An Approximate Model for Bidders in Sequential Automated Auctions *

Kumaara Velan and Erol Gelenbe

Intelligent Systems and Networks Group
Electrical and Electronic Engineering Department
Imperial College
London SW7 2BT, UK
{kumaara.velan,e.gelenbe}@imperial.ac.uk

Abstract. In this paper, we propose a probabilistic model to study the interaction of bidder and seller agents in sequential automated auctions. We consider a designated "special bidder" (SB) who arrives at an auction and observes the ongoing activities among a number of bidders and the seller jointly, as a stochastic system that is parameterised by the rate of the bidding and selling events. The auction is modelled as a random process with discretised state-space, where the state-space represents the recently attained price for the good. We distinguish the statistical properties of the SB from that of the others, and isolate the system states that denote the desirable outcomes for the SB. In this manner, we define the measures that are of interest to the SB: its winning probability, the average time it takes to purchase an item, and its expected savings with respect to the maximum payable. For tractability, we introduce an approximately equivalent model that yields convenient and closed form expressions of these measures with minimal loss in the accuracy. We examine the effects of the SB's time to bid, and study how decisions may be taken to balance the trade-offs between its interests.

1 Introduction

Computerised auctions are increasingly prevalent in the Internet-based economy, where a significant part of transactions are conducted by automated buyers and sellers acting on behalf of their human counterparts. Since physical restrictions do not apply in this setting, the buyers can move freely between sites, and hence, have access to a large choice of identical or substitutable goods that may be of interest. In many instances, the buyers will be aware of possible sales of similar items in future [3], probably by a different seller at a different place. This knowledge of future availability of similar items adds flexibility to the buyers

* This research was undertaken as part of the ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Networks) project and is jointly funded by a BAE Systems and EPSRC (Engineering and Physical Research Council) strategic partnership (EP/C548051/1).

decisions; one may choose to restrain oneself from bidding aggressively and pull back under intense competition, in hope of securing a better deal later.

We can also imagine the good on sale being non-exhaustible. For instance, in allocation problems [4] where the resource on offer is not actually sold, but rented out for some period, and is returned to the seller for reuse. This concept may find application in computer networks where the bandwidth is rented to potential users, or in disaster management applications where the resource may be an ambulance, fire engine, or help personnel, that is dispatched to the most needy. Again, the buyers will have to balance their need to consume the resource immediately against their willingness to pay, possibly a higher cost, for the impatience.

In this paper, we generalise the approach in [1], to analyse systems where bidders with distinguishable statistical properties engage in an auction. Suppose a "special bidder" (SB) arrives and observes ongoing activities of the other bidders and the seller combined, as a stochastic system that is parameterised by the rate of the occuring events. The SB's problem is then to adjust its decisions with respect to the rest of the agents so as to maximise its benefits. To this end, we propose a mathematical model to characterise the performance measures that are of interest to the SB, and study its trade-offs.

In the following parts, we first outline the model to be studied and detail the mathematical formalisation in Sect. 2. Then, we introduce an approximately equivalent model that yields tractability, and provide analytical solutions for this case. In Sect. 3, we define the performance measures that are of interest to the SB and the seller, where we show that the approximate solutions allow us to derive these measures in neat closed forms that facilitate straightforward computation. Numerical examples are given in Sect. 4, where we illustrate the model's predictions and evaluate the accuracy of the approximation approach. Finally, we draw conclusions in Sect. 5, and suggest future work.

2 An Ascending Auction with Discretised Increments

We consider an ascending auction, where the possible valuations of successive bids are increasing by one unit price with respect to the previous offer. The analysis follows directly if the increments take on any size, so long as they are discrete and fixed in advance; but this case is not discussed here. Suppose there are $n+1$ participating bidders, one of whom is the SB. Each of the "other" bidders raises the price at some rate β_2, while the SB raises the price at rate $\beta_1 c$; these parameters model the average of the random delay before the price is incremented. Here we let the rest of the bidders collectively share a common statistical behaviour that is distinct from the SB, but it will be easy to extend the model to distinguish behaviours within smaller groups among the other bidders.

We assume that, during the auction, the bidders will not raise the price by any more than what is minimally necessary to stay as the winning bidder; in fact this is the standard approach used by the eBay's proxy bidding agent [5]. The good on sale will have some valuation $v > 0$ that represents the maximum price anyone is willing to pay, and all bidding activities stop once this maximum is achieved. To start with, we let v take a known and fixed valuation, and subsequently, show how the results generalise when it is replaced by a random variable V.

Upon receiving a bid offer, the seller may deliberate on accepting it while waiting for a better deal. This duration is modelled as a random variable with average δ^{-1}. If, when an offer is being thus considered, a higher bid is made, the seller rejects the old bid to consider the new offer. The rejected bidder will rejoin the rest of the bidders, and may attempt bidding again. On the other hand if, at some point during the auction, the waiting time lapses without interruption from new bids, then the auction concludes with a sale to the owner of the highest bid and enters a random "rest" period with an average r^{-1}. This is simply the overhead duration until the next auction can start. For instance, in the case where the item is rented out, instead of sold, this rest time will include the rent duration before it is returned for another round of auction.

Following that, the auction restarts and repeats itself for an infinite number of times, each auction being independent of the previous. All the discussed random variables are modelled via exponential distributions with parameters of their respective rates, and they maintain their statistical properties between auctions.

2.1 The Mathematical Model

Let $\{ t_i : i = 1, 2, \dots \}$, be the sequence of instants when the i–th auction starts, and assume $t_1 = 0$. For any of the auctions, using a fixed v, we model the system as a continuous time Markov process $\{ X(t) : t \geq 0 \}$ [6], which takes valuation from the state-space $\{ 0, O(l), R(l), A(O,l), A(R,l) \}$ where $1 \leq l \leq v$. The valuation $X(t_i + t) = 0$ denotes the case when the auction has restarted for the i–th round and is yet to receive any bids after some time $t \geq 0$, so that $t_{i+1} = \inf\{ t : t > t_i, X(ti + 1) = 0 \}$ defines the restarting instants. Using these reference points, we can describe all the state valuations: The state $X(t_i + t) = RO(l)$, if at time $t_i + t$ for $t > 0$ during the i–th auction $(t_i + t < t_{i+1})$, the price has attained valuation l for $1 \leq l \leq v$, where the l–th bid was placed by the SB (disregarding all previous $l-1$ bids). The state $X(t_i + t) = A(R, l)$ corresponds to the case at time $t_i + t$ for $t > 0$ during the i–th auction $(t_i + t < t_{i+1})$, the seller has accepted the bid placed by the SB at price l for some $1 \leq l \leq v$. The state $X(t_i + t) = A(O, l)$, if at time $t_i + t$ for $t > 0$ during the i–th auction $(t_i + t < t_{i+1})$, the price has attained valuation l for $1 \leq l \leq v$, where the l–th bid was not placed by the SB. The state $X(t_i + t) = O(l)$, if at time $t_i + t$ for $t > 0$ during the i–th auction $(t_i + t < t_{i+1})$, the seller has accepted the highest bid at price l for some $1 \leq l \leq v$, which was not owned by the SB.

When the auction is active, any bidder, with the exception of the winning bidder, may place a bid at its respective rate. Thus, at any price $1 \leq l \leq v - 1$, only n bidders will be active, while at price 0 all $n + 1$ will be involved. No bidders will raise the price beyond valuation v.

The transition from any state $U(l)$ to the corresponding state $A(U, l)$, where $U = O, R$ and $1 \leq l \leq v$, marks the event of a sale at price l, and occurs at rate

δ; while the transition from any state $A(U,l)$, where $U = O, R$ and $1 \le l \le v$, to state 0 denotes the event of an auction restart, and occurs at rate r.

Let $P(\cdot)$ represent the stationary probability distribution for the Markov chain with state-space $\{\,0, O(l), R(l), A(O,l), A(R,l) : 1 \le l \le v\,\}$. Then these probabilities will satisfy the following system balance equations:

$$
\begin{aligned}
&P(O(1))((n-1)\beta_1 + \beta_2 + \delta) = n\beta_1 P(0), \\
&P(R(1))(n\beta_1 + \delta) = \beta_2 P(0), \\
&P(O(l))((n-1)\beta_1 + \beta_2 + \delta) = (n-1)\beta_1 P(O(l-1)) \\
&\qquad\qquad\qquad\qquad\qquad\qquad + n\beta_1 P(R(l-1)), \quad 2 \le l \le v-1, \\
&P(R(l))(n\beta_1 + \delta) = \beta_2 P(O(l-1)), \quad 2 \le l \le v-1, \\
&P(O(v))\delta = (n-1)\beta_1 P(O(v-1)) + n\beta_1 P(R(v-1)), \\
&P(R(v))\delta = \beta_2 P(O(v-1)), \\
&P(A(O,l))r = \delta P(O(l)), \quad 1 \le l \le v, \\
&P(A(R,l))r = \delta P(R(l)), \quad 1 \le l \le v, \\
&1 = P(0) + \sum_{U=O,R} \sum_{l=1}^{v} \big[P(U(l)) + P(A(U,l))\big].
\end{aligned}
\tag{1}
$$

Solving this model will result in closed form, but lengthy and unwieldy expressions of the probabilities of interest. Instead, in pursuit of an analytically tractable model, we will introduce an approximation that will simplify the problem.

2.2 Approximate Model for a Large Number of Bidders

When the number of bidders (other than the SB) is very large, i.e. $n >> 1$, the model is considerably simplified. Starting from (1) we can write:

$$
\begin{aligned}
&a = \frac{n\beta_1}{(n-1)\beta_1 + \beta_2 + \delta} \approx \frac{n\beta_1}{n\beta_1 + \beta_2 + \delta}\,, \qquad b = \frac{\beta_2}{n\beta_1 + \delta}\,, \qquad \text{and} \\
&P(l) = P(O(l)) + P(R(l)),
\end{aligned}
$$

so that, the stationary probabilities related to the other bidders become

$$
\begin{aligned}
&P(O(1)) = aP(0), \\
&P(O(l)) = aP(l-1), \quad 2 \le l \le v-1, \\
&P(O(v))\delta = n\beta_1 P(v-1), \\
&P(A(O,l))r = \delta P(O(l)), \quad 1 \le l \le v,
\end{aligned}
\tag{2}
$$

and the stationary probabilities related to the SB become

$$\begin{aligned}
&P(R(1)) = bP(0), \\
&P(R(l)) = bP(O(l-1)), \quad 2 \le l \le v-1, \\
&P(R(v))\delta = \beta_2 P(O(v-1)), \\
&P(A(R,l))r = \delta P(R(l)), \quad 1 \le l \le v.
\end{aligned} \tag{3}$$

Writing down the probability for $P(0)$, and from the law of total probability, we have

$$1 = P(0) + \sum_{U=0,R} \sum_{l=1}^{v} \left[P(U(l)) + P(A(U,l)) \right].$$

To proceed, we first define the sequences

$$H(l) = \begin{cases} a, & l = 1; \\ a\left[H(l-1) + G(l-1)\right], & 2 \le l \le v-1; \\ \dfrac{n\beta_1}{\delta}\left[H(l-1) + G(l-1)\right], & l = v, \end{cases}$$

and

$$G(l) = \begin{cases} b, & l = 1; \\ bH(l-1), & 2 \le l \le v-1; \\ \dfrac{\beta_2}{\delta} H(l-1), & l = v, \end{cases}$$

so that the stationary probabilities can be written as

$$\begin{aligned}
P(R(l)) &= G(l)P(0), \\
P(O(l)) &= H(l)P(0), \\
P(A(R,l)) &= \frac{\delta}{r} G(l)P(0), \\
P(A(O,l)) &= \frac{\delta}{r} H(l)P(0).
\end{aligned} \tag{4}$$

Then, after some algebraic manipulation, we have the probability that the auction has started and is waiting for bids:

$$P(0) = \frac{\delta r}{(\beta_2 + n\beta_1)(\delta + r) + \delta r}\,. \tag{5}$$

Now, using a simple substitution of $G(l)$ by its valuation defined as a function of $H(l-1)$, we obtain the roots for the recurrent sequence $H(l)$:

$$R_{1,2} = \frac{1}{2}\left[\ a \pm \sqrt{a^2 + 4ab}\ \right].$$

As a consequence, we can write $H(l)$ in a closed form solution:

$$H(l) = \frac{R_1^{l+1} - R_2^{l+1}}{R_1 - R_2}\,, \quad 1 \leq l \leq v-1, \tag{6}$$

and at the boundary $l = v$, the solution involves a different set of coefficients:

$$H(v) = \left[\frac{R_1^v - R_2^v + b(R_1^{v-1} - R_2^{v-1})}{R_1 - R_2}\right] \frac{n\beta_1}{\delta}\,. \tag{7}$$

Since $G(l)$ is defined as a function of $H(l-1)$, its solutions are easily derived from the above:

$$\begin{aligned} G(1) &= b, \\ G(l) &= \left[\frac{R_1^l - R_2^l}{R_1 - R_2}\right] b, \quad 2 \leq l \leq v-1, \\ G(v) &= \left[\frac{R_1^v - R_2^v + b(R_1^{v-1} - R_2^{v-1})}{R_1 - R_2}\right] \frac{n\beta_1}{\delta}\,. \end{aligned} \tag{8}$$

We know that by the truth of the inequalities $a, b > 0$, $a^2 + 4ab > 0$, and $R_1 - R_2 = \sqrt{a^2 + 4ab} > 0$, the given solutions exist and take real valuations. In addition, when the maximum valuation v is only known in terms of some general probability distribution $p(v) = \mathrm{Prob}[V = v]$, the solutions are directly computed as expectations with respect to the random variable.

3 Performance Measures of Interest

In this section, we define and obtain the performance measures of interest, both for the SB and the seller. First we will derive the average duration of an auction, τ. The probability $P(0)$ can be interpreted as the ratio of the average time the process spends in state 0 to the average time between two consecutive reentries to that state. The first is simply the average time the auction waits for the first bid after having restarted, i.e. $[\beta_2 + n\beta_1]^{-1}$, and the latter is τ. Thus, using (5) we get

$$\begin{aligned} P(0) &= \frac{\text{Average time spent in state 0}}{\tau}\,, \\ \tau &= \frac{\delta r + (\delta + r)(\beta_2 + n\beta_1)}{\delta r(\beta_2 + n\beta_1)}\,. \end{aligned} \tag{9}$$

In any auction, the *probability that the SB wins the item, as opposed to any one of the other bidders*, represented by π, is the conditional probability that the auction ends in a state $A(R, l)$ for some $1 \leq l \leq v$, given that a sale is made. The latter is the probability of any event from the set $\{A(U, l) : U = O, R \text{ and } 1 \leq l \leq v\}$ occurring. Thus

$$
\begin{aligned}
\pi &= \frac{\sum_{l=1}^{v} P(A(R,l))}{\sum_{l=1}^{v}[P(A(R,l)) + P(A(O,l))]} \\
&= \left[\sum_{l=1}^{v} P(A(R,l))\right] \cdot \left[\frac{\delta r + (n\beta_1 + \beta_2)(\delta + r)}{\delta(n\beta_1 + \beta_2)}\right].
\end{aligned} \tag{10}
$$

Suppose we are interested in the *average time that the SB waits until it succeeds in purchasing an item*, which we denote by $\psi(v)$. This is readily derived from the above measures. Note that, due to the auctions being independent and identical processes, the winning probabilities are identical across the auctions. Thus this measure is the average of multiples of τ, probabilistically weighted by the number of auction rounds it takes to win:

$$
\psi(v) = \frac{1}{r\sum_{l=1}^{v} P(A(R,l))}. \tag{11}
$$

The final measure of interest for the SB is the *average savings it generates with respect to the maximum that it is willing to pay, given that it wins the item*, or

$$
\phi(v) = \frac{\sum_{l=1}^{v}(v-l)P(A(R,l))}{\sum_{l=1}^{v} P(A(R,l))}, \tag{12}
$$

which we will refer to as the SB's "expected payoff".

On the other hand, the seller's expected income I from an auction is computed as the average of the closing price l weighted by the probability of a sale at that price, regardless of the identity of the winner:

$$
I = \frac{\sum_{l=1}^{v} l[P(A(R,l)) + P(A(O,l))]}{\sum_{l=1}^{v}[P(A(R,l)) + P(A(O,l))]}. \tag{13}
$$

Over many such auctions, the seller's interest may be to maximise its *income per unit time* $\iota = \frac{I}{\tau}$.

3.1 Approximate Solutions

Using the solutions from (6), (7) and (8), we can readily derive the probability that the SB wins the good:

$$
\begin{aligned}
\sum_{l=1}^{v} P(A(R,l)) = P(0)\Bigg[\frac{\delta b}{r} &+ \frac{\delta b}{r(R_1 - R_2)}\left[\frac{R_1^v - R_1^2}{R_1 - 1} - \frac{R_2^v - R_2^2}{R_2 - 1}\right] \\
&+ \frac{\beta_2}{r}\left[\frac{R_1^v - R_2^v}{R_1 - R_2}\right]\Bigg].
\end{aligned} \tag{14}
$$

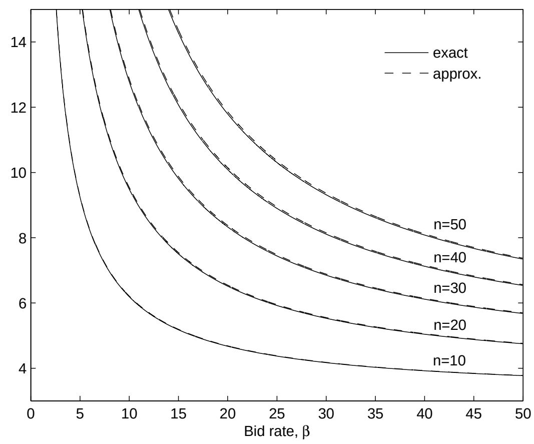

Fig. 1. The SB's expected time to win against bidding rate β_2, comparing exact solution with approximation result for various number of bidders. Other parameters are $\delta = 0.5$, $r = 1$, $\beta_1 = 1$, and $V \sim U(80, 100)$.

Similarly, the expected price paid by the SB when it makes the purchase is

$$\sum_{l=1}^{v} lP(A(R,l)) = P(0)\Bigg[\frac{\delta b}{r} + \frac{v\beta_2}{r}\Big[\frac{R_1^v - R_2^v}{R_1 - R_2}\Big] + \frac{\delta b}{r}\frac{1}{R_1 - R_2}\Big[\frac{(v-1)R_1^{v+1} - vR_1^v - R_1^3 + 2R_1^2}{(R_1-1)^2} - \frac{(v-1)R_2^{v+1} - vR_2^v - R_2^3 + 2R_2^2}{(R_2-1)^2}\Big]\Bigg]. \tag{15}$$

The ability to express these two quantities in a neat and closed form is an important gain and justifies our approximation approach; this is not possible using the exact model (1). We can now express the quantities $\psi(v)$ and $\phi(v)$ in a lengthy but convenient form that can be directly computed using formulas (11) and (12), respectively.

4 Numerical Examples

We will present the model's predictions and assess the accuracy of the approximation approach via some numerical examples. We let all parameters of the system be fixed, where their valuations capture the stochastic nature of the activities, and examine the impact of the SB's bidding rate β_2 on the quantities $\psi(v)$ and $\phi(v)$.

In the following, we compare the results that are obtained by solving the original model (1) numerically and exactly, with that evaluated from the approximate analytical solutions (14) and (15). Figure 1 confirms, as intuition

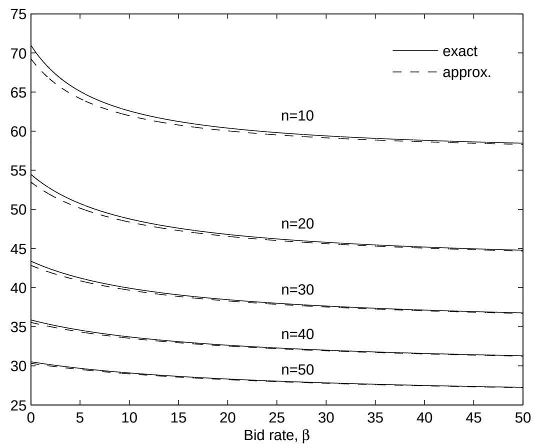

Fig. 2. The SB's expected payoff against bidding rate β_2, comparing exact solution with approximation result for various number of bidders. Other parameters are $\delta = 0.5$, $r = 1$, $\beta_1 = 1$, and $V \sim U(80, 100)$.

suggests, that by bidding at a high rate (the $x - axis$), the SB can expect to win the item in a shorter time (the $y - axis$). Also, the slope $d\psi/d\beta_2 < 0$ for all β_2, and it increases with β_2. The increase is more rapid for smaller valuations of n, and accentuates the effect of diminishing returns for the SB.

The quantity $\phi(v)$ is shown in Fig. 2. The intensity of the competition that the SB faces, represented by different valuations of n, determines the levels of payoffs that the SB may expect. Within these levels, however, increasing the SB's rate of bidding adversely affects its expected payoff given that it wins. In other words, when n is fixed, $\phi(v)$ drops with β_2; this drop is more significant for smaller valuations of n.

Notably, in all the results, the approximations are close to the exact solutions. In all cases, the expected time to win is overestimated, while the expected payoff is underestimated, by the approximate solutions. This can be intuitively explained. In reality, the cumulative bidding rate of the other bidders is conditional upon the event that the highest bidder is the SB, in which case it takes a multiplicative factor of n; or the event that the highest bidder is one of themselves, in which case it takes a multiplicative factor of $n-1$, because the winning bidder refrains from further participation. By approximating, we have "lumped" these events together, taking the cumulative rate as a factor of n. Therefore we have modelled a process more competitive than the actual, and, consequently, the results will be less favourable to the SB: the expected payoff is lower, and the expected time to win is higher than the actual case.

Interestingly, there is an observable trade-off here. Increasing the bid rate will result in the SB buying the item quickly at the expense of some payoff. If the SB is pressed for time and needs to purchase the item quickly, this may be a good option. On the other hand, for the "bargain hunters", buying the item at the best possible price may be a priority, and the time spent, perhaps, is of less importance. Thus, clearly, the SB will have to balance its willingness to forgo some payoff against the prospect of winning the item sooner.

5 Conclusions

In this paper, we propose a discrete state-space continuous time probability model to study auctions that proceed with unit increments, where the state-space represents the attained price. We distinguish the SB from the others, and isolate the states corresponding to the desirable outcomes for the SB, so that we can compute the measures of interest for the SB. These measures are: the probability that the SB wins rather than one of the other bidders, the average time the SB waits to be able to purchase an item, and its expected savings with respect to the maximum payable. For tractability, we introduce an approximation that allows us to conveniently express these measures in closed form, with minimal loss in the accuracy.

With this model, we can quantitatively relate the speed at which the SB raises the price to its measures of interest, in the presence of varying degrees of competition from the other bidders and the seller's urgency to make a sale. There are a few interesting extensions and applications of this model that we will pursue in our future work. In particular, we can characterise the time constraints faced by the bidders, and examine optimal decisions under this formulation. Also, the possibility and the benefits of bidders moving between different auction sites can be considered.

References

1. Gelenbe, E.: Analysis of automated auctions. In Levi, A., Savas, E., Yenigün, H., Balcisoy, S., Saygin, Y., eds.: ISCIS 2006. LNCS. Volume 4263., Heidelberg, Springer (2006) 1–12
2. Gelenbe, E.: Analysis of single and networked auctions. Accepted for publication ACM Trans. Internet Technol. (2008)
3. Zeithammer, R.: Forward-looking bidding in online auctions. J. Marketing Res. **43**(3) (2006) 462–476
4. Gagliano, R.A., Fraser, M.D., Schaefer, M.E.: Auction allocation of computing resources. Commun. ACM **38**(6) (1995) 88–102
5. eBay: http://ebay.co.uk
6. Medhi, J.: Stochastic Processes. J. Wiley & Sons. (1994)
7. Gelenbe, E., Mitrani, I.: Analysis and Synthesis of Computer Systems. Academic Press, New York and London (1980)
8. David, E., Rogers, A., Jennings, N.R., Schiff, J., Kraus, S., Rothkopf, M.H.: Optimal design of english auctions with discrete bid levels. ACM Trans. Internet Technol. **7**(2) (2007) Article 12 (34 pages).
9. Krishna, V.: Auction Theory. 1st. edn. Academic Press, San Diego (2002)
10. Rothkopf, M.H., Harstad, R.M.: Modeling competitive bidding: A critical essay. Manage. Sci. **40**(3) (1994) 364–384