Auction Mechanisms for Efficient Advertisement Selection on Public Displays (Extended Abstract)[1]

Terry Payne Ester David Nicholas R. Jennings Matthew Sharifi

University of Southampton, SO17 1BJ, United Kingdom
{*trp,ed,nrj,mns203*}*@ecs.soton.ac.uk*

1 Background

Public electronic displays are increasingly used to provide information such as alerts, event details, or notifications within public environments such as airports, city centres, and retail stores. Such systems typically utilise a variety of delivery methods to maximise the number of different adverts displayed, and thus increase their overall exposure to target audiences. However, these methods are typically naïve and do not take into account the composition of the current audience. In contrast, *interactive* public displays support communication with users through active use of their handheld devices such as PDAs or phones. 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 to identify their presence. This fails to work in public spaces, when no prior knowledge exists about the audience in front of a display, and where such displays need to react to the presence of several users simultaneously.

We have developed *BluScreen*, an intelligent public display that detects nearby users through their Bluetooth-enabled PDA or phone, 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 have developed a multi-agent auction-based mechanism to efficiently select an advert for each advertising time slot. 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.

2 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. Specifically, we have implemented a repetitive,

[1] This is an extended abstract of [1].

second-price sealed-bid auction that takes place before each of the advertising cycles. 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.

To identify users, *BluScreen* detects Bluetooth-enabled devices in the vicinity of the display, through attached Bluetooth sensors. These encounters are recorded as collocation events in terms of location and duration. The duration of each 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.

An advertising agent, a_j, uses a heuristic strategy for generating its valuation (expected utility) for the next advertising cycle, C^{i+1}. 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 set of detected devices found in front of the screen at the end of C^i, termed $end(C^i)$.

Using this information, an advertising agent, a_j, searches through its history to check if the devices in set $end(C^i)$ were exposed to its advert in the past, and, if so, for how long. This is then used to generate the valuation for the next advertising cycle, and submit this as a bid. Formally, a_j's valuation for C^{i+1} is:

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

3 Results

To evaluate the bidding strategy described in [1], a simulation was developed that models device behaviour in terms of their likelihood of arriving at, and subsequently remaining at, a *BluScreen* display given the currently displayed advertisement. Two alternate selection methods were compared with the auction: *Round-Robin* selection and *Random* selection (described in [1]). The former is a familiar mechanism used to repeatedly cycle through a number of adverts in order.

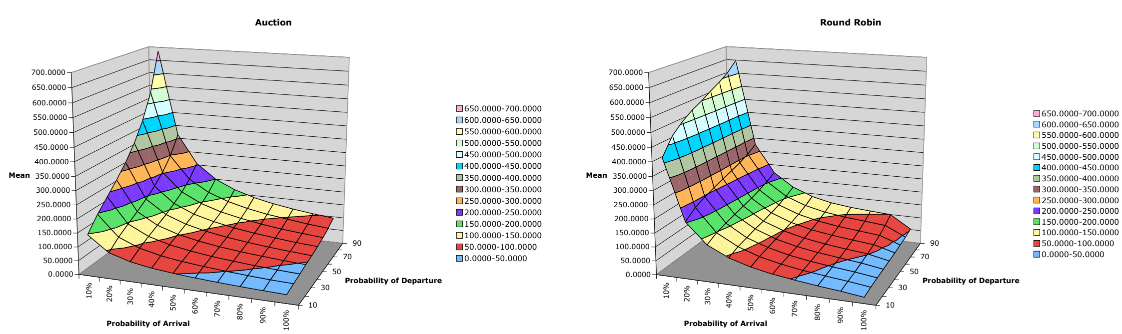

Figure 1: *The graphs show how the Auction selection method compares with the Round-Robin selection method with different types of user behaviour, in terms of their arrival and departure probabilities (i.e. the x and y axes respectively). The z axis illustrates the number of advertising cycles required before all the users (50 in these experiments) observe all 10 adverts within the experiment.*

Figure 1 illustrates the performance of the *Auction* selection method (described in Section 2) with the *Round-Robin* selection method with different types of user behaviour. The Auction method requires significantly fewer advertising cycles than Round-Robin for cases where the departure probability is less than the arrival probability, due to the predictive nature of the valuation. On average, the auction requires, 36% fewer advertising cycles to display all the adverts to each user, when compared to the *Round-Robin* approach.

References

[1] Terry Payne, Ester David, Nicholas R. Jennings, and Matthew Sharifi. Auction mechanisms for efficient advertisement selection on public displays. In *Proceedings of 17th European Conference on Artificial Intelligence (ECAI-06), Riva del Garda, Italy*, pages 285–289. IOS Press, 2006.

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