

An Advanced Bidding Agent for Advertisement Selection on Public Displays

Alex Rogers¹, Esther David², Terry R. Payne¹ and Nicholas R. Jennings¹

¹ Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.
{acr, trp, nrj}@ecs.soton.ac.uk

² Ashkelon College, Ashkelon, Israel.
astrdod@ash-college.ac.il

ABSTRACT

In this paper we present an advanced bidding agent that participates in first-price sealed bid auctions to allocate advertising space on *BluScreen* – an experimental public advertisement system that detects users through the presence of their Bluetooth enabled devices. Our bidding agent is able to build probabilistic models of both the behaviour of users who view the adverts, and the auctions that it participates within. It then uses these models to maximise the exposure that its adverts receive. We evaluate the effectiveness of this bidding agent through simulation against a range of alternative selection mechanisms including a simple bidding strategy, random allocation, and a centralised optimal allocation with perfect foresight. Our bidding agent significantly outperforms both the simple bidding strategy and the random allocation, and in a mixed population of agents it is able to expose its adverts to 25% more users than the simple bidding strategy. Moreover, its performance is within 7.5% of that of the centralised optimal allocation despite the highly uncertain environment in which it must operate.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Intelligent agents

General Terms

Algorithms, Design, Theory

Keywords

bidding agent, auction, public display

1. INTRODUCTION

Electronic displays are increasingly being used within public environments, such as airports, city centres and retail stores, in order to advertise commercial products, or to entertain and inform passers-by. Recently, researchers have begun to investigate how the content of such displays may be varied dynamically over time in order to increase its variety, relevance and exposure [9]. Particular research attention has focused on the need to take into account the dynamic nature of the display's audience, and to this end, a number

of interactive public displays have been proposed. These displays have typically addressed the needs of a closed set of known users with pre-defined interests and requirements, and have facilitated communication with these users through the active use of handheld devices such as PDAs or phones [3, 7]. As such, these systems assume prior knowledge about the target audience, and require either that a single user has exclusive access to the display, or that users carry specific tracking devices so that their presence can be identified [6, 11]. However, these approaches fail to work in public spaces, where no prior knowledge regarding the users who may view the display exists, and where such displays need to react to the presence of several users simultaneously.

By contrast, Payne et al. have developed an intelligent public display system, named *BluScreen*, that detects and tracks users through the Bluetooth enabled devices that they carry with them everyday [8]. Within this system, a decentralised multi-agent auction mechanism is used to efficiently allocate advertising time on each public display. Each advert is represented by an individual advertising agent that maintains a history of users who have already been exposed to the advert. This agent then seeks to acquire advertising cycles (during which it can display its advert on the public displays) by submitting bids to a marketplace agent who implements a sealed bid auction. The value of these bids is based upon the number of users who are currently present in front of the screen, the history of these users, and an externally derived estimate of the value of exposing an advert to a user.

In this paper, we present an advanced bidding agent that significantly extends the sophistication of this approach. In particular, we consider the more general setting in which it is impossible to determine an a priori valuation for exposing an advert to a user. This is likely to be the case for *BluScreen* installations within private organisations where the items being advertised are forthcoming events or news items of interest to employees and visitors, and thus have no direct monetary value (indeed in this case bidding is likely to be conducted in some virtual currency). In addition, it is also likely to be the case within new commercial installations where limited market experience makes estimating a valuation impossible. In both cases, it is more appropriate to assume that an advertising agents will be assigned a total advertising budget, and that it will have a limited period of time in which to spend this budget (particularly so where the adverts are for forthcoming events). The advertising agent is then simply tasked with using this budget to maximum effect (i.e. to achieve the maximum possible advert exposure within this time period).

Now, in order to achieve this goal, the advertising agent must be capable of modelling the behaviour of the users in order to predict the number who will be present in any future advertising cycle. In addition, it must also understand the auction environment in which

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.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.
Copyright 2007 IFAAMAS.

it competes, in order that it may make best use of its limited budget. Thus, in developing an advanced bidding agent that achieves this, we advance the state of the art in four key ways:

1. We enable the advertising agents to model the arrival and departure of users as independent Poisson processes, and to make maximum likelihood estimates of the rates of these processes based on their observations. We show how these agents can then calculate the expected number of users who will be present during any future advertising cycle.
2. Using a decision theoretic approach we enable the advertising agents to model the probability of winning any given auction when a specific amount is bid. The cumulative form of the gamma distribution is used to represent this probability, and its parameters are fitted using observations of both the closing price of previous auctions, and the bids that that advertising agent itself submits.
3. We show that our explicit assumption that the advertising agent derives no additional benefit by showing an advert to a single user more than once, causes the expected utility of each future advertising cycle to be dependent on the expected outcome of all the auctions that precede it. We thus present a stochastic optimisation algorithm based upon simulated annealing that enables the advertising agent to calculate the optimal sequence of bids that maximises its expected utility.
4. Finally, we demonstrate that this advanced bidding strategy outperforms a simple strategy with none of these features (within an heterogenous population the advertising agents who use the advanced bidding strategy are able to expose their adverts to 25% more users than those using the simple bidding strategy), and we show that it performs within 7.5% of that of a centralised optimiser with perfect knowledge of the number of users who will arrival and depart in all future advertising cycles.

The remainder of this paper is organised as follows: Section 2 discusses related work where agents and auction-based marketplaces are used to allocated advertising space. Section 3 describes the prototype BluScreen system that motivates our work. In section 4 we present a detailed description of the auction allocation mechanism, and in section 5 we describe our advanced bidding strategy for the advertising agents. In section 6 we present an empirical validation of our approach, and finally, we conclude in section 7.

2. RELATED WORK

The commercial attractiveness of targeted advertising has been amply demonstrated on the internet, where recommendation systems and contextual banner adverts are the norm [1]. These systems typically select content based upon prior knowledge of the individual viewing the material, and such systems work well on personal devices where the owner's preferences and interests can be gathered and cached locally, or within interactive environments which utilise some form of credential to identify the user (e.g. e-commerce sites such as Amazon.com).

Attempts to apply these approaches within the real world have been much more limited. Gerding et al. present a simulated system (CASy) whereby a Vickrey auction mechanism is used to sell advertising space within a modelled electronic shopping mall [2]. The auction is used to rank a set of possible advertisements provided by different retail outlets, and the top ranking advertisements are selected for presentation on public displays. Feedback is provided through subsequent sales information, allowing the model to build up a profile of a user's preferences. However, unlike the BluScreen



Figure 1: A deployed BluScreen prototype.

system that we consider here, it is not suitable for advertising to many individuals simultaneously, as it requires explicit interaction with a single user to acquire the user's preferences.

By contrast, McCarthy et al. have presented a prototype implementation of a system (GroupCast) that attempts to respond to a group of individuals by assuming a priori profiles of several members of the audience [7]. User identification is based on infrared badges and embedded sensors within an office environment. When several users pass by the display, a centralised system compares the user's profiles to identify common areas of interest, and content that matches this common interest is shown.

Thus, whilst CASy is a simulated system that allows advertisers to compete for the attention of single user, GroupCast is a prototype system that detects the presence of groups of users and selects content to match their profiles. Despite their similarities, neither system addresses the settings that interests us here: how to allocate advertising space between competing advertisers who face an audience of multiple individuals about whom there is no a priori profile information. Thus, in the next section we describe the prototype BluScreen system that motivates our work.

3. THE BLUSCREEN PROTOTYPE

BluScreen is based on the notion of a scalable, extendable, advertising framework whereby adverts can be efficiently displayed to as many relevant users as possible, within a knowledge-poor environment. To achieve these goals, several requirements have been identified:

1. Adverts should be presented to as diverse an audience as possible, whilst minimising the number of times the advert is presented to any single user.
2. Users should be identified by existing ubiquitous, consumer devices, so that future deployments within public arenas will not require uptake of new hardware.
3. The number of displays should be scalable, such that adverts appear on different displays at different times.
4. Knowledge about observed behaviour and composition of the audience should be exploited to facilitate inference of user interests which can be exploited by the system.

To date, a prototype systems that addresses the first two goals has been demonstrated [8]. This system uses a 23 inch flat-screen display deployed within an office environment to advertise events and news items. Rather than requiring the deployment of specialised hardware, such as active badges (see [11] for details), BluScreen detects the presence of users in the vicinity of each display through the Bluetooth-enabled devices that they carry with them everyday¹.

¹Devices must be in discovery mode to detectable.

Device Type	Unique Samples	Devices
Occasional	< 10	135
Frequent	10 – 1000	70
Persistent	> 1000	6

Table 1: Number of Bluetooth devices observed at different frequencies over a six month sample period.

This approach is attractive since the Bluetooth wireless protocol is characterised by its relative maturity, market penetration, and emphasis on short-range communication. Table 1 summarises the number of devices detected by this prototype installation over a six month period. Of the 212 Bluetooth devices detected, approximately 70 were detected regularly, showing that Bluetooth is a suitable proxy for detecting individuals in front of the screen.

In order to achieve a scalable and extendable solution a multi-agent systems design philosophy is adopted whereby a number of different agents types interact (see figure 2). The interactions of these agents are implemented through a web services protocol², and they constitute a decentralised marketplace that allocates advertising space in an efficient and timely manner. In more detail, the responsibilities of each agent types are:

Bluetooth Device Detection Agent: This agent monitors the environment in the vicinity of a BluScreen display and determines the number and identity of any Bluetooth devices that are close by. It keeps historical records of the arrival and departure of Bluetooth devices, and makes this information available to advertising agents as requested.

Marketplace Agent: This agent facilitates the sale of advertising space to the advertising agents. A single marketplace agent represents each BluScreen display, and access to this screen is divided into discrete advertising cycles of fixed duration. Before the start of each advertising cycle, the marketplace agent holds a sealed-bid auction (see section 4 for more details). The winner of this auction is allocated access to the display during the next cycle.

Advertising Agent: This agent represents a single advert and is responsible for submitting bids to the marketplace agent in order that it may be allocated advertising cycles, and thus, display its advert to users. It interacts with the device detection agent in order to collect information regarding the number and identity of users who are currently in front of the display. On the basis of this information, its past experiences, and its bidding strategy, it calculates the value of the bid that it should submit to the marketplace agent.

Thus, having described the prototype BluScreen system, we next go on to describe the details of the auction mechanism that we consider in this work, and then the advanced bidding agent that operates bids within this auction.

4. THE AUCTION MECHANISM

As described above, BluScreen is designed to efficiently allocate advertising cycles in a distributed and timely manner. Thus, one-shot sealed bid auctions are used for the market mechanism of the marketplace agent. In previous work, each advertising agent was assumed to have an externally derived estimate of the value of exposing an advert to a user. Under this assumption, a second-price sealed-bid auction was shown to be effective, since advertis-

²This is implemented on a distributed Mac OS X based system using the Bonjour networking protocol for service discovery.

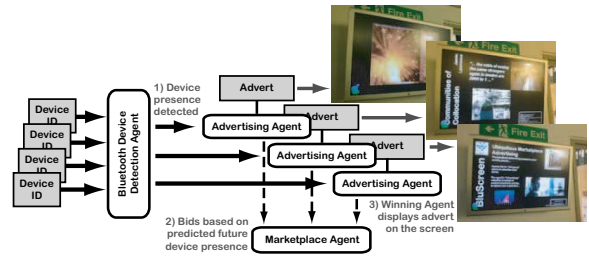


Figure 2: The BluScreen agent architecture for a single display.

ing agents have a simple strategy of truthfully bidding their valuation in each auction [8].

However, as described earlier, in this paper we consider the more general setting in which it is impossible to determine an a priori valuation for exposing an advert to a single user. This may be because the BluScreen installation is within a private organisation where what is being advertised (e.g. news items or forthcoming events) has no monetary value, or it may be a new commercial installation where limited market experience makes estimating such a valuation impossible. In the absence of such a valuation, the attractive economic properties of the second-price auction can not be achieved in practise, and thus, in our work there is no need to limit our attention to the second-price auction. Indeed, since these auctions are actually extremely rare within real world settings [10], in this work we consider the more widely adopted first-price auction since this increases the applicability of our results.

Thus, in more detail, we consider an instance of a BluScreen installation with a single display screen that is managed by a single marketplace agent³. We consider that access to the display screen is divided into discrete advertising cycles, each of length t_c , and a first-price sealed bid auction is held immediately prior to the start of each advertising cycle. The marketplace agent announces the start and deadline of the auction, and collects sealed bids from each advertising agent. At the closing time of the auction the marketplace agent announces to all participants and observers the amount of the winning bid, and informs the winning advertising agent that it was successful (the identity of the winning advertising agent is not announced to all observers). In the case that no bids are placed within any auction, a default advert is displayed.

Having described the market mechanism that the marketplace agent implements, we now go on to describe and evaluate an advanced bidding strategy for the advertising agents to adopt.

5. ADVANCED BIDDING STRATEGY

As described above, we consider the case that the advertising agents do not have an externally derived estimate of the value of exposing the advert to a single user. Rather, they have a constrained budget, B , and a limited period of interest during which they wish to display their advert. Their goal is then to find the appropriate amount to bid within each auction in this period, in order to maximise the exposure of their advert.

In attempting to achieve this goal the advertising agent is faced with a high level of uncertainty about future events. It will be uncertain of the number of users who will be present during any advertising cycle since even if the number of users currently present

³This assumption of having a single BluScreen instance is made to simplify our task of validating the correctness and the efficiency of the proposed mechanism and strategy, and generalising these results to the case of multiple screens is the aim of our future work.

is known, some may leave before the advert commences, and others may arrive. Moreover, the amount that must be bid to ensure that an auction is won is uncertain since it depends on the number and behaviour of the competing advertising agents.

Thus, we enable the agent to use its observations of the arrival and departure of users to build a probabilistic model, based upon independent Poisson processes, that describes the number of users who are likely to be exposed to any advert. In addition, we enable the agent to observe the outcome of previous advertising cycle auctions, and use the observations of the closing price, and the success or otherwise of the bids that it itself submitted, to build a probabilistic model of the bid required to win the auction. The agent then uses these two models to calculate its expected utility in each advertising cycle, and in turn, determine the optimal sequence of bids that maximises this utility given its constrained budget. Having calculated this sequence of bids, then the first bid in the sequence is actually used in the auction for the next advertising cycle. However, at the close of this cycle, the process is repeated with a new optimal sequence of bids being calculated in order take to account of what actually happened in the preceding auction (i.e. whether the bid was successful or not, and how many users arrived or departed).

Thus, in the next three subsections we describe these two probabilistic models, and their application within the bidding strategy of the advertising agent.

5.1 Predicting the Number of Users

In order to predict the number of users that will be present in any future advertising cycle, it is necessary to propose a probabilistic model for the behaviour of the users. Thus, our advanced bidding strategy assumes that their arrival and departures are determined by two independent Poisson processes⁴ with arrival rate, λ_a , and departure rate, λ_d . This represents a simple model that is commonly applied within queuing theory⁵ [5], yet is one that we believe well describes the case where BluScreen displays are placed in communal areas where people meet and congregate. Given the history of users' arrivals and departures obtained from the device detection agent, the advertising agent makes a maximum likelihood estimation of the values of λ_a and λ_d .

In more detail, if the advertising agent has observed n users arriving within a time period t , then the maximum likelihood estimation for the arrival rate λ_a is simply given by:

$$\lambda_a = \frac{n}{t} \quad (1)$$

Likewise, if an agent observes n users each with a duration of stay of t_1, t_2, \dots, t_n time periods, then the maximum likelihood estimation for the departure rate λ_d is given by:

$$\frac{1}{\lambda_d} = \frac{1}{n} \sum_{i=1}^n t_i \quad (2)$$

⁴Given a Poisson distribution with rate parameter λ , the number of events, n , within an interval of time t is given by:

$$P(n) = \frac{e^{-\lambda t} (\lambda t)^n}{n!}$$

In addition, the probability of having to wait a period of time, t , before the next event is determined by:

$$P(t) = \lambda e^{-\lambda t}$$

⁵Note however that in queuing theory it is typically the arrival rate and service times of customers that are modelled as Poisson processes. Our users are not actually modelled as a queue since the duration of their stay is independent of that of the other users.




Figure 3: Example showing how to predict the number of users who see an advert shown in an advertising cycle of length t_c , commencing at time t in the future.

In environments where these rates are subject to change, the agent can use a limited time window over which observations are used to estimate these rates. Alternatively, in situations where cyclic changes in these rates are likely to occur (i.e. changing arrival and departure rates at different times of the day, as may be seen in areas where commuters pass through), the agent can estimate separate values over each hour long period.

Having estimated the arrival and departure rate of users, and knowing the number of users who are present at the current time, the advertising agent is then able to predict the number of users who are likely to be present in any future advertising cycle⁶. Thus, we consider the problem of predicting this number for an advertising cycle of duration t_c that starts at a time t in the future, given that n users are currently present (see figure 3). This number will be composed of three factors: (i) the fraction of the n users that are initially present who do not leave in the interval, $0 \leq \tau < t$, before the advertising cycle commences, (ii) users that actually arrive in the interval, $0 \leq \tau < t$, and are still present when the advertising cycle actually commences, and finally, (iii) users that arrive during the course of the advertising cycle, $t \leq \tau < t + t_c$.

Now, considering case (i) above, the probability of one of the n users still being present when the advertising cycle starts is given by $\int_t^\infty \lambda_d e^{-\lambda_d \tau} d\tau = e^{-\lambda_d t}$. Thus we expect $n e^{-\lambda_d t}$ of these users to be present. In case (ii), we expect $\lambda_a t$ new users to arrive before the advertising cycle commences, and the probability that any of these will still be there when it actually does so is given by $\frac{1}{t} \int_0^t e^{-\lambda_d (t-\tau)} d\tau = \frac{1}{\lambda_d t} [1 - e^{-\lambda_d t}]$. Thus we expect $\frac{\lambda_a}{\lambda_d} [1 - e^{-\lambda_d t}]$ of these users to be present. Finally, in case (iii) we expect $\lambda_a t_c$ users to arrive during the course of the advertising cycle. Thus, the combination of these three factors gives an expression for the expected number of users who will be present within an advertising cycle of length t_c , that commencing at time t in the future, given that there are n users currently present:

$$\mathcal{N}_{n,t} = n e^{-\lambda_d t} + \frac{\lambda_a}{\lambda_d} [1 - e^{-\lambda_d t}] + \lambda_a t_c \quad (3)$$

Note that as t increases the results become less dependent upon the initial number of users, n . The mean number of users present at any time is simply λ_a / λ_d , and the mean number of users exposed to an advert in any advertising cycle is given by $\lambda_a (t_c + \frac{1}{\lambda_d})$.

5.2 Predicting the Probability of Winning

In addition to estimating the number of users who will be present in any advertising cycle, an effective bidding agent must also be able to predict the probability of it winning an auction given that it submits any specified bid. This is a common problem within bidding agents, and approaches can generally be classified as game theoretic or decision theoretic. Since our advertising agents are unaware of the number or identity of the competing advertising

⁶Note that we do not require a user to be present for the entire advertising cycle in order to be counted as present.

agents, the game theoretic approach is precluded. Thus, we take a decision theoretic approach similar to that adopted within continuous double auctions where bidding agents estimate the market price of goods by observing transaction prices [4].

Thus, our advertising agents uses a parameterised function to describe the probability of winning the auction given any submitted bid, $P(b)$. This function must have support $[0, \infty)$ since bids must be positive. In addition, we expect it to exhibit by an 's' shaped curve whereby the probability of winning an auction is small when the submitted bid is very low, the probability is close to one when the bid is very high, and there is a transition point that characterises the change from a losing to a winning bid. To this end, we use the cumulative form of the gamma distribution for this function:

$$P(b) = \frac{\gamma(k, b/\theta)}{\Gamma(k)} \quad (4)$$

where $\Gamma(k)$ is the standard gamma function, and $\gamma(k, b/\theta)$ is the incomplete gamma function. This function has the necessary properties described above, and has two parameters, k and θ . The transition point where $P(b) = 0.5$ is given by $k\theta$ and the sharpness of the transition is described by $k\theta^2$. In figure 4 we show examples of this function for three different values of k and θ .

The advertising agent chooses the most appropriate values of k and θ by fitting the probability function to observations of previous auctions. An observation is a pair $\{b_i, o_i\}$ consisting of the bid, b_i , and an auction outcome, o_i . Each auction generates at least one pair in which b_i is equal to the closing price of the auction, and $o_i = 1$. In addition, another pair is generated for each unsuccessful bid submitted by the advertising agent itself, and in this case $o_i = 0$. Thus, having collected N such pairs⁷, the agent finds the values of k and θ by evaluating:

$$\arg \min_{k, \theta} \sum_{i=1}^N \left[o_i - \frac{\gamma(k, b_i/\theta)}{\Gamma(k)} \right]^2 \quad (5)$$

This expression can not be evaluated analytically, but can be simply found using a numerical gradient descent method whereby the values of k and θ are initially estimated using their relationship to the transition point described above. The gradient of this expression is then numerically evaluated at these points, and new estimates of k and θ calculated by making a fixed size move in the direction of maximum gradient. This process is repeated until k and θ have converged to an appropriate degree of accuracy.

5.3 Expected Utility of an Advertising Cycle

The goal of the advertising agent is to gain the maximum exposure for its advert given its constrained budget. We define the utility of any advertising cycle as the expected number of users who will see the advert for the first time during that cycle, and hence, we explicitly assume that no additional utility is derived by showing the advert to any user more than once⁸. Thus, we can use the results of the previous two sections to calculate the expected utility of each advertising cycle remaining within the advertising agent's period of

⁷In the case that no unsuccessful bids have been observed, there is no evidence of where the transition point between successful and unsuccessful bids is likely to occur. Thus, in this case, an additional pair with value $\{\alpha \min(b_1 \dots b_n), 0\}$ is automatically created. Here $\alpha \in [0, 1]$ determines how far below the lowest successful bid the advertising agent believes the transition point to be. We have typically used $\alpha = 0.5$ within our experiments.

⁸As noted before, we assume that a user has seen the advert if they are present during any part of the advertising cycle, and we do not differentiate between users who see the entire advert, or users who see a fraction of it.

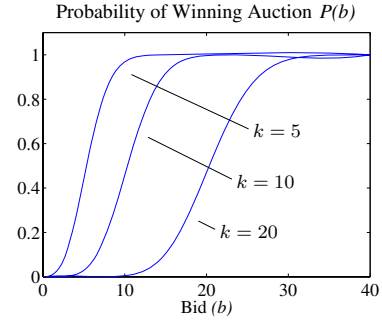


Figure 4: Cumulative gamma distribution representing the probability of winning an auction ($\theta = 1$ and $k = 5, 10 \& 20$).

interest. In the first advertising cycle this is simply determined by the probability of the advertising agent winning the auction, given that it submits a bid b_1 , and the number of users who are currently in front of the BluScreen display, but have not seen the advert before, is n . Thus, the expected utility of this advertising cycle is simply described by:

$$u_1 = P(b_1)\mathcal{N}_{n,0} \quad (6)$$

Now, in the second advertising cycle, the expected utility will clearly depend on the outcome of the auction for the first. If the first auction was indeed won by the agent, then there will be no users who have yet to see the advert present at the start of the second advertising cycle. Thus, in this case, the expected number of new users who will see the advert in the second advertising cycle is described by $\mathcal{N}_{0,0}$ (i.e. only newly arriving users will contribute any utility). By contrast, if the first auction was not won by the agent, then the expected number of users who have yet to see the advert is given by \mathcal{N}_{n,t_c} where t_c is the length of the preceding advertising cycle (i.e. exactly the case described in section 5.1 where there are n users initially present and the advertising cycle starts at a time t_c in the future). Thus, the expected utility of the second advertising cycle is given by:

$$u_2 = P(b_2) [P(b_1)\mathcal{N}_{0,0} + (1 - P(b_1))\mathcal{N}_{n,t_c}] \quad (7)$$

We can generalise this result by noting that the number of users expected to be present within any future advertising cycle will depend on the number of cycles since an auction was last won (since at this point the number of users who are present but have not seen the advert must be equal to zero). Thus, we must sum over all possible ways in which this can occur, and weight each by its probability. Hence, the general case for any advertising cycle is described by the rather complex expression:

$$u_i = P(b_i) \left[\sum_{j=1}^{i-1} \mathcal{N}_{0,(i-j-1)t_c} P(b_j) \prod_{m=j+1}^{i-1} (1 - P(b_m)) + \mathcal{N}_{n,(i-1)t_c} \prod_{m=1}^{i-1} (1 - P(b_m)) \right] \quad (8)$$

Thus, given this expression, the goal of the advertising agent is to calculate the sequence of bids over the c remaining auctions, such that the total expected utility is maximised, whilst ensuring that the remaining budget, B , is not exceeded:

$$\arg \max_{b_1 \dots b_c} \sum_{i=1}^c u_i \quad \text{such that} \quad \sum_{i=1}^c b_i = B \quad (9)$$




Figure 5: Total expected utility of the advertising agent over a continuous range of values of b_1 for a number of discrete values of budget, B , when there are just two auction cycles.

Having calculated this sequence, a bid of b_1 is submitted in the next auction. Once the outcome of this auction is known, the process repeats with a new optimal sequence of bids being calculated for the remaining advertising cycles of the agent's period of interest.

5.4 Optimal Sequence of Bids

Solving for the optimal sequence of bids expressed in equation 9 can not be performed analytically. Instead we develop a numerical routine to perform this maximisation. However, it is informative to initially consider the simple case of just two auctions.

5.4.1 Two Auction Example

In this case the expected utility of the advertising agent is simply given by $u_1 + u_2$ (as described in equations 6 and 7), and the bidding sequence is solely dependent on b_1 (since $b_2 = B - b_1$). Thus, we can plot the total expected utility against b_1 and graphically determine the optimal value of b_1 (and thus also b_2).

To this end, figure 5 shows an example calculated using parameter values $\lambda_a = 1/120$, $\lambda_d = 1/480$ and $t_c = 120$. In this case, we assume that $k = 10$ and $\theta = 1$, and thus, given that $k\theta$ describes the midpoint of the cumulative gamma distribution, a bid of 10 represents a 50% chance of winning any auction (i.e. $P(10) = 0.5$). In addition, we assume that $n = \lambda_a/\lambda_d = 4$, and thus the initial number of users present is equal to the mean number that we expect to find present at any time. The plot indicates that when the budget is small, then the maximum utility is achieved at the extreme values of b_1 . This corresponds to bidding in just one of the two auctions (i.e. $b_1 = 0$ and $b_2 = B$ or $b_1 = B$ and $b_2 = 0$). However, as the budget increases, the plot passes through a transition whereby the maximum utility occurs at the midpoint of the x-axis, corresponding to bidding equally in both auctions (i.e. $b_1 = b_2 = B/2$). This is simply understood by the fact that continuing to allocate the budget to a single auction results in diminishing returns as the probability of actually winning this auction approaches one.

In this case, the plot is completely symmetrical since the number of users present at the start is equal to its expected value (i.e. $n = \lambda_a/\lambda_d$). If however, $n < \lambda_a/\lambda_d$ the plot is skewed such that when the budget is small, it should be allocated to the second auction (since more users are expected to arrive before this advertising cycle commences). Conversely, when $n > \lambda_a/\lambda_d$ the entire budget should be allocated to the first auction (since the users who are currently present are likely to depart in the near future). However, in both cases, a transition occurs whereby given sufficient budget it is preferable to allocate the budget evenly between both auctions⁹.

⁹In fact, one auction is still slightly preferred, but the difference in

```

temp ← 1
rate ← 0.995
bold ← initial random allocation
Uold ← Evaluate(bold)
WHILE temp > 0.0001
  i, j ← random integer index within b
  t ← random real number between 0 and bi
  bnew ← bold
  binew ← biold - t
  bjnew ← bjold + t
  Unew ← Evaluate(bnew)
  IF rand < exp((Unew - Uold)/temp) THEN
    bold ← bnew
    Uold ← Unew
  ENDIF
  temp ← temp × rate
ENDWHILE

```

Figure 6: Stochastic optimisation algorithm to calculate the optimal sequence of bids in the general case of multiple auctions.

5.4.2 General Case

In general, the behaviour seen in the previous example characterises the optimal bidding behaviour of the advertising agent. If there is sufficient budget, bidding equally in all auctions results in the maximum expected utility. However, typically this is not possible and thus utility is maximised by concentrating what budget is available into a subset of the available auction. The choice of this subset is determined by a number of factors. If there are very few users currently present, it is optimal to allocate the budget to later auctions in the expectation that more users will arrive. Conversely, if there are many users present, a significant proportion of the budget should be allocated to the first auction to ensure that it is indeed won, and these users see the advert. Finally, since no utility is derived by showing the advert to a single user more than once, the budget should be allocated such that there are intervals between showings of the advert, in order that new users may arrive.

Now, due to the complex form of the expression for the expected utility of the agent (shown in equation 8) it is not possible to analytically calculate the optimal sequence of bids. However, the inverse problem (that of calculating the expected utility for any given sequence of bids) is easy. Thus, we can use a stochastic optimisation routine based on simulated annealing to solve the maximisation problem. This algorithm starts by assuming some initial random allocation of bids (normalised such that the total of all the bids is equal to the budget B). It then makes small adjustments to this allocation by randomly transferring the budget from one auction to another. If this transfer results in an increase in expected utility, then it is accepted. If it results in a decrease in expected utility, it might still be accepted, but with a probability that is determined by a temperature parameter. This temperature parameter is annealed such that the probability of accepting such transfers decreases over time. In figure 6 we present this algorithm in pseudo-code.

6. EVALUATION

In order to evaluate the effectiveness of the advanced bidding strategy developed within this paper we compare its performance to three alternative mechanisms. One of these mechanisms represents a simple alternative bidding strategy for the advertising agents, whilst the other two are centralised allocation mechanisms that represent

expected utility between this and an even allocation is negligible.

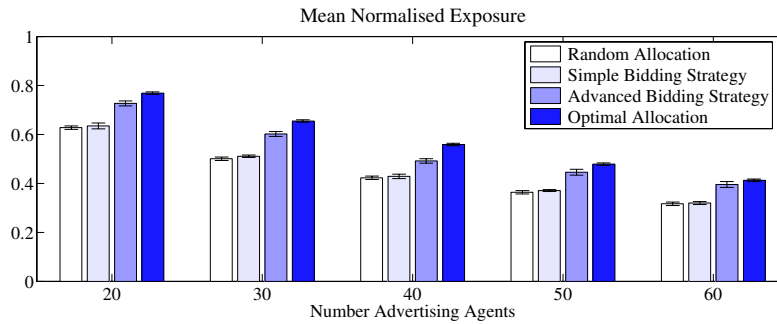


Figure 7: Comparison of four different allocation mechanisms for allocating advertising cycles to advertising agents. Results are averaged over 50 simulation runs and error bars indicate the standard error in the mean.

the upper and lower bounds to the overall performance of the system. In more detail, the four mechanisms that we compare are:

Random Allocation: Rather than implementing the auction mechanism, the advertising cycle is randomly allocated to one of the advertising agents.

Simple Bidding Strategy: We implement the full auction mechanism but with a population of advertising agents that employ a simple bidding strategy. These advertising agents do not attempt to model the users or the auction environment in which they bid, but rather, they simply evenly allocate their remaining budget over the remaining advertising cycles.

Advanced Bidding Strategy: We implement the full auction mechanism with a population of advertising agents using the probabilistic models and the bidding strategy described here.

Optimal Allocation: Rather than implementing the auction mechanism, the advertising cycle is allocated to the advertising agent that will derive the maximum utility from it, given perfect knowledge of the number of users who will arrive and depart in all future advertising cycles.

Using these four alternative allocation mechanisms, we ran repeated simulations of two hours of operation of the entire BluScreen environment for a default set of parameters whereby the arrival and departure rate of the users are given by $\lambda_a = 1/120s$ and $\lambda_d = 1/480s$, and the length of an advertising cycle is 120s. Each advertising agent is assigned an advert with a period of interest drawn from a Poisson distribution with a mean of 8 advertising cycles, and these agents are initially allocated a budget equal to 10 times their period of interest. For each simulation run, we measure the mean normalised exposure of each advert. That is, the fraction of users who were detected by the BluScreen display during the period of interest of the advertising agent who were actually exposed to the agent’s advert. Thus a mean normalised exposure of 1 indicates that the agent managed to expose its advert to all of the users who were present during its period of interest (and a mean normalised exposure of 0 means that no users were exposed to the advert).

Figure 7 shows the results of this experiments. We first observe the general result that as the number of advertising agents increases, and thus the competition between them increases, then the mean normalised exposure of all allocation mechanisms decreases. We then observe that in all cases, there is no statistically significant improvement in using the simple bidding strategy compared to random allocation ($p > 0.25$ in Student’s t-test). Since this simple bidding strategy does not take account of the number of users present, and in general, simply increases its bid price in each auction until it

does in fact win one, this is not unexpected. However, in all cases the advanced bidding strategy does indeed significantly outperform the simple bidding agent ($p < 0.0005$ in Student’s t-test), and its performance is within 7.5% of that of the optimal allocation that has perfect knowledge of the number of users who will arrive and depart in all future advertising cycles.

In addition, we present results of experiments performed over a range of parameter values, and also with a mixed population of advertising agents using both the advanced and simple bidding strategies. This is an important scenario since advertisers may wish to supply their own bidding agents, and thus, a homogeneous population is not guaranteed. In each case, keeping all other parameters fixed, we varied one parameter, and these results are shown in figure 8. In general, we see the similar trends as before. Increasing the departure rate causes a decrease in the mean normalised exposure since advertising agents have less opportunities to expose users to their adverts. Increasing the period of interest of each agent decreases the mean normalised exposure, since more advertising agents are now competing for the same users. Finally, increasing the arrival rate of the users causes the results of the simple and advanced bidding strategies to approach one another, since the variance in the number of users who are present during any advertising cycle decreases, and thus, modelling their behaviour provides less gain. However, in all cases, the advanced bidding strategy significantly outperforms the simple one ($p < 0.0005$ in Student’s t-test). On average, we observe that advertising agents who use the advanced bidding strategy are able to expose their adverts to 25% more users than those using the simple bidding strategy.

Finally, we show that a rational advertising agent, who has a choice of bidding strategy, would always opt to use the advanced bidding strategy over the simple bidding strategy, regardless of the composition of the population that it finds itself in. Figure 9 shows the average normalised exposure of the advertising agents when the population is composed of different fractions of the two bidding strategies. In each case, the advanced bidding strategy shows a significant gain in performance compared to the simple bidding strategy ($p < 0.0005$ in Student’s t-test), and thus, gains improved exposure over all population compositions.

7. CONCLUSIONS

In this paper, we presented an advanced bidding strategy for use by advertising agents within the BluScreen advertising system. This bidding strategy enabled advertising agents to model and predict the arrival and departure of users, and also to model their success within a first-price sealed bid auction by observing both the bids that they themselves submitted and the winning bid. The ex-




Figure 8: Comparison of an evenly mixed population of advertising agents using simple and advanced bidding strategies over a range of parameter settings. Results are averaged over 50 simulation runs and error bars indicate the standard error in the mean.

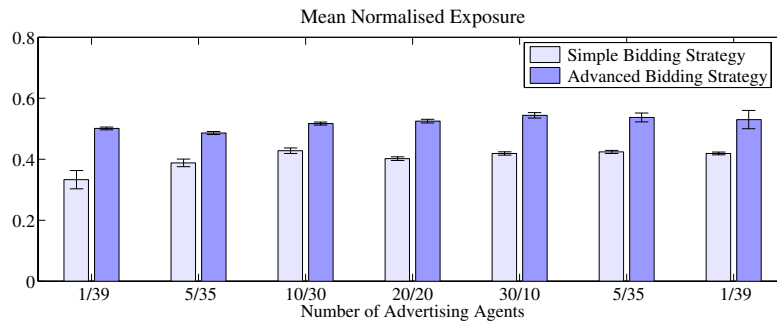


Figure 9: Comparison of an unevenly mixed population of advertising agents using simple and advanced bidding strategies. Results are averaged over 50 simulation runs and error bars indicate the standard error in the mean.

pected utility, measured as the number of users who the advertising agent exposes its advert to, was shown to depend on these factors, and resulted in a complex expression where the expected utility of each auction depended on the success or otherwise of earlier auctions. We presented an algorithm based upon simulated annealing to solve for the optimal bidding strategy, and in simulation, this bidding strategy was shown to significantly outperform a simple bidding strategy that had none of these features. Its performance closely approached that of a central optimal allocation, with perfect knowledge of the arrival and departure of users, despite the uncertain environment in which the strategy must operate.

Our future work in this area consists of extending this bidding strategy to richer environments where there are multiple interrelated display screens, where maintaining profiles of users allows a richer matching of user to advert, and where alternative auction mechanisms are applied (we are particularly interested in introducing a ‘pay per user’ auction setting similar to the ‘pay per click’ auctions employed by internet search websites). This work will continue to be done in conjunction with the deployment of more BluScreen prototypes in order to gain further real world experience.

8. ACKNOWLEDGEMENTS

The authors would like to thank Heather Packer and Matthew Sharifi (supported by the ALADDIN project - www.aladdinproject.org) for their help in developing the deployed prototype.

9. REFERENCES

- [1] A. Amiri and S. Menon. Efficient scheduling of internet banner

- advertisements. *ACM Transactions on Internet Technology*, 3(4):334–346, 2003.
- [2] S. M. Bohte, E. Gerding, and H. L. Poutre. Market-based recommendation: Agents that compete for consumer attention. *ACM Transactions on Internet Technology*, 4(4):420–448, 2004.
- [3] K. Cheverst, A. Dix, D. Fitton, C. Kray, M. Rouncefield, C. Sas, G. Saslis-Lagoudakis, and J. G. Sheridan. Exploring bluetooth based mobile phone interaction with the hermes photo display. In *Proc. of the 7th Int. Conf. on Human Computer Interaction with Mobile Devices & Services*, pages 47–54, Salzburg, Austria, 2005.
- [4] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, (22):1–29, 1998.
- [5] D. Gross and C. M. Harris. *Fundamentals of Queueing Theory*. Wiley, 1998.
- [6] J. Hightower and G. Borriella. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, 2001.
- [7] J. F. McCarthy, T. J. Costa, and E. S. Liongosari. Unicast, outcast & groupcast: Three steps toward ubiquitous, peripheral displays. In *Proc. of the 3rd Int. Conf. on Ubiquitous Computing*, pages 332–345, Atlanta, USA, 2001.
- [8] T. R. Payne, E. David, M. Sharifi, and N. R. Jennings. Auction mechanisms for efficient advertisement selection on public display. In *Proc. of the 17th European Conf. on Artificial Intelligence*, pages 285–289, Trentino, Italy, 2006.
- [9] A. Ranganathan and R. H. Campbell. Advertising in a pervasive computing environment. In *Proc. of the 2nd Int. Workshop on Mobile Commerce*, pages 10–14, Atlanta, Georgia, USA, 2002.
- [10] M. Rothkopf, T. Teisberg, and E. Kahn. Why are vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.
- [11] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, 1992.