Auction Mechanisms for Efficient Advertisement Selection on Public Displays

Terry Payne and Ester David and Nicholas R. Jennings and Matthew Sharifi[1]

Abstract. Public electronic displays can be used as an advertising medium when space is a scarce resource, and it is desirable to expose many adverts to as wide an audience as possible. Although the efficiency of such advertising systems can be improved if the display is aware of the identity and interests of the audience, this knowledge is difficult to acquire when users are not actively interacting with the display. To this end, we present *BluScreen*, an intelligent public display, which selects and displays adverts in response to users detected in the audience. Here, users are identified and their advert viewing history tracked, by detecting any Bluetooth-enabled devices they are carrying (e.g. phones, PDAs, etc.). Within *BluScreen* we have implemented an agent system that utilises an auction-based marketplace to efficiently select adverts for the display, and deployed this within an installation in our Department. We demonstrate, by means of an empirical evaluation, that the performance of this auction-based mechanism when used with our proposed bidding strategy, efficiently selects the best adverts in response to the audience presence. We benchmarked our advertising method with two other commonly applied selection methods for displaying adverts on public displays; specifically the *Round-Robin* and the *Random* approaches. The results show that our auction-based approach, that utilised the novel use of Bluetooth detection, outperforms these two methods by up to 64%.

1 Introduction

Public electronic displays[2] are increasingly being used to provide information to users, to entertain (e.g. showing news bulletins), or to advertise products within public environments such as airports, city centres, and retail stores. Within these displays, advertisers typically utilise a variety of delivery methods to maximise the number of different adverts displayed, and thus increase their overall exposure to target audiences [8]. However, these methods are typically naïve and do not take into account the current audience.

On the other hand, a number of *interactive* public displays have been proposed that support communication with a user through active use of handheld devices such as PDAs or phones, or to a closed set of known users with pre-defined interests and requirements [3, 7]. Such 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 [4, 10] so that their presence can be identified. These approaches fail to work in public spaces, where no prior knowledge exists regarding users who may view the display, and where such displays need to react to the presence of several users simultaneously.

[1] University of Southampton, UK, email:{trp, ed, nrj, mns203}@ecs.soton.ac.uk

[2] An electronic display can change the advert presented over time.

In contrast, we have developed an intelligent public display that utilises a novel approach to detect nearby users, in order to improve the selection of adverts for display. The goal of the selection is to maximise the exposure of as many adverts as possible to as wide an audience as possible (i.e. to maximise the number of distinct adverts seen by the population of users). In doing so, the main advantage of our system design is that it achieves this goal without: (i) any prior knowledge on the audience, (ii) the need for any specific action by the user, or (iii) the need for any client-based software. Moreover, unlike interactive public displays, our detection technology facilitates an awareness of several devices simultaneously.

As no direct feedback is received from the audience and the only knowledge available is based on the past observations of user presence, one of the key challenges of our system is to predict which advert is likely to gain the highest exposure during the next advertising cycle. To approximate this prediction, our system utilizes history information of past users' exposure to certain sets of adverts (so that we don't repeat material they have already seen), along with the information about what users are currently viewing on the display. In particular, we developed a multi-agent auction-based mechanism to efficiently select an advert for each advertising time slot. In this system, each agent represents a stakeholder that wishes to advertise, and it is provided with a bidding strategy that utilises a heuristic to predict future advert exposure, based on the expected audience composition.

In order to evaluate our design, we deployed a prototype of the system in our department (where it advertises information about research projects, press releases and announcements), and developed a simulator which models user behaviour. Here we report empirical results that show that the auction is more efficient in selecting adverts to maximise exposure of this advertising material to a set of users in the shortest time possible. Specifically, the auction requires, on average, 36% fewer advertising cycles to display all the adverts to each user, when compared to the *Round-Robin* approach (or 64% fewer adverts when compared to the *Random* selection approach).

In more detail, this paper advances the state of the art by:

1. Deploying the *BluScreen* prototype which synergistically combines a public screen with the novel detection of nearby handheld devices using Bluetooth wireless technology.
2. Developing a multi-agent auction-based marketplace for effectively marketing adverts on the *BluScreen* prototype.
3. Devising a novel heuristic-based bidding strategy that can be used by the agents in the auction mechanism.
4. Benchmarking our method against two commonly used selection strategies, *Round-Robin* and *Random*, to demonstrate that our method can efficiently display the best set of adverts given the current audience.

The remainder of this paper is organised as follows: Section 2 discusses related work and motivates the use of agents and auction-based marketplaces within *BluScreen*. The deployed system and use of Bluetooth-devices is discussed in Section 3 and the underlying architecture and auction mechanism are described in Sections 4 and 5. In Section 6, we describe the device simulation and our experimental testbed, and present empirical results. Section 7 concludes.

2 Related Work

Targeted advertising has become prevalent through personal advertising systems, such as recommendation systems or web-based banner adverts [1]. These 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). Such approaches work well when advertising to an individual user, where a rich user profile exists a priori; in contrast, *BluScreen* selects adverts based on the presence of several users, where no profile data is available.

CASy [2] extends this targeted advertising metaphor by using a Vickrey auction mechanism to sell advertising space within a *modelled* electronic shopping mall. The auction was used to rank a set of possible advertisements provided by different retail outlets, and select the top ranking advertisements 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. Although a number of different consumer models were evaluated, both with static and adaptive bidding strategies, the system has not been deployed, and is not suitable for advertising to many individuals simultaneously, as it requires explicit interaction with a single user to initially acquire the user's interest and preferences.

The Hermes Photo Display [3] is an example of a community-based display that interacts with users via Bluetooth-enabled phones. Although not used for advertising, users can share photos with each other by either uploading them from a phone to the display, or downloading shared photos to their phone. No specific client software is required on the phone, as the system utilises the fact that many modern mobile phones can share multimedia via Bluetooth. Although this display can be used to share photos; unlike *BluScreen*, the presented content is static, and requires direct user interaction (i.e. via a touch-screen) to browse the photos.

Groupcast [7] was a project that responded to the local audience within a corporate environment to display bespoke media. It had the advantage of knowing a priori the profiles of several members of the audience, and thus could exploit this pre-defined knowledge. User identification was based on an infrared badge system and embedded sensors within an office environment. When several users passed by the display, Groupcast compared the user's profiles to identify common areas of interest. Whereas *BluScreen* selects the best content to maximise exposure to the current audience, Groupcast would try to present content that matched this common interest. However, acquiring such content a priori, and determining high fidelity user profiles were found to be major stumbling blocks, and the mutual content selection was finally abandoned.

3 The BluScreen Prototype

The notion that user presence in front of the intelligent display can be identified remotely, has been explored using several methods within different pervasive projects. However, most require the deployment of specialised hardware, such as *Active Badges*, [4, 10]. The Bluetooth wireless protocol, characterised by its relative maturity, market penetration[3], and emphasis on short-range communication (e.g. 2-3 meters), has emerged as an alternative, generic mechanism for identifying or interacting with users through small, handheld devices.

BluScreen utilises Bluetooth-enabled devices as proxies for identifying users. Bluetooth sensors attached to an intelligent screen can be used to identify these devices[4] in the local environment, and can record the encounters as a collocation event in terms of location and duration. The duration of this collocation event is assumed to relate to a possible level of interest in the displayed material; e.g. a user who is interested in the current advertising material will linger at the display during the advert.

A *BluScreen* prototype has been developed and deployed to evaluate the feasibility of the auction-based approach. A 23 inch flat-screen display and Bluetooth sensor were deployed outside an office adjacent to the corner of two corridors and an exit (thus maximising visibility to individuals moving within both corridors). The environment was scanned for Bluetooth devices every 20 seconds[5], and these scans were logged over six months to determine whether or not Bluetooth devices could be used as a proxy for human presence. Twelve adverts were generated, describing a range of topics including research projects, upcoming events, and general information.

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.

Table 1 summarises the number of devices detected, split into three clusters: *occasional*, whereby devices are observed for only a short time (such as visitors, etc.); *frequent* representing devices that regularly pass by the display; and *persistent* which represents devices that are regularly found in proximity to the sensor (such as laptops, etc.) that do not represent individuals passing the screen. Of the 212 unique Bluetooth devices detected, approximately 70 were detected with some degree of regularity, suggesting that Bluetooth detection could be used as a proxy for individuals passing in front of a screen.

4 The Agent Architecture

Several types of agents have been designed within the *BluScreen* architecture that we developed (illustrated in Figure 1):

Bluetooth Device-Detection Agent This agent monitors the environment by periodically sampling it to determine if there are any Bluetooth devices detectable. Current samples are then compared to previous samples to determine if new devices are present or existing devices have left the environment, and to update the exposure rate of existing devices.

Advertising Agent Each Advertising Agent represents a single *Advertisement* and is responsible for effectively purchasing advertising space by generating bids based on their predicted valuation of the upcoming advertising cycle. This predicted valuation may be

[3] In October, 2005, ElectronicsWeekly.com reported that "*...Bluetooth is expected to ship around 270 million units during 2005...*"

[4] Only Bluetooth devices in discovery mode are detectable.

[5] The choice of a 20 second scanning cycle was determined by evaluating different scanning cycle lengths with varying numbers of nearby devices.

based on the expected exposure to an arbitrary audience (i.e. being seen by as many users as possible) or could be more selective (e.g. by identifying whether or not users have a preference for the type of material contained in the agent's advertisement). However, our deployed prototype described here currently implements the former method. Thus, each advertising agent maintains a history of successful advertising cycles (i.e. for which cycles an agent won an auction, and thus displayed its advert on the *BluScreen* display), and also the time (and duration) of each Bluetooth-device that was detected by that display during each advertising cycle.

Marketplace Agent The marketplace agent facilitates the sale of advertising time slots. A repetitive second-price sealed-bid auction [9] is used as a selection mechanism (this choice is justified below) to determine which advertising agent (if any) will display its advert on the next time slot, and its corresponding payment.

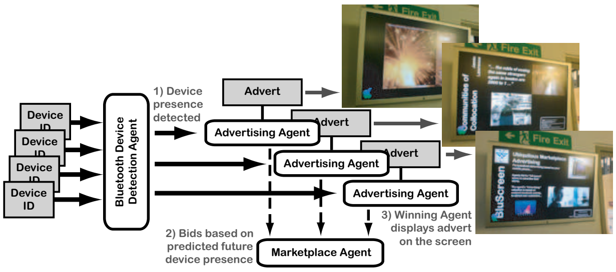

Figure 1. The BluScreen Agent Architecture for a single intelligent display

5 Auction Mechanism

BluScreen is designed to support a scaleable and a dynamic advertising framework, while maximising the exposure of as many adverts as possible to as wide an audience as possible, within a knowledge-poor environment. The main principle of our design is to distribute the control of the content displayed, in a way that no single entity can dictate who will advertise next. In contrast, the system, as a whole, will decide who will be the most profitable agent (i.e., expected to gain the highest exposure by displaying its advert in the next advertising cycle) and therefore will be awarded the facility of advertising in that cycle. Note that although the advertising agents are self interested (i.e. aim to maximise their own exposure to as large an audience as possible using their fixed budget), this does not contradict our desired overall design goal, since in our context, we assume that each agent will bid a value that reflects its expected exposure from displaying on the next cycle. Specifically, we have implemented a repetitive, second-price sealed-bid auction that takes place before each of the advertising cycles[6]. The winner of each auction is the advertising agent who placed the highest bid, and it is charged for the second highest bid. As truth-revealing has been shown to be a dominant strategy in this case, the effective local decisions of each individual agent contribute towards effective overall system.

In this section we start by describing the BluScreen model and the mechanism we developed, including the strategies of the advertising agents given their valuation for the next advertising cycle's auction. However, as our system is a knowledge-poor environment where no direct feedback is collected from the users (instead only a limited history and the current status of the detected devices is available), in the final part of this section we provide the agent with a heuristic method to generate this value.

5.1 The Auction Model

In our model we assume that there is a single *BluScreen* instance available for advertising which is managed by the marketplace agent[7]. The advertising time is divided into discrete intervals which we term the *advertising cycle*, C^i. During each such cycle, only one agent can display its corresponding advert, and during this time only this agent has the ability to collect information about devices which were detected by the *BluScreen*. However, at the end of each advertising cycle, the marketplace agent informs all the advertising agents about any devices detected in front of the screen at that time.

Each advertising agent, a_j, is assumed to be self interested, and to be managed by a different stake-holder which credits it with a fixed budget for advertising. For each advertising cycle, each advertising agent generates its private value (discussed in detail below) and places a corresponding bid in terms of some currency. Once an agent has fully spent its budget, it is deleted from the system. On the other hand, at any point in time, new agents may be added to the system. The goal of each advertising agent is to utilise its budget in a way that maximises its exposure.

The marketplace agent acts as the auctioneer and runs an auction before each of the advertising cycles. As part of the auctioneer duties, it sets a reservation price that represents the minimum acceptable bid[8]. In the case that none of the agents place bids that meet this price, a default advert is displayed. The marketplace agent is also assumed to be trusted by the advertising agents. This feature is necessary for the proper behaviour of the system, as in a second-price sealed-bid auction no one observes the whole set of bids except the auctioneer.

5.2 Bidding Strategy for the Advertising Agents

In our model, agents participate in a *repetitive*, second-price sealed bid auction whilst facing budget constraints. For non-repetitive variants of this auction, it has been shown that the bidding strategy:

$$\beta(a_j) = min\,\{budget(a_j), v\} \tag{1}$$

is a ***dominant strategy*** [6], where $budget(a_j)$ is the current balance of agent a_j and v is its private valuation. Now, it is easy to see that this holds for the repetitive case as well, as in each round of the second-price sealed-bid auction, a new advertising cycle is auctioned. However, there exists a dependence with previous rounds in terms of the devices that were exposed in the past, and those that are currently in front of the screen. However, this dependence is concealed in the valuation of the bidder for the next advertising cycle.

[6] We chose this mechanism for several reasons. As one of the critical requirements of this system is to operate in a timely manner, only *one-shot* auctions were considered. Now, *first-price* and *second-price* sealed-bid auctions are the most popular forms of such auctions. Here we chose the second-price for its superiority in terms of its economically desired properties. In particular, it has a dominant strategy of truth revealing. Such a property is highly desired as it frees the agents from requiring expensive and costly strategies in terms of time and computation. In our context, this is even more significant as the system should react in a timely manner. Moreover, the second-price sealed-bid auction is efficient in the sense that the auctioned item, in our case the right to advertise on the next advertising cycle, is awarded to the bidder that values it the most. Our choice can also be supported by the fact that the continuously repeated second-price sealed-bid auction is also being used for web advertising using search engines like Google and Overture [5].

[7] This assumption is made to simplify our task of validating the correctness and the efficiency of the proposed mechanism and strategy. This assumption will be relaxed in future work.

[8] The reservation price is used to avoid selling the advertising space for free, and to ensure that when competition is low, there is a minimum income.

But, once the agent has generated its valuation for the next advertising cycle, its dominant strategy is truth telling.

Although the agents utilise truth telling as a dominant strategy, they still face a very challenging problem of being able to estimate their ***expected exposure***, as they are provided only with information about past and current audience composition (in terms of detectable Bluetooth devices), and from that alone they need to estimate their future exposure. In the next subsection we describe a heuristic method we developed for the bidder valuation estimation task.

5.3 Heuristics for Estimating Private Values

Here we provide an advertising agent, a_j, with a heuristic strategy for generating its valuation (expected utility) for the next advertising cycle, C^{i+1}. Recall that the main target of an agent a_j is to maximise its exposure to as large an audience as possible, given its financial constraints. The core principle of this heuristic is to utilise the information an agent has about past devices' exposure, to predict the future expected exposure. Specifically, each time an advertising agent a_j has to make a decision about its valuation for the next cycle C^{i+1}, it has two types of information on which to base its decision:

(i) history observation, $H(a_j)$, of exposed devices which were collected during the advertising cycles it won, $WonCycles(a_j)$, in the past, $H(a_j) = \{(C^t, d, x)\}$ where $C^t \in WonCycles(a_j)$, d is the device id, and x is the exposure duration; and

(ii) the current set of detected devices which were in front of the screen at the end of C^i, termed $end(C^i)$.

Using this information, an advertising agent, a_j, assumes that the devices that will be in front of the screen in C^{i+1} are $end(C^i)$. Therefore, we propose that a_j will search through its history to check if these devices were exposed to its advert in the past, and, if so, for how long[9]. If a device was exposed to the same advert during several different advertising cycles, then we propose to consider only the one in which it was exposed to the most. Formally, a_j's valuation for C^{i+1} is:

$$v(a_j, C^{i+1}) = \sum_{d \in end(C^i)} 1 - max\left\{x \middle| \left\{C^t, d, x\right\} \in H(a_j)\right\}^{10} \tag{2}$$

Now, we believe this is the best heuristic an agent can use, given its limited information it faces currently. We intend to incorporate more sophisticated mechanisms for generating statistics within future versions of *BluScreen*, that can provide the agents with more relevant information, such as the rate of detected devices as a function of the time in the day. Given that information, agents will be able to apply learning methods to improve their predictions.

6 Empirical Evaluation

To evaluate the bidding strategy described in the previous Section within a controlled environment, a simulation was developed. This is based on the same architecture as the deployed *BluScreen* instance, but it additionally supports the modelling of users in the audience, in terms of Bluetooth devices present[11].

The *BluScreen* simulation modelled devices' behaviour (i.e. as a proxy for a user) in terms of their likelihood of arriving at, and subsequently remaining at, a *BluScreen* display given the currently displayed advertisement. The *device model* was defined with the following assumptions:

- Device presence is measured in discrete sample-intervals, to simulate observations made by the Bluetooth sensor;
- The duration of an advert is assumed to be equal to a whole number of sample-intervals;
- An advert is considered as *fully seen* only when a device has been present for the whole duration of the advert;
- Devices can arrive at any point during an advertising cycle, whereby the probability that a device will arrive is P_{arrive};
- The probability a device may leave without observing the subsequent advert is P_{depart}. A device will automatically leave if it has *fully seen* the subsequent advert;
- Both P_{depart} and P_{arrive} assume a uniform distribution.

Two alternate selection methods were compared with the auction: ***Round-Robin*** selection and ***Random*** selection. The former is a familiar mechanism used to repeatedly cycle through a number of adverts (in order). Given our device model, this approach represents the optimal selection mechanism when $P_{depart} = 0$ and $P_{arrive} = 1$. The latter mechanism randomly selects a new advert to display at the beginning of each new advertising cycle, independently of what has been previously selected. This method was selected as a baseline against which the other worst-case methods can be compared.

To simplify this analysis, each experiment consisted of ten adverts of equal duration (i.e. six samples which is equivalent to a 2 minute advert), represented by ten advertising agents. In each experiment, adverts are selected and presented until every device has *fully seen* all ten adverts. To ensure statistical stability, the results of each experiment are averaged over 10,000 test runs, and the mean results are reported. Where necessary, a Student's t-test is used to confirm statistical significance at the 95% confidence level.

6.1 Experimental Results

In our first experiment, we examine the behaviour of each selection mechanism as the number of devices present increases (Figure 2). The number of devices, N_d, was varied between 1 and 100, and the device behaviour was defined using $P_{depart} = 0.05$ and $P_{arrive} = 0.5$. The graph supports the hypothesis that advert selection using the *BluScreen auction* is statistically more significant than either *Round-Robin* or *Random* (assuming the modelling parameters defined above). Specifically, as N_d increases, there is a corresponding exponential rise in the number of required adverts. The mean number of adverts required by the *BluScreen auction* is lower than the *Round-Robin* selection mechanism or the *Random* selection mechanism for all numbers of devices tested; e.g. for $N_d = 50$, the auction method required a mean of 40.24 ± 0.06 advertising cycles to display all 10 adverts to all the devices, compared to *Round-Robin* (64.16 ± 0.18) or *Random* (108.50 ± 0.28). This suggests that although each selection method can scale, a single device will be exposed to all the adverts in typically 36% fewer advertising cycles than

[9] If a device only observes part of an advert, a_j still has an incentive to expose its advert to this device, as we make the assumption that the advert may be dynamic (e.g. contain video). Therefore a device is only exposed to all the information contained in the advert if it observes the advert for the whole advertising cycle. However, having observed a partial advert will diminish the devices' contribution towards bids in future cycles.

[10] In the case where a device is not exposed to any of the previous displays, we assume the default value of x is zero.

[11] Although the deployed prototype could be used to test the auction mechanism based on observed real-world events, factors such as the location of the screen, variances in the advertised material, and the number of detectable devices can affect the resulting performance.

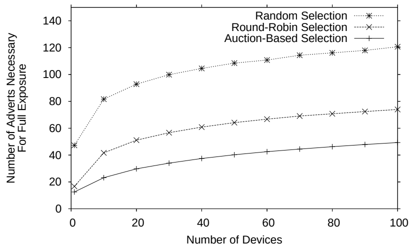

Figure 2. The effect of varying the number of observed devices; $P_{depart} = 0.05$, $P_{arrive} = 0.5$.

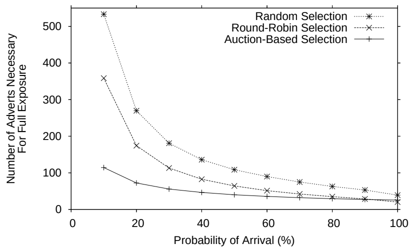

Figure 3. The effect of varying the probability of a device arriving during an advertising cycle. $P_{depart} = 0.05$, $N_d = 50$.

the *Round-Robin* approach, or 64% fewer advertising cycles than the *Random* approach when using the *BluScreen auction*.

The second two experiments therefore explore how the efficiency of the different selection mechanisms change as the device model is varied. In each of these experiments, $N_d = 50$. In particular, Figure 3 plots the results for $0 < P_{arrival} \leq 1$. These results show that for low values of $P_{arrival}$, the number of adverts required is significantly lower for *BluScreen auction* than for either of the other methods considered. However, as the probability increases towards 100%, the efficiency of both *BluScreen auction* and *Round-Robin* converge. This result is expected, as when all devices are guaranteed to be present, *BluScreen auction* will select each advert in turn (depending on the tie-breaking strategy (e.g. random) used by the auctioneer), and thus exhibit the same behaviour as *Round-Robin*.

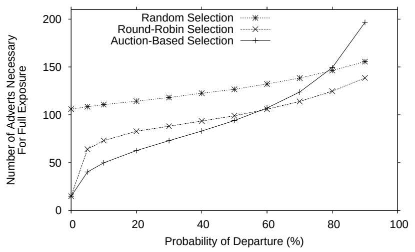

Figure 4. The effect of varying the probability of a device leaving the audience during an advertising cycle. $P_{arrive} = 0.5$, $N_d = 50$.

One artefact of assuming low departure probabilities is that once a device has arrived, it is unlikely to depart, and hence is more likely to remain through a complete cycle of adverts (in the case of *Round-Robin*), or its presence be accurately predicted (when using *BluScreen auction*). Therefore, as P_{depart} increases, the efficiency of both approaches should degrade. Figure 4 illustrates this. As one would expect, the performance of both *BluScreen auction* and *Round-Robin* are statistically equivalent when $P_{depart} = 0$. However, the number of adverts required increases sharply in the range $P_{depart} = [0.01 \ldots 0.1]$; and then rises almost linearly until $P_{depart} = 0.5$. Beyond this value, the rise becomes logarithmic. This rise is more significant for the auction approach, and may be due to poor estimates in future device presence, leading to erroneous private value predictions. In contrast, the rise in the number of adverts required for *Random* is almost linear; this result is not surprising as the selection of each advert is independent of the audience composition, and hence is unaffected by the departure of a device.

7 Conclusions

In this paper, we addressed the problem of unobtrusively providing audience composition data to an intelligent public display, to select the optimal advert that maximises exposure of a set of adverts. To this end, we proposed the novel use of Bluetooth wireless technology to detect nearby handheld devices, and designed an auction-based marketplace and bidding heuristic which predicted the level of expected exposure of an advert given the detected audience and knowledge of previous exposure the audience members had to different adverts. A prototype public display that utilises this marketplace was implemented and deployed in our Department. The logs generated by this system indicated that a large number of Bluetooth-enabled devices could be detected, and that a significant subset was detected on a regular basis. Moreover, we implemented a simulation environment to facilitate controlled evaluation of the auction mechanism, when compared with other commonly used advert selection mechanisms, and showed that in most cases the auction-based approach outperforms the other commonly used approaches. This approach has particular relevance for advertising in environments where users only linger for short periods, but often return (such as railway platforms or office entrances).

In future work, we plan to extend our approach to work with multiple displays, to increase the exposure of adverts to a larger community, and to build user preferences based on the duration users spend viewing different adverts that are categorised (thus taking into account inferred interest, as well as exposure).

REFERENCES

[1] Ali Amiri and Syam Menon, 'Efficient scheduling of internet banner advertisements', *ACM Trans. on Internet Tech.*, **3**(4), 334–346, (2003).

[2] Sander M. Bohte, Enrico Gerding, and Han La Poutre, 'Market-based recommendation: Agents that compete for consumer attention', *ACM Trans. on Internet Tech.*, **4**(4), 420–448, (2004).

[3] Keith Cheverst, Alan Dix, Daniel Fitton, Chris Kray, Mark Rouncefield, Corina Sas, George Saslis-Lagoudakis, and Jennifer G. Sheridan, 'Exploring bluetooth based mobile phone interaction with the hermes photo display', in *MobileHCI '05: Proc. 7th Int. Conf. on Human computer interaction with mobile devices & services*, pp. 47–54, (2005).

[4] Jeffrey Hightower and Gaetano Borriella, 'Location systems for ubiquitous computing', *IEEE Computer*, **34**(8), 57–66, (2001).

[5] Brendan Kitts and Benjamin Leblanc, 'Optimal bidding on keyword auctions', *Electronic Market*, **14**(3), 186–201, (2004).

[6] Vijay Krishna, *Auction Theory*, Academic Press, 2002.

[7] Joseph F. McCarthy, Tony J. Costa, and Edy S. Liongosari, 'Unicast, outcast & groupcast: Three steps toward ubiquitous, peripheral displays', in *UbiComp '01: Proc. 3rd Int. Conf. on Ubiquitous Computing*, pp. 332–345, London, UK, (2001). Springer-Verlag.

[8] Anand Ranganathan and Roy H. Campbell, 'Advertising in a pervasive computing environment', in *WMC '02: Proc. 2nd Int. Wkshp. on Mobile commerce*, pp. 10–14, New York, NY, USA, (2002). ACM Press.

[9] W. Vickrey, 'Counterspeculation, Auctions and Competitive Sealed Tenders', *Journal of Finance*, 8–37, (1961).

[10] Roy Want, Andy Hopper, Veronica Falcao, and Jonathan Gibbons, 'The active badge location system', *ACM Trans. on Information Systems*, **10**(1), 91–102, (1992).