Hoen et al. 2002]). In the extended agent architecture of the prototype different market segments are handled by sub agents (which can run on different machines). In all, we do not perceive the somewhat increased communication as a significant problem, but rather as an issue that can easily be addressed in the process of framework-engineering if necessary.

6.2 Open Problems and Future Research

We investigated the concept of CASy for several basic models. The results we describe here show that the market-based approach yields excellent buyer-seller matching given adaptation of the bids made by the sellers. The ACE simulations have been carried out to demonstrate feasibility and learnability of the concept, as these simulations showed effective matching for different auction types and consumer behavior models. It is also interesting to investigate how software agents can be developed for more advanced settings: one such example would be the extension

of the simulations to a dynamical market, with sellers changing their profiles, or sellers entering and leaving the market. For ACE feasibility and learnability studies, methods that can deal with such dynamic environments are only just starting to emerge. As we demonstrated that the steady state version of the problem per-se is both effective and learnable, we would expect that dynamic versions of the problem would also be learnable, but the effectiveness is then rather dependent on the speed and quality of the machine learning techniques employed by the shop agents as well as the actually chosen models for the dynamic environment. E.g., for methods such as neural networks, the introduction of dynamics into the market will mean that additional complexity in terms of effective (commercial) exploration/exploitation strategies has to be introduced. At this point we leave the investigation of dynamization of the system as an interesting problem for future research.

Other points that need to be addressed in future work should be concerned with taking account of the role of (local) ontologies, of marketing and data-mining techniques, and of partial consumer information. Furthermore, in this work, we placed an emphasis on the $N+1$ -price auction with single sealed bids. Other types of auctions could be further investigated, for example addressing the possible feedback given on bids of other participants (e.g. multi-round auctions) or to address the revenue of the central manager.

From the consumer's point of view, we have interpreted the expenditures of a consumer at a shop as a measure for his interest in the shop. CASy gives priority to suppliers with the largest expected payoffs for a given consumer. This thus leads to optimization of the revenue of the collection of shops in the shopping mall as a whole. In the case that within a certain business sector, the ratio between expenditures and payoff is similar for the suppliers in the sector, this means that CASy completely reacts on the interest of an individual consumer. However, across different sectors, there may be differences or anomalies, leaving the extension of CASy with additional (monetary) correction mechanisms to avoid such anomalies as an interesting open problem. This is part of our future work.

Finally, 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. Also, which application areas are more suited for the existing recommender systems, and which for the CASy system.

7. CONCLUSION

In this paper, we have presented a competitive distributed system, CASy, for allocating consumer attention space (Section 2). By evolutionary simulation as in agent-based computational economics (ACE), we have shown the conceptual feasibility of the system (Sections 3 and 4). 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 while suppliers could learn their niche in the market. Furthermore, we designed adaptive software agents for suppliers for the above models and showed that these performed successfully (Section 5). Finally, we reflected on the advantages, opportunities, and further open problems concerning the proposed system (Section 6).




Fig. 14. Iteration loop of the evolutionary algorithm. First, offspring are created from the parental population using mutation and recombination operators. The best candidates of the union of parents and offspring are then selected to be the parents in the next iteration.

EA	Parental population size (μ)	25
Parameters	Offspring population size (λ)	25
	Selection scheme	$(\mu + \lambda)$ -ES
	Mutation model	self-adaptive
	Initial standard deviations ($\sigma_i(0)$)	0.1
	Minimum standard deviation (ϵ_σ)	0.025

Table V. Default settings of the evolutionary system.

APPENDIX

An evolutionary algorithm (EA) is used to evolve the bidding strategies of the supplier agents. The implementation is based on “evolution strategies” (ES), a branch of evolutionary computation that traditionally focuses on real-coded problems [Bäck 1996].³ Below we first outline the general functioning of the EA and we then provide further details.

The different stages within an iteration of the evolutionary algorithm are depicted in Fig. 14. The system initially starts with a randomly initialized “parental” population for each supplier within the system. Each individual within a population is defined by its chromosome, which represents the bidding strategy used (further details are provided below).

In the next stage (see Fig. 14), “offspring” individuals are created. An offspring agent is generated in two steps. First, an individual in the parental population is (randomly, with replacement) selected. This agent’s strategy is then mutated to create a new offspring individual (the mutation model is specified below).

In the final stage of the iteration (see Fig. 14), the fittest individuals are selected as the new “parents” for the next iteration (the selection procedure is explained in detail below). The fitness of an individual is determined by the so-called fitness function (details are provided below), and is an indication of the performance of the individual’s bidding strategy.

This final step completes one iteration (or “generation”) of the EA. All relevant settings of the evolutionary system are listed in Table V.

³The widely-used genetic algorithms (GAs) are more tailored toward binary-coded search spaces [Holland 1975; Mitchell 1996; Goldberg 1989].

Selection model. Selection is performed using the $(\mu + \lambda)$ -ES selection scheme [Bäck 1996]. In conventional notation, μ is the number of parents and λ is the number of generated offspring ($\mu = \lambda = 25$, see Table V). The μ survivors with the highest fitness are selected (deterministically) from the union of parental and offspring agents. The $(\mu + \lambda)$ -ES selection scheme is an example of an “overlapping generations” model, in which successful agents can survive for multiple generations. In general, EAs use two additional operators: mutation and recombination. These operators are explained in detail below.

Fitness evaluation. The fitness function measures the performance of a bidding strategy. The actual performance naturally depends on the context, i.e. the profiles of the visiting consumers and the bidding strategies used by the opponents (viz. the competing shops). In order to obtain an adequate indication of the performance, the fitness measure is based on both aspects. Firstly, the total amount of profits made when a number of different virtual customers visit the mall is measured. The fitness of a strategy is then equal to the average profits obtained in several trials with different opponent strategies. The fitness of the opponent strategies is determined concurrently.

We now give a more detailed description of the steps used to determine the fitness of the suppliers’ bidding strategies.

- (1) For each of the suppliers combine the offspring and parent population into a single larger population. We now have m populations, one for each supplier.
- (2) Reset all previously made profits.
- (3) Select randomly a single strategy from each population. These bidding strategies are used by the suppliers in the competition. If the competitor is set to *random* (as in Section 4.2), however, the strategies are evaluated against random bidding strategies.
- (4) Let a number of consumers with different profiles visit the shopping mall in a sequential order. We use a fixed set of consumers that are evenly distributed over the profile space (this reduces stochastic variation in measuring the performance of the strategies).
- (5) For each consumer the supplier obtains feedback on the obtained profits. When a consumer visits the mall the following steps determine the profits:
 - (a) Each supplier bids on the consumer using the selected strategy and given the consumer’s profile. The strategy is basically a function which maps the consumer profile to a bid. Below, the details on the strategy representation are described.
 - (b) The mall manager agent (MMA) selects the winners and determines the advertising costs, as described in Section 2.6. Only suppliers who bid higher than zero will participate in the negotiation.
 - (c) The MMA shows the list of selected suppliers to the consumer, who decides how much to buy. The purchase amount is determined by the consumer profile and consumer behavior models described in Sections 3.4 and 3.5.
- (6) The total profits (purchases minus advertising costs) for each strategy are then stored for later reference.




Fig. 15. Piece-wise linear bidding strategy as learned by co-evolution. The bidding strategy determines the bid value for any consumer profile.

- (7) If the profits of a strategy have been determined a pre-set, fixed number of trials (and the strategy has thus been tested against different opponent strategies), this strategy is removed from the population.
- (8) The process is repeated from step 2 until all the populations are empty.
- (9) The fitness for each strategy then equals the average profit obtained in each of the trials.

Bidding strategy. In general terms, a supplier's bidding strategy is a function which returns a bid value given the consumer profile. Within the set-up of the simulation the profile has either one or two dimensions. In case of a single dimension, the strategy is represented using a piece-wise linear function that returns the bid given a value along the consumer-profile axis. An example of such a bidding strategy is given in Fig. 15. For a two-dimensional consumer profile, the strategy is represented by triangular planes. See for an example Fig. 16.

Strategy encoding. The bidding strategies are encoded on the so-called chromosomes of an agent. A chromosome consists of a collection of "genes". Each gene is a real value with a range between 0 and 1. In case of a one-dimensional profile, these genes together form a piece-wise linear function, which maps a consumer profile onto a value to bid. The chromosome then contains (x, y) coordinates for each of the defining points (the number of defining points is a parameter in the simulation), where x is the consumer profile and y the bidding value. The bidding values for the edges of the consumer profile are always specified within the chromosome. The bidding value for a given consumer profile is then calculated by interpolation between two points neighboring of the consumer profile on each side.

For a two-dimensional consumer profile, the strategy is represented by triangular planes. The strategy is constructed using Delaunay triangulation of the (three-dimensional) defining points. The bidding value is then determined by interpolation between the three vertexes of the triangle containing the given consumer profile.

Initialisation. The chromosomes are initialized at the beginning of each EA run by drawing a random number in the unit interval for each gene (from a uniform




Fig. 16. An example of an evolved bidding strategy for a two-dimensional consumer profile.

distribution).

Mutation and recombination model. Mutation operates directly on the “chromosome” of an agent. The offspring’s genes x_i are created by adding a zero-mean Gaussian variable with a standard deviation σ_i to each corresponding gene of the parent [Bäck 1996; van Bragt et al. 2000]. All offspring genes with a value larger than unity (or smaller than zero) are set equal to unity (respectively zero). In our simulations, we use an elegant mutation model with self-adaptive control of the standard deviations σ_i [Bäck 1996, pp. 71-73][van Bragt et al. 2000]. This model allows the evolution of both the genes and the corresponding standard deviations at the same time. More formally, an agent consists of object variables $[x_0, \dots, x_{l-1}]$ and ES-parameters $[\sigma_0, \dots, \sigma_{l-1}]$ in this model.

The mutation operator first updates an agent’s ES-parameters σ_i into σ'_i -values in the following way:

$$\sigma'_i := \sigma_i \exp[\tau' N(0, 1) + \tau N_i(0, 1)], \quad (1)$$

where τ' and τ are the so-called “global” and “individual” learning rates, and $N(0, 1)$ denotes a normally distributed random variable having expectation zero and standard deviation one. The indices i in N_i indicates that the variable is sampled anew for each value of i . We use commonly recommended settings for these parameters.⁴ After the strategy parameters have been modified, the object variables are mutated:

$$x'_i := x_i + \sigma'_i N_i(0, 1). \quad (2)$$

The initial standard deviations $\sigma_i(0)$ are set to a value of 0.1 (see Table V). The

⁴Namely, $\tau' = (\sqrt{2l})^{-1}$ and $\tau = (\sqrt{2\sqrt{l}})^{-1}$ [Bäck 1996, p. 72], where l is the length of the chromosome.

particular value chosen for $\sigma_i(0)$ is not expected to be crucial, because the self-adaptation process (consisting of the parameter updating and the chromosome selection process) rapidly scales the step sizes into the proper range. To prevent complete convergence of the population, we force all standard deviations to remain larger than a small value $\varepsilon_\sigma = 0.025$ [Bäck 1996, pp. 72–73] (see Table V).

Initially we focus on mutation-based models in this paper. In a later stage, recombination can be used to ascertain whether performance is increased.

REFERENCES

- ALKEMADE, F., LA POUTRÉ, H., AND AMMAN, H. 2003. Intermediaries in an electronic trade network. In *Proceedings of the fourth ACM Conference on Electronic Commerce (EC'03)*. June 9–12, San Diego, USA, The ACM Press. to appear.
- BÄCK, T. 1996. *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, New York and Oxford.
- BLATTBERG, R. AND DEIGHTON, J. 1991. Interactive marketing: Exploiting the age of addressability. *Sloan Management Review* 33, 5–14.
- BOHTE, S., GERDING, E., AND LA POUTRÉ, H. 2001. Competitive market-based allocation of consumer attention space. Tech. Rep. SEN-R0131, CWI, Amsterdam.
- BOHTE, S., KOK, J., AND LA POUTRÉ, H. 2002. Spike-prop: error-backpropagation in multi-layer networks of spiking neurons. *Neurocomputing* 48, 1–4 (October), 17–37.
- BOHTE, S. M., GERDING, E. H., AND LA POUTRÉ, H. 2001. Competitive market-based allocation of consumer attention space. In *Proceedings of the 3rd ACM Conference on Electronic Commerce (EC-01)*, M. Wellman, Ed. The ACM Press, 202–206.
- CHENG, J. AND WELLMAN, M. 1998. The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics* 12, 1–24.
- CHIRCU, A. AND R.J., K. 1999. Strategies for internet middlemen in the intermediation/disintermediation/reintermediation cycle. *Electronic Markets* 9, 2, 109–117.
- CLEARWATER, S., Ed. 1995. *Market based Control of Distributed Systems*. World Scientific Press, Singapore.
- DASGUPTA, P. AND MASKIN, E. 2000. Efficient auctions. *The Quarterly Journal of Economics* 115, 341–388.
- GELLMANN, R. 1996. Disintermediation and the internet. *Government Information Quarterly* 13, 1, 1–8.
- GERDING, E., VAN BRAGT, D., AND POUTRÉ, H. L. 2000. Multi-issue negotiation processes by evolutionary simulation: Validation and social extensions. *Computational Economics*. Forthcoming; a preliminary version appeared as Technical Report SEN-R0024, CWI, Amsterdam, 2000.
- GIBNEY, M., JENNINGS, N., VRIEND, N., AND GRIFFITHS, J. 1999. Market-based call routing in telecommunications networks using adaptive pricing and real bidding. In *Proceedings of the 3rd International Workshop on Intelligent Agents for Telecommunication Applications (IATA-99)*, Stockholm, Sweden, August 9–10 1999. 50–65.
- GOLDBERG, D. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading.
- GREENWALD, A. AND KEPHART, J. 1999. Shopbots and pricebots. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*. Vol. 1. 506–511.
- GUTTMAN, R., MOUKAS, A., AND MAES, P. 1998. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review Journal*, 147–159.
- HAECKEL, S. 1998. About the nature and future of interactive marketing. *Journal of Interactive Marketing* 12, 63–71.
- HESKES, T. 1997. Practical confidence and prediction intervals. In *Advances in Neural Information Processing Systems*, M. C. Mozer, M. I. Jordan, and R. Petsche, Eds. Vol. 9. The MIT Press, 176–183.

- HOLLAND, J. 1975. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. The University of Michigan Press/Longman Canada, Ann Arbor.
- HTTP://WWW.GOOGLE.COM. Google search engine.
- HTTP://WWW.OVERTURE.COM. Overture.
- HTTP://WWW.WORLDMALL.COM. World Mall, online shopping portal.
- KEPHART, J., HANSON, J., AND GREENWALD, A. 2000. Dynamic pricing by software agents. *Computer Networks* 32, 6, 731–752.
- KIRMAN, A. AND VRIEND, N. 2001. Evolving market structure: An ACE model of price dispersion and loyalty. *Journal of Dynamics and Control* 25, 459–502.
- KOBSA, A. 2001. Tailoring privacy to users' needs. In *Proceedings of the 8th International Conference on User Modeling, Sonthofen, Germany*. Springer Verlag.
- KRULWICH, B. 1996. The BargainFinder agent: Comparison price shopping on the internet. In *Agents, Bots and Other Internet Beasts*, J. Williams, Ed. SAMS.NET publishing.
- MITCHELL, M. 1996. *An Introduction to Genetic Algorithms*. The MIT Press, Cambridge, MA.
- MONTGOMERY, A. 2001. Applying quantitative marketing techniques to the internet. *Interfaces* 32, 90–108.
- MONTGOMERY, A. AND SRINIVASAN, K. 2002. Learning about customers without asking.
- MOUKAS, A., GUTTMAN, R., AND MAES, P. 1998. Agent-mediated Electronic Commerce: An MIT Media Laboratory Perspective. In *Proceedings of the International Conference on Electronic Commerce, Seoul, Korea*.
- NISSEN, M. 2000. Agent-based supply chain disintermediation versus re-intermediation: Economic and technological perspectives. *International Journal of Intelligent Systems in Accounting, Finance and Management* 9, 237–256.
- NWANA, H. AND NDUMU, D. 1999. A perspective on software agents research. *The Knowledge Engineering Review* 14, 2, 1–18.
- OWEN, M., LEE, L., STEWARD, S., AND THOMAS, D. 1999. Multi-agent trading environment. *BT Technology Journal* 17, 33–43.
- PEPPERS, D., ROGERS, M., AND DORF, B. 1999. Is your company ready for one-to-one marketing? *Harvard Business Review* 77, 151–162.
- SCHAFER, J., KONSTAN, J., AND RIEDI, J. 1999. Recommender systems in e-commerce. In *Proceedings of the first ACM Conference on Electronic Commerce (EC'99)*. ACM Press, 158–166.
- SHAFFER, G. AND ZHANG, Z. 2002. Competitive one-to-one promotions. *Management Science* 48, 9, 1143–1160.
- SUTTON, R. AND BARTO, A. 1998. *Reinforcement learning: An introduction*. MIT-press, Cambridge, MA.
- 'T HOEN, P. J., BOHTE, S. M., GERDING, E., AND LA POUTRÉ, H. 2002. Implementation of a competitive market-based allocation of consumer attention space. In *Lecture Notes in Computer Science: Proceedings of the 4th Workshop on Agent Mediated Electronic Commerce.*, W. Walsh, Ed. Vol. 2531. Springer.
- TESFATSION, L. 2001. Introduction to the special issue on agent-based computational economics. *Journal of Economic Dynamics and Control* 25, 3/4, 281–293.
- VAN BRAGT, D., GERDING, E., AND LA POUTRÉ, H. 2000. Equilibrium selection in alternating-offers bargaining models: The evolutionary computing approach. Tech. Rep. SEN-R0013, CWI, Amsterdam. Presented at the 6th International Conference on Computing in Economics and Finance (CEF'2000).
- VAN BRAGT, D. AND LA POUTRÉ, H. 2001. Generating efficient automata for negotiations - an exploration with evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECKO-2001), San Francisco, July 2001*.
- VAN BRAGT, D., VAN KEMENADE, C., AND LA POUTRÉ, H. 2001. The influence of evolutionary selection schemes on the iterated prisoner's dilemma. *Computational Economics* 17, 253–263.
- VARIAN, H. 1995. Economic mechanism design for computerized agents. In *First USENIX Workshop on Electronic Commerce*. New York, 13–21.
- ACM Transactions on Internet Technology, Vol. V, No. N, Month 20YY.

- VERMEER, B. 2001. Data Quality and Data Alignment in E-Business. Ph.D. thesis, Eindhoven University of Technology.
- VICKREY, W. 1961. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* 16, 8–37.
- VRIEND, N. 1995. Self-organization of markets: An example of a computational approach. *Computational Economics* 8, 205–231.
- WEISS, G., Ed. 1999. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. The MIT Press, Cambridge, MA.

Received Month Year; revised Month Year; accepted Month Year