Heuristic Bidding Strategies for Multiple Heterogeneous Auctions

David C. K. Yuen[1], Andrew Byde[2], Nicholas R. Jennings[1]

Abstract. This paper investigates utility maximising bidding heuristics for agents that participate in multiple heterogeneous auctions, in which the auction format and the starting and closing times can be different. Our strategy allows an agent to procure one or more items and to participate in any number of auctions. For this case, forming an optimal bidding strategy by global utility maximisation is computationally intractable, and so we develop two-stage heuristics that first provide reasonable bidding thresholds with simple strategies before deciding which auctions to participate in. The proposed approach leads to an average gain of at least 24% in agent utility over commonly used benchmarks.

1 Introduction

The growing number of online auction sites liberates individual shoppers from searching only within their surrounding community. Moreover, if the item is popular, it is not uncommon to find hundreds of auctions offering the same good[3]. However, it can be difficult for a human to keep track of more than a few auctions. But, *global bidders* that can accept information and participate in a large number of auctions are likely to receive extra benefit when compared with bidders that only operate locally in one auction [8]. When all of these factors are taken together, it is apparent that there is a growing need to develop autonomous agent bidding strategies that can operate over multiple auctions.

To this end, we seek to devise a best response bidding strategy for global bidders participating in multiple auctions. Each such bidder may have a demand for one or more units of identical private-value goods. Our primary interest is in multiple heterogeneous auctions, because it is the most prevalent case in practice, in which the auction format and the starting and closing times can be different. We assume each seller auctions one item and these items are perfect substitutes. Due to its complex nature, most analytical studies of heterogeneous multiple auctions are based on the simplified case of homogeneous auction types that either operate sequentially [9] or completely in parallel [8]. For our more general setting, heuristic approaches are the norm [1, 2, 3, 4, 5] and are the approach we adopt here.

Now, to develop a bidding strategy that maximises the expected utility, it is theoretically possible to model the multiple auctions as a Markov Decision Process (MDP) and calculate the expected utility by backward induction [2, 3]. The utility evaluation can be simplified by adopting pessimistic estimation [4], in which an agent sets its bids using no information on the ordering of the future auctions. However, issues associated with computational intractability limit these approaches to only a low number of auctions. In contrast, Anthony *et al.* [1] devised a heuristic bidding strategy that combines the effects of the desire to obtain a bargain, the deadline and the desperateness for obtaining an item into a single formula. However, their bidding tactic ignores the fact that some of the remaining auctions can be more favourable to the global agent (e.g. have fewer local bidders). Relatedly, the bidding strategy developed by Dumas *et al.* [5] consists of a probabilistic bid learning and planning stage. However, their algorithm considers simultaneous auctions as incompatible and would bid in only one of them.

Against this background, this paper develops a two-stage heuristic approach to approximate the best response bidding strategy for a global bidder. In the first stage, a threshold heuristic is employed to compute a maximum bid or threshold for each auction. Then in the second stage, the agent decides whether it should participate in each of the available auctions using an auction selection heuristic that exploits the bidding thresholds calculated from the first stage. A number of threshold and auction selection heuristics are developed under this flexible framework. In developing these heuristics, this work advances the state of art in a number of ways. First, we devise several two-stage heuristics that enable the agent to improve its utility further by bidding in more auctions than its demand. This tactic requires the agent to manage the excess procurement risk carefully and is an area ignored by previous bidding heuristics other than those using MDPs or exhaustive search approaches. Second, we provide a pseudo-polynomial time auction selection heuristic that closely approximates the near optimal solutions of an exponential time exhaustive search algorithm from the literature [4]. Third, we extend the use of equal bidding thresholds to maximise the agent utility from complementary goods [8] to our perfect substitute setting and verify that the agent's utility received is very close to the global maximum (at least for simultaneous second-price auctions).

The paper is structured as follows. Section 2 details the heuristics applied to set the threshold and select the auctions to participate in. Section 3 empirically evaluates the strate-

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

[2] HP Labs, Filton Road, Stoke Gifford, Bristol, BS34 8QZ, UK. andrew.byde@hp.com

[3] To illustrate the scale of this problem, within eBay alone, more than a thousand auctions were selling Apple's 4GB iPod mini at the time of writing.

gies and compares them with other heuristics proposed in the literature. Section 4 concludes.

2 The Multiple Auction Bidding Strategy

Before detailing the bidding heuristics, we formalise the multiple auction setting being studied. Here an auction a can be characterised by a tuple $(t^s(a), t^c(a), N(a), \theta(a))$ consisting of the starting time, closing time, number of local bidders and auction format for auction a. Notice that since $N(a)$ is the number of *local* bidders in an auction; the total number (including the agent itself) is $N(a)+1$. The auction format can be one of the standard four single-sided types, English, Dutch, first-price (FPSB) or second-price sealed bid (SPSB), denoted as {ENG, DUT, FPSB, SPSB}. We introduce the degree of overlap $\lambda \in [0, 1]$:

$$\begin{array}{lcl} \text{time with at least 1 running auction:} & T_{occ} & = |\{t \in \mathbb{N} | \exists a : t^s(a) \leq t < t^c(a)\}| \\ \text{degree of overlap:} & \lambda & = \dfrac{\sum_a (t^c(a) - t_s(a)) - T_{occ}}{T_{occ}\,(M-1)} \end{array}$$

to characterise a set of M auctions by the starting and closing times. Purely sequential and purely simultaneous auctions are the two extreme cases with $\lambda = 0$ and 1 respectively.

The agent's valuation V for obtaining multiple items is linear with random coefficient v per item up to a maximum of k items: the agent derives value $V(x) = v \min(k, x)$ from x items, and the utility of the agent for obtaining x items at cost Π is quasi-linear: $U(x, \Pi) = V(x) - \Pi$.

We assume that all bidders have identically distributed independent item-valuations v with distribution $F(x) = Prob(v \leq x)$, $f(x) = F'(x)$. We will use the uniform distribution in $[0, 1]$ as an example throughout the text, but the heuristics themselves are equally applicable to other distributions unless specified otherwise. For a set of M auctions, our goal is to devise a bidding strategy $b = (b_1, b_2, \cdots, b_M)$ where b_a gives the threshold (i.e. the maximum bid that should be placed) for auction a. In sealed bid auctions, the threshold is the bid placed, in Dutch auctions, the agent bids when the current price falls the threshold, and in English auctions, the agent bids until the current price reaches the threshold. The threshold is set to zero if the auction selection heuristic indicates that the agent should not participate in that auction.

At any point in time, the bidding strategy considers the subset of auctions A^{av} that is *available*. An auction is *unavailable* if it is either closed or is of type ENG and the current price is already higher than the agent's threshold. We use $M^{av} = |A^{av}|$ for the number of available auctions. The agent is expected to keep track of the number of acquired items. At time T, the agent has an extra capacity of k_T items $= k - x_T$ where x_T is the number of items already acquired.

2.1 The Threshold Heuristics

The threshold set required by the auction selection stage is generated using one of the heuristics developed in this section. In all cases, the setting of the threshold is affected by the differences between individual auctions (e.g. the number of bidders and the auction format) and by the availability of the item in future auctions. Since we want to identify the most influential factors to the received utility, each of our heuristics examines only a subset of the above factors. Specifically, we develop two heuristics. The *single auction dominant* heuristic assigns different thresholds according to the auction format and the number of bidders in each auction. The *equal threshold* heuristic assigns the same threshold to every auction and adjusts this according to the availability of the item in future auctions. Some heuristics in the literature have also examined the bidding history [5] and user preferences [1] when adjusting the bidding level. But, with our modular two-stage heuristic design, these external factors can easily be included later by extending or replacing the existing heuristics.

2.1.1 The Single Auction Dominant Heuristic

This heuristic sets a different threshold for each of the auctions according to the auction format and the number of bidders. Specifically, the threshold is assigned as if it were the only auction available and the agent had valuation v for a single item. We choose the single auction dominant bidding heuristic not only because its analytical solution is well known [7], but also because of its ability to adjust the bid for different auction types and numbers of bidders with little computational overhead. The threshold is the same as the true value for second price mechanisms (ENG, SPSB) and for the first price mechanisms (DUT, FPSB), it becomes [7]:

$$b_a = v - F^{-N(a)}(v) \int_0^v F^{N(a)}(x)dx$$

which is $v\,N(a)/(N(a)+1)$ if the private values of the bidders follow a uniform distribution.

2.1.2 The Equal Threshold Heuristic

This heuristic sets a single threshold, b_{eq}, for all auctions. This choice is justified because the global bidder's utility in multiple auctions selling complementary goods can be maximised with equal-bid pairs [8] and our brute-force best bidding strategy search upon simultaneous auctions with unit agent demand ($k = 1$) seems to suggest that placing equal bids is also suitable for our perfect substitute setting[4].

Our goal is to find the equal bid threshold that approximately maximises expected utility:

$$\begin{aligned} E(U|b) &= E(V|b) - E(\Pi|b) \\ &= vE(\min(k, X)|b) - \sum_{a \in A^{av}} E(\Pi(a)|b) \end{aligned} \quad (1)$$

To do this, we adopt a procedure that estimates the likely number of units won (X) and the likely payments (Π) separately. Now, the expected payment $E(\Pi|b)$ is the sum of payments for each auction and the expected payment $\Pi(a)$ in a particular auction a from using threshold $b(a)$ is:

$$E(\Pi(a)|b(a)) = \int_0^{b(a)} yF^{N(a)-1}(y)f(y)dy \quad (2)$$

For valuations v that are uniformly distributed in $[0, 1]$ this gives an expected payment of $N(a)\,b^{N(a)+1}/(N(a)+1)$. However, the valuation term is non-linear with respect to the number of units won, and so the situation is more complicated.

[4] It should be noted that the optimality of equal bids is not universal even for simultaneous SPSB auctions with identical bidder distributions. For example, in the extreme case in which the valuation approaches the upper range, it is best to bid one's true value for k requested items and zero for the rest.

Thus to estimate the number of units the agent might win by using a particular bid threshold, we replace the available current and future auctions with simultaneous SPSB auctions. We believe this is a reasonable assumption for two reasons. First, the expected revenue of second-price auctions is only slightly lower than those of first-price ones. Second, unlike currently running English (or Dutch) auctions in which we know the closing price must be larger (or smaller for Dutch) than the current price, we do not have such information regarding future auctions. So, in terms of information availability, it resembles the situations in simultaneous sealed bid auctions.

Now, to measure the number of bidders $\bar{N}$ in the set of available auctions A^{av} we use the harmonic mean $\frac{1}{\bar{N}} = \frac{1}{M^{av}} \sum_{a \in A^{av}} \frac{1}{N(a)}$ because the agent is less likely to win in auctions with many bidders and so should give less weight to their consideration.

In so doing, we assume that the events of winning each auction are independent and identically distributed, so that the number of items the agent will win has a binomial distribution. The probability of an agent with threshold b winning in any second-price auction with n other local bidders (who will bid their valuation) is $p(b) = F^n(b)$. The probability of winning x items out of M^{av} auctions is:

$$Prob(X = x|b) = \left(\begin{array}{c} M^{av} \\ x \end{array} \right) p(b)^x (1 - p(b))^{M^{av} - x}$$

Combining this with:

$$E(\min(k, X)|b) = \sum_{x \leq k} x Prob(X = x|b) \; + \; k Prob(X > k|b)$$

allows us to estimate the value term in (1).

Since the agent is not compensated for any extra units acquired, its utility drops rapidly if the purchase quota is exceeded accidentally. However, to identify the equal threshold b_{eq} that approximates the maximum utility, a simple one dimensional *golden section search* [10] is sufficient because the utility curve has only a single maximum.

2.2 The Auction Selection Heuristics

The auction selection heuristics determine which available auctions the agent should participate in, on the basis of the collection of bidding thresholds calculated using a threshold heuristic. With our flexible architecture, it is possible to introduce heuristics that satisfy different user preferences [1] at a later stage. But for now, we intend to develop heuristics that approximate well to the global utility maximum, which is the goal for most users. Specifically, we describe two such heuristics: an *exhaustive search* heuristic inspired by [4], and a much more computationally tractable *knapsack utility approximation* heuristic.

2.2.1 Exhaustive Search Selection

Byde *et al.* [4] describe an exhaustive search procedure for bid decision making. Their algorithm subsumes the selection of auctions into the choice of thresholds: ideally *all* tuples of thresholds would be tested, which includes the "trivial" zero thresholds that are equivalent to not participating in an auction at all. In our context, where thresholds are determined by a separate process, exhaustive search consists of testing almost all subsets of available auctions. Nevertheless, some heuristic properties of different auction formats are exploited to prune the search space. For example, in English auctions, we assume that an agent prefers to bid in auction a_1 if its current price is lower than that of a_2 while the expected chance of winning is higher. In sealed bid auctions, the agent refrains from submitting a bid until the deadline is imminent so that it can learn from more auctions before making a decision.

The expected value $E(V|S, b)$ of bidding only in a set of auctions S is calculated using the sum over all possible sets $W \subseteq S$ of auctions that might be won:

$$E(V(X)|S, b) = \sum_{W \subseteq S} \left(V(|W|) \prod_{a \in W} p_a(b(a)) \prod_{a \notin W} (1 - p_a(b(a))) \right) \quad (3)$$

where $p_a(b(a))$ is the probability of winning auction a given threshold $b(a)$. Expected payments are given by (2) as above.

Clearly this heuristic is expensive to calculate both because of the exhaustive search over the set of auctions S in which to bid, and because of the exponential complexity of (3).

2.2.2 Knapsack Utility Approximation

The development of this heuristic is inspired by our observation that the utility is non-linear with respect to the number of wins. Specifically, each additional acquired item provides extra utility until the maximum demand k is reached. Beyond that point, each additional item incurs a large marginal loss because the agent does not receive any extra compensation for the increased payment. Now, since the number of wins is a random variable and the utility plummets if the agent acquires more than k items, it is crucial to carefully control the risk of exceeding the demand limit. Given this, the rationale for the heuristic proposed is to identify a subset of available auctions that minimises the expected payment (i.e. locally maximises the utility), while maintaining the expected number of wins within a certain safe level. The knapsack algorithm was chosen because its solution can be approximated in pseudo-polynomial time by dynamic programming [6].

In more detail, the heuristic works as follows. First, the agent (re)-evaluates its capacity for extra items. For planning purposes, the extra capacity $\hat{k}_T = k_T - H$ should be reduced by the number of holding auctions H. A holding auction is any sealed bid auction that the agent has submitted a bid to and has yet to receive the results or any English auction in which it is currently leading. The intuition here is that the agent has a chance to win once it holds a bid in an auction. Therefore, it should temporarily reduce the demand limit to prevent accidentally exceeding its purchase quota. Second, by making a similar argument to that of Section 2.1.2 when we simplified the auction setting, we model the multiple auctions as if they are all SPSB and run simultaneously. Given this, the threshold set generated by the threshold heuristics is converted into a M^{av}-tuple $b = (b(a_1), b(a_2), \ldots, b(a_{M^{av}}))$ with its first element corresponding to the threshold of the auction with the least number of bidders, and so on. If the agent selects n auctions to participate in, it will submit n non-zero bids out of a choice of M^{av} assuming the agent has a preference for auctions with fewer bidders. The expected utility is calculated for $k \leq n \leq M^{av}$. The number of auctions to participate in,

n_{opt}, that gives the highest expected utility is then identified:

$$b_n = (\underbrace{b(a_1), b(a_2), \ldots, b(a_n)}_{n}, \underbrace{0, \ldots, 0}_{M^{av}-n})$$

$$n_{opt} = \arg\max_{n \geq \hat{k}_T} E(U|b_n) \quad (4)$$

Now, the agent could just decide to participate in the n_{opt} auctions with the least number of bidders. However, this is not likely to be a good enough strategy because the simplified model ignores the differences between auction formats and discards a lot of useful information obtained after the auction starts (such as the current price of the English and Dutch auctions). On the other hand, with the exception of the case when M is small, it is impractical to calculate the exact utility as it requires the probability calculation to be repeated for each of the 2^M possible winning combinations. In contrast, if we assume the outcomes of the auctions are independent, the expected number of wins can easily be calculated as the sum of the winning probabilities of the individual auctions. By maintaining the expected number of wins at a sufficiently low level, we can mitigate the risk of accidentally purchasing too many items even if we place more than $\hat{k}_T$ bids.

We now turn to the issue of finding the optimal expected number of wins. Assuming the agent bids in n_{opt} auctions with threshold $b_{opt} = (b(a_1), b(a_2), \ldots, b(a_{n_{opt}}), 0, \ldots, 0)$, the expected number of wins that gives highest utility[5] can be calculated as:

$$\hat{x}_{opt} = E(X|b_{opt}) = \sum_{i=1}^{n_{opt}} p_a(b(a_i)) + x_T$$

Finally, we can apply a knapsack algorithm to find the combination of auctions that minimises the payment, while keeping the predicted number of winning items less than or equal to the optimal number of winning items $\hat{x}_{opt}$. Generally speaking, the goal of the knapsack algorithm is to maximise the combined value of items packed into a bag, providing the size constraint has not been violated. Here the optimal winning number $\hat{x}_{opt}$ is taken as the weight constraint and each auction in the subset A^{av} is a candidate item available for selection. For auction a, the weight is the chance of winning $p_a(b_a)$ if bidding with the threshold $b(a)$ found by the heuristics in Section 2.1. The value for participating in auction a is defined as minus the expected payment, so that maximising value minimises payment. The knapsack algorithm sets the bidding thresholds to zero to indicate the auctions that the agent should not participate in. The best selection of auctions to participate in is re-evaluated at each time step.

The expected payment is given by (2), but, for current English (or Dutch) auctions, the expected payment is set to the current price x_t plus (or minus for Dutch auction) the minimum bid increment. It is because the current price update at each round presents a prime opportunity for the agent to hunt for bargains. For example, if there is a more favourable auction later in which the agent is planning to bid up to $b(L)$ with an expected payment of $\pi(L)$, there is no harm in bidding lower than $b(L)$ for the current English or Dutch auction. It may get an even better deal by chance.

[5] It is important to note that the expected number of wins $E(X)$ is different from the upper demand limit k. For example, the expected number of wins can be much smaller than the upper demand limit if the buyer has a low valuation.

3 Empirical Evaluation

This section evaluates the heuristics we have devised. We have introduced two threshold heuristics – single auction dominant (DOM) and equal threshold (EQT), and two auction selection heuristics – exhaustive search selection (ES) and knapsack utility approximation (KS). In particular, we are interested in evaluating how the combined heuristics (DOM-ES, DOM-KS, EQT-ES, EQT-KS) perform. As benchmark, we use a random strategy (RND) that chooses k auctions at the start then bids locally in them and a greedy strategy (GRD) that participates in the k_T auctions with the least number of bidders at time T. Both RND and GRD place bids following the dominant strategy for a single auction. In our experiment, we consider two scenarios: (1) simultaneous SPSB auctions, (2) unrestricted settings, in which the number of bidders, the auction types, and the starting and closing times are all randomised. The first scenario is considered because simultaneous SPSB auctions require only one round of decision. Thus this simplified setting allows us to search for the optimal bidding strategy using a global maximisation technique[6] and compare that with our heuristics. The second scenario is actually our focus because it is the one that most closely represents reality.

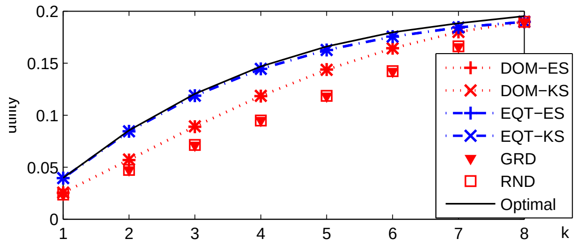

Figure 1. Agent utility for simultaneous SPSB auctions($M = 8, N = 5$).

The first environment evaluates the agent utility for a set of 8 simultaneous SPSB auctions. Specifically, Figure 1 shows the average utility as agent demand k varies from 1 to 8. As can be seen, the combined heuristics using EQT (EQT-ES and EQT-KS) perform better than those using DOM and are within 3% of the optimal. Now, the DOM heuristic adjusts the bidding level in accordance to factors affecting each single auction (i.e. number of bidders and auction format), whereas EQT considers the future item availability from all the remaining auctions. Thus the latter effect dominates, especially for larger M, and so explains the observed improvement when using EQT. Although DOM is less successful, by using our ES or KS decision heuristics, 15-25% of the utility loss can still be recovered when compared with either GRD or RND heuristics.

The second scenario evaluates the unrestricted setting. We examine the received agent utility for the various heuristics with $M = 12$ and $k = 3$. For our results to generalise well to practical scenarios, M should be as large as possible. However, as we shall see later (see Figure 3), ES rapidly becomes computationally intractable as M increases; thus we set $M = 12$ as a compromise value that has a reasonable number of auctions but that is still tractable for ES. According to the ANOVA 2-factor test, all of our heuristics (DOM-KS, DOM-ES, EQT-KS,

[6] Simulated annealing is applied as a global maximisation technique to identify the best bids. However, since it requires up to 1000 seconds to run even for this simplified setting, it is not a practical solution to our problem.

EQT-ES) are significantly better than the benchmark heuristics[7]. Amongst our heuristics, DOM-ES is significantly better than DOM-KS, whereas EQT-ES performs similarly to EQT-KS[8]. On average, our best strategies (EQT-KS, EQT-ES) improve the agent utility by 24% and 27% when compared with the GRD benchmark.

In more detail, the set of auctions can be characterised by the degree of overlap λ and the percentage that are sealed-bid (%SB). To assess how these factors affect our heuristics, the auction sets are classified into one of the four categories, C1 – C4[9], in terms λ and $\%SB$. The agent utility follows the order of EQT-ES $\geq$ EQT-KS $\geq$ DOM-KS $\geq$ DOM-KS $>$ GRD $>$ RND and the ranking is the same for each of the four categories.

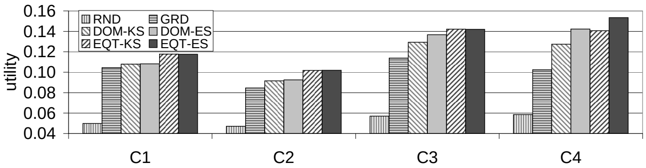

Figure 2. Agent utility for unrestricted auction setting. ($M = 12, k = 3$, $N = [5, 10]$ uniformly distributed)

In addition, our heuristics compare increasingly favourably[10] to the GRD/RND benchmarks when $\lambda \geq 0.5$ and/or %SB < 0.5. This is because many auctions are started or closed at about the same time if λ is large. Thus an agent can hunt for bargains by bidding in many auctions, while not taking too large a risk of accidentally purchasing too many items. When the percentage of English and Dutch auctions is high (i.e. when %SB is low), the agent receives more information in the bidding process. In short, both the careful management of excess procurement risk and improved utilisation of available information require much intelligence, and, thus, the more sophisticated heuristics are beneficial.

The DOM and EQT threshold heuristics require comparatively little computation because they adopt a simplified auction setting. However, the time taken to evaluate the more complex auction selection heuristics is an important practical consideration, especially when many decisions have to be made at the same time (which often occurs near the closing time of a set of simultaneous auctions). To this end, Figure 3 shows the average computation time for the heuristics when running on a modern PC. The ES heuristic is known to have exponential time complexity and it fails to keep up with practical requirements when M is larger than about 12. In contrast, the KS heuristic has been shown to remain acceptable for M up to at least 200 auctions with its pseudo-polynomial time complexity.

To summarise our results, the agent's utility is affected strongly by the threshold heuristic and EQT is the preferred choice because it is better than DOM for each of the four categories of multiple auctions examined in the second scenario. For the auction selection heuristic, the computation time is an extra factor to consider in addition to the received agent utility. Here the ES heuristic can tolerate bad selection of bidding thresholds and its use should be considered if its computation time is deemed acceptable. But for a larger number of auctions ($M > 12$) the KS heuristic is more suitable because the sacrifice in performance when compared with ES is small (DOM-KS) or in some cases insignificant (EQT-KS).

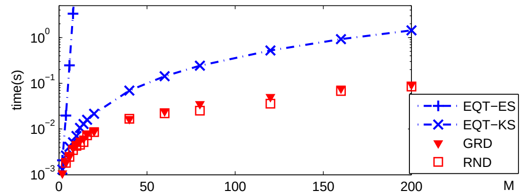

Figure 3. Average computation time for different heuristics.

4 Conclusions

This paper has developed and evaluated a novel two-stage heuristic as a bidding strategy for multiple heterogeneous auctions. In the first phase, a set of bidding thresholds is generated by the threshold heuristics, before a more sophisticated decision heuristic is applied to decide which subset of auctions to participate in. Our empirical evaluation shows that EQT is a better threshold heuristic than DOM and for the auction selection heuristic, KS is an acceptable choice especially for larger numbers of auctions. At present, the heuristics developed assume that there is only one global bidder in the market. However, as a future direction, we intend to examine the Nash equilibrium strategy and analyse the change in best response when more global agents operate in the system.

ACKNOWLEDGEMENTS

This research was undertaken as part of the EPSRC funded project on Market-Based Control (GR/T10664/01). This is a collaborative project involving the Universities of Birmingham, Liverpool and Southampton and BAE Systems, BT and HP.

REFERENCES

[1] P. Anthony and N.R. Jennings, 'Developing a bidding agent for multiple heterogeneous auctions', *ACM Trans on Internet Technology*, **3**, 185–217, (2003).

[2] C. Boutilier, M. Goldszmidt, and B. Sabata, 'Continuous value function approximation for sequential bidding policies', in *Proc. of UAI*, pp. 81–90, (1999).

[3] A. Byde, 'A dynamic programming model for algorithm design in simultaneous auctions', in *Electronic Commerce: Second International Workshop, WELCOM 2001*, volume LNCS 2232, pp. 152–163, (2001).

[4] A. Byde, C. Priest, and N. R. Jennings, 'Decision procedures for multiple auctions', in *AAMAS'02*, pp. 613–620, (2002).

[5] M. Dumas, L. Aldred, G. Governatori, and A.H.M. ter Hofstede, 'Probabilistic automated bidding in multiple auctions', *Journal of Electronic Commerce Research*, **5**, 25–49, (2005).

[6] M. R. Garey and D. S. Johnson, *Computers and Intractability*, W.H. Freeman, 1979.

[7] V. Krishna, *Auction Theory*, Academic Press, 2002.

[8] V. Krishna and R.W. Rosenthal, 'Simultaneous auctions with synergies', *Games and Economic Behavior*, **17**, 1–31, (1996).

[9] P. Milgrom and R.J. Weber, *The Economic Theroy of Auctions*, volume 2, chapter A Theory of Auctions and Competitive Bidding, II, 179–194, Edward Elgar, 2000.

[10] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 2nd edn., 1992.

---

[7] Comparing our worst heuristic (DOM-KS: $\bar{u} = 0.113$) to the best benchmark heuristic (GRD: $\bar{u} = 0.101$), the null hypothesis that they perform the same is rejected at the 95% confidence level (CI) ($F = 46.6 >$ critical value $F_c = 3.86$).

[8] The null hypothesis that DOM-ES ($\bar{u} = 0.120$) and DOM-KS ($\bar{u} = 0.113$) perform the same is rejected ($F = 9.28 > F_c = 3.86$). But, for EQT-ES ($\bar{u} = 0.129$) and DOM-KS ($\bar{u} = 0.125$), the null hypothesis that they perform the same is accepted at the 95% CI ($F = 3.32 < F_c = 3.86$).

[9] C1: $\lambda < 0.5$ and %SB≥ 0.5, C2: $\lambda \geq 0.5$ and %SB≥ 0.5, C3: $\lambda < 0.5$ and %SB< 0.5, C4: $\lambda \geq 0.5$ and %SB< 0.5.

[10] Our worst heuristic, DOM-KS, gains 3, 8, 14 and 24% higher utility than GRD for category C1-C4 respectively.