Competitive Bidding in a Certain Class of Auctions

Mathias Johansson[,†]

Signals and Systems Group, Uppsala University, Sweden,
e-mail: mj@signal.uu.se
Dirac Research AB, Sweden, e-mail: mj@dirac.se

Abstract. We consider the problem of determining the amount to bid in a certain type of auctions in which customers submit one sealed bid. The bid reflects the price a customer is willing to pay for one unit of the offered goods. The auction is repeated and at each auction each customer requests a certain amount of goods, an amount that we call the capacity of the customer and that varies among customers and over time. At each auction, only the customer with the largest bid-capacity product obtains any goods. The price paid by the winner equals his/her bid-capacity product, and the amount of goods obtained in return equals the winner's capacity. The auction is repeated many times, with only limited information concerning winning bid-capacity products being announced to the customers. This situation is motivated in for example wireless communication networks in which a possible way of obtaining a desired service level is to use dynamic pricing and competitive bidding. In this application, the capacity is typically uncertain when the bid is made. We derive bidding rules and loss functions for a few typical service requirements.

Keywords: Competitive bidding, Bayesian inference, Maximum entropy, Auctions

1. INTRODUCTION

We here consider a bidding situation in which customers compete for a resource which can only be used by one customer at a time. The resource carries a certain utility, the *capacity* of the resource, which varies over time and is different for each competitor. For instance, the capacity may in a mobile telecommunications network be the time-varying data rate over the communications channel.

Each competitor submits one sealed bid. After all bids have been collected, a winner is announced who gets access to the resource for a certain time period and thereby receives goods according to that customer's capacity. For the next period, a new auction is carried out again under similar circumstances.

If the winning bid was q and the capacity of the winning customer was c, the winning customer pays qc monetary units, i.e. q is the price per unit utility. The auctioneer's income for each auction is thus qc, and the winning customer is the one with maximum price-capacity product qc.

Our problem set-up is the following:

- Different bidders u may have different capacities c_u
- Each bidder u reports its own capacity c_u to the auctioneer along with its bid q_u. Both values are hidden for other customers.
- Although all information reported to the auctioneer is sealed, a bidder obtains some

implicit information regarding other bidders' capacities and bids from how many times the bidder wins the auction. The bidder does however not know who wins an auction that is not won by the bidder, nor, in that case, the winning price-capacity product.

- The auctioneer knows all bidders' capacities and bids.

The question we seek to answer is then: What is the best bid that a customer can make? Clearly, the answer depends on the customer's need for capacity, and – having established a loss function describing this – any information at hand that can assist in reaching a decision. This type of problem was considered by Friedman [1] in 1956, and a similar strategy as the one we will use here was suggested. Friedman considers the objective of bidding for maximum expected profit in a scenario where a government agency invites a large number of companies in the same industry to bid for contracts. Friedman notes that "the difficulty in determining the expected profit lies in determining ... the probability of winning as a function of the amount bid". He suggests the use of histograms of bids from old auctions, assuming that all previous bids are made public after an auction. In our scenario, we do not assume knowledge of all previous bids. In many auctions, only the winning bids are announced and then Friedman's method fails to determine a probability distribution for the other customers' bids. From our present understanding of probability theory as logic, however, the solution is straightforward. As always, a probability distribution should not reflect old frequencies but carry all information, and lack thereof, that we actually have concerning the unknown event. In our specific scenario, the information we assume to be in possession of will lead to a maximum entropy problem. In general, additional information should be processed through Bayes' rule.

The problem formulation has a motivation in mobile communications, where it has been suggested that one way of obtaining different service levels is to use dynamic pricing and competitive bidding. There, the capacity is the bit rate that a user can receive or transmit data with. In such a network, prices would decrease when the network is under-utilized or the user is near a base station, and *vice versa*. Before proceeding to the technical derivations, let us first take a concrete example of the auctioning procedure.

Example: Two users compete for access to a wireless communications channel. In any time slot (on the order of milliseconds), only one user may access the channel. To maximize revenues, the base station transmits data to the user which pays the highest total price for access. At a certain time slot t, user 1 can receive $c_1(t) = 100$ bits of data and decides to bid $q = 0.1$ per bit, i.e. the total bid-capacity product is 10, whereas user 2 has a capacity $c_2(t) = 80$ and bids $q_2(t) = 0.2$ per bit, giving the bid-capacity product 16. The bids and capacities are transmitted to the base station on a separate control channel (the bids would typically not be updated over a number of consecutive periods, thereby alleviating the need for a high-rate feedback channel). Neither user knows the other user's bid or capacity. The base station receives the information from both users and awards the next slot to user 2 who has the highest bid-capacity product. Finally, on a regular basis the base station broadcasts some aggregate statistics of the winning bid-capacity products, which will be described in later parts of the paper. The users adjust their bids according to this information and the process is repeated.

2. A BAYESIAN STRATEGY FOR COMPETITIVE BIDDING

Our approach is to minimize the expected loss conditional on the limited information I available to the customer. Let a particular customer u's probability that he or she will have the largest bid-capacity product of all customers be denoted by $P(u \mid I)$. Then $P(u \mid I)$ is equal to the probability that the customer v with the largest bid-capacity product of all other customers has a lower bid-capacity product than customer u. Let q_v denote the bid of v, c_v the corresponding capacity, and $y = q_v c_v$ the largest bid-capacity product among all customers except u. We can then find the probability that u wins as follows: first determine the probability that $y < c_u q_u$ assuming knowledge of c_u, i.e. $\int_0^{c_u q_u} P(y \mid c_u I) dy$. Then multiply this with the probability distribution for c_u given I to obtain the joint probability for c_u and $y < c_u q_u$. Integrating the result over all possible capacities c_u, we have

$$P(u \mid I) = \int P(c_u \mid I) \int_0^{c_u q_u} P(y \mid c_u I) dy dc_u \; . \tag{1}$$

In order to determine this probability distribution we must first find the probability distribution for c_u and that for y. We will consider a general case in which the capacities c_u may be unknown in advance, as that is often the case in mobile communications. If the capacity is already known the solution simplifies straightforwardly.

Assume that there are K different possible capacities c_k. We suppose further that each customer stores the number of time slots that each capacity c_k could be used during a recent time window. If nothing else than these numbers are known, the probability that the customer's capacity will be c_k is then the expected frequency with which that capacity will be used. According to Laplace's rule of succession, see [2], Chapter 18, the probability for having the capacity c_k is

$$P(c_k \mid I) = \frac{n_k + 1}{N + K} \; , \tag{2}$$

where n_k is the number of time slots over the last N records that capacity c_k (but not higher) could be attained.

Now, the distribution $P(y \mid I)$ of the other customers' best price-capacity product depends heavily on the information I that customer u possesses. We will here assume that the auctioneer periodically broadcasts the expected winning price-capacity product for the coming period along with a measure of the prediction uncertainty. The simplest such scheme would consist of recording the average of the most recent winning price-capacity products and its variance. More advanced schemes include determining a model for the time evolution of price-capacity products. Here, we will assume that an expectation is available along with a variance for the prediction. These two quantities are broadcast to all users at regular intervals.

With no other knowledge than the mean and the variance of a variable, the least biased probability distribution according to the maximum entropy principle is Gaussian. Thus, we shall take

$$P(y \mid I) = \frac{1}{\sqrt{2\pi}\sigma_y} \exp\left\{ -\frac{1}{2\sigma_y^2}(y - \mu_y)^2 \right\} \; , \tag{3}$$

with μ_y denoting the expectation of y, and σ_y^2 the variance of the distribution. Here, by not truncating the distribution at zero we have assumed that the variance of the distribution is not too large compared to the mean, so that the tail of the distribution below $y = 0$ is negligible. It should also be pointed out that we are told the mean and the variance of *all* winning price-capacity products, which includes those times when customer u won. However, we should actually determine a distribution for the winning price-capacity products of all customers *except* u. Below, we discuss how to adjust the mean and the variance to subtract out the contributions from customer u. It is however not clear in general that this distribution, having excluded one of the components, should also be Gaussian. We have good reason to use a Gaussian distribution if there are many bidders with independently and symmetrically varying price-capacity products around some mean. Now, the bids are not logically independent since all customers base their decisions on partly the same information. On the other hand, the capacity variations will often, for instance in the mobile communications scenario described above, be independent among customers, which to some extent will have a "randomizing" effect on the price-capacity products. Nonetheless, we may argue that a correlated distribution might be a better model. We will leave this alternative as a topic for future research, and here continue to work with the Gaussian model.

Inserting (2) and (3) into (1) (replacing the integral over c_u with a sum, reflecting that c_u is discrete) we obtain

$$\begin{aligned} P(u \mid I) &= \sum_{k=1}^{K} \frac{n_k+1}{N+K} \int_{-\infty}^{q_u c_k} \frac{1}{\sqrt{2\pi}\sigma_y} \exp\left\{ -\frac{(y-\mu_y)^2}{2\sigma_y^2} \right\} dy \\ &= \sum_{k=1}^{K} \frac{n_k+1}{N+K} \times \frac{1}{2} \mathrm{erfc}\left(\frac{\mu_y - q_u c_k}{\sqrt{2}\sigma_y} \right) \end{aligned} \tag{4}$$

where $\mathrm{erfc}(x) = 1 - \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt$ is the complementary error function.

2.1. Typical loss functions

Different customers may have different service demands. We here propose a number of loss functions that are intended to reflect typical requirements. The loss functions would moreover often be supplemented by a constraint on the maximum allowed bid.

2.1.1. Constant demand

A customer u wishing to obtain a certain amount ϕ_u of goods over the coming N time slots should use

$$L(q_u, x_u(q_u)) = |x_u(q_u) - \phi_u| \ , \tag{5}$$

where $x_u(q_u)$ is the actual amount of goods that the user will obtain for q_u monetary units.

2.1.2. *Price-performance ratio*

A customer u may wish to increase his bid if that bid would result in a significantly increased amount of delivered goods. In some sense, the price-performance ratio should be optimized. A possible formalization is the following: A price increase of 1 unit is acceptable given that the amount of goods obtained then increases by at least a factor a. Then the following loss function should be used.

$$L(q_u, x_u(q_u)) = \frac{a^{q_u}}{\max(x_u(q_u), b)} \;, \tag{6}$$

where $x_u(q_u)$ is the actual amount of goods that the customer will obtain for q_u monetary units. If $x_u(q_u) > b$ then an increased bid, $q_u \rightarrow q_u + 1$ will result in a lower loss if and only if $x_u(q_u+1) > a x_u(q_u)$, because then we obtain

$$L(q_u+1, x_u(q_u+1)) = \frac{a^{q_u+1}}{x_u(q_u+1)} < \frac{a^{q_n}}{x_u(q_u)} = L(q_u, x_u(q_u)) \;. \tag{7}$$

The formulation (6) also includes a minimum acceptable delivery size; if the user is to pay more than 0 monetary units per bit then the throughput must satisfy $x_u(q_u)/a^{q_u} > b$.

For example, if the customer requires at least an amount of 50 units per time slot, and if a price raise of 1 unit is acceptable only if the obtained goods then double, the loss function is $2^{q_u}/\max(x_u(q_u), 50)$.

2.2. Making the decision – expectations and computations

The expected throughput $\langle x_u(q_u) \rangle$ per time slot as a function of the bid q_u is

$$\langle x_u(q_u) \rangle = \sum_{k=1}^{K} c_k \times \frac{n_k+1}{N+K} \times \frac{1}{2} \mathrm{erfc}\left(\frac{\mu_y - q_u c_k}{\sqrt{2}\sigma_y} \right) \;. \tag{8}$$

Similarly, the expected loss using the loss function (5) is

$$\langle L(q_u) \rangle = \sum_{k=1}^{K} |c_k - \phi_u| \times \frac{n_k+1}{N+K} \times \frac{1}{2} \mathrm{erfc}\left(\frac{\mu_y - q_u c_k}{\sqrt{2}\sigma_y} \right) \;. \tag{9}$$

The expected loss using (6) involves determining the expectation of $1/x_u$ for the Gaussian-distributed uncertainty of x_u, an expectation which is not available in closed form. We shall instead use the expected value of x_u directly in (6), thus obtaining a suboptimal solution that does not fully account for our actual uncertainty in making the bid. The estimated loss $\hat{L}(q_u)$ is then

$$\hat{L}(q_u) = \frac{a^{q_u}}{\max(\langle x_u(q_u) \rangle, b)} \;, \tag{10}$$

where $\langle x_u(q_u) \rangle$ is defined in (8).

Recall that y is the winning price-capacity product of all customers *except* customer u. In calculating the best bid, a customer must therefore adjust the variance and the mean of the distribution for the best price-capacity product since these quantities are broadcast and based on all customers. These adjustments are quite difficult to carry out for a customer who has been awarded all or almost all resources over the last period. Usually, however, we would expect that there are many different customers who obtain at least some goods, and then the following adjustments may be used.

The average μ_y is estimated from the broadcast value μ_w (the average of the winning bids) by

$$\mu_y = \frac{l\mu_w - q_u(t-1)x_u(t-1)}{l - l_u} \tag{11}$$

where l is the number of time slots between consecutive price updates, l_u is the number of time slots that customer u won, and $q_u(t-1)x_u(t-1)$ is the sum of customer u's price-capacity products for the l_u time slots that were won by customer u in the previous period of l slots.

Similarly, the variance is estimated by

$$\sigma_y^2 = \frac{l\sigma_w^2 - l_u\sigma_u^2(t-1)}{l - l_u} \tag{12}$$

where σ_u^2 is the sample variance for the price-capacity product of customer u in the slots that this customer won.

In order to compute the minimum of either of the two expected loss expressions (9) and (10) a numerical one-dimensional search is carried out using e.g. the Nelder-Mead simplex algorithm [3].

3. EXAMPLES

We now consider the performance of the scheme outlined in this paper based on simulations of the mobile communications scenario mentioned in the Introduction. Assume one transmitting base station and $U = 4$ competing users. With a periodicity of $n = 20$ time slots, i.e. every 20th auction, each mobile user updates its bid and submits it to the base station. An upper limit on the bid, $q_u \leq 5$, is also assumed. There are $K = 4$ different transmission rates, i.e. capacities, and each user determines and tells the base station the capacity that can be used in the next time slot based on channel measurements. The base station then transmits exclusively in each time slot to the user with the highest price-capacity product. There are four different capacities (bits per time slot),

$$c_1 = 0 \quad c_2 = 74 \quad c_3 = 92 \quad c_4 = 106. \tag{13}$$

In the simulation, the actual capacities for each user in each time slot are drawn from a random number generator, with equal but independent statistics (quantized Gaussian with mean 80 and standard deviation 20) for each user. For more details on how to determine proper capacity levels in a network, please see Chapter 6 of [4]. The capacity probabilities (2) are updated continuously as more data becomes available.

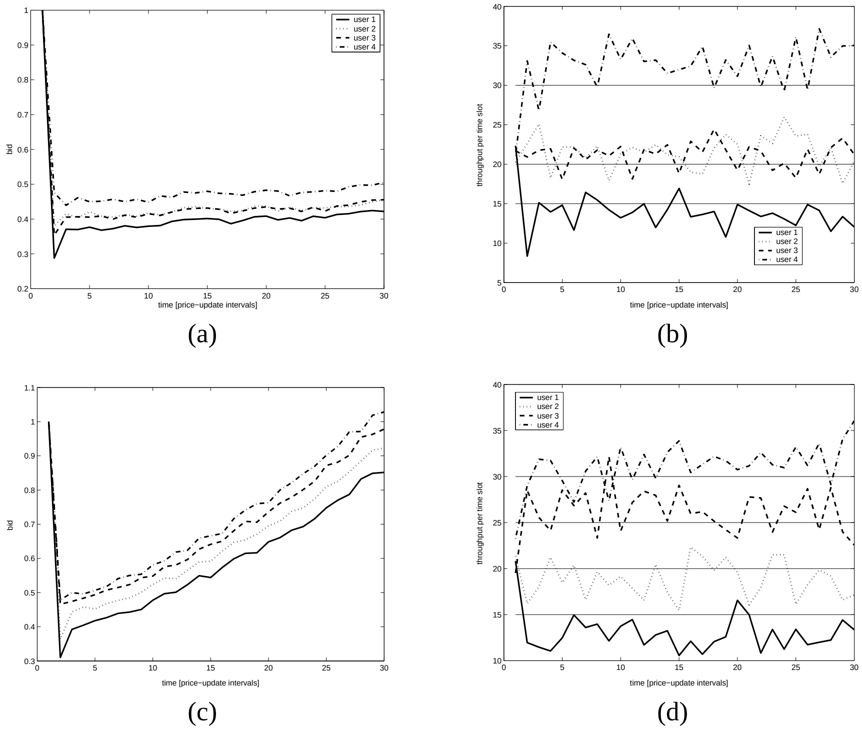

FIGURE FIGURE 1. (a) The evolution of the bids for the four users with desired rates 15, 20, 20 and 30 respectively. (b) The obtained throughput per time slot for the four users. (c) The evolution of the bids for the four users with desired rates 15, 20, 25 and 30 respectively. (d) The obtained throughput per time slot for the four users.

We first consider a case where all four users have a desired rate per time slot according to

$$\phi_1 = 15 \quad \phi_2 = 20 \quad \phi_3 = 20 \quad \phi_4 = 30 \tag{14}$$

and attempt to minimize (9). Figures 1 (a) and (b) show the resulting bids and obtained throughput per time slot from this test in a simulation lasting for 600 repeated auctions (i.e. 30 price-update intervals). The plotted results are averages from 25 simulations. The average obtained rates over the entire simulated period were found to be close to the desired rates:

$$\bar{x}_1 = 14 \quad \bar{x}_2 = 21 \quad \bar{x}_3 = 21 \quad \bar{x}_4 = 33. \tag{15}$$

Under otherwise similar circumstances, Figures 1 (c) and (d) show the bids and the obtained capacities when the desired rate of user 3 was increased to 25 bits per time slot, yielding a more competitive setting. Here, we see that the prices tend to increase because the users have trouble obtaining the desired quality of service. The average obtained capacity per time slot over the entire simulated period now becomes

$$\bar{x}_1 = 13 \quad \bar{x}_2 = 19 \quad \bar{x}_3 = 26 \quad \bar{x}_4 = 31. \tag{16}$$

In a similar setting as the previous one, we now let user 1 minimize the approximate expectation (10) of the price-performance-related loss with $a = 2$ and $b = 8$. Recall that

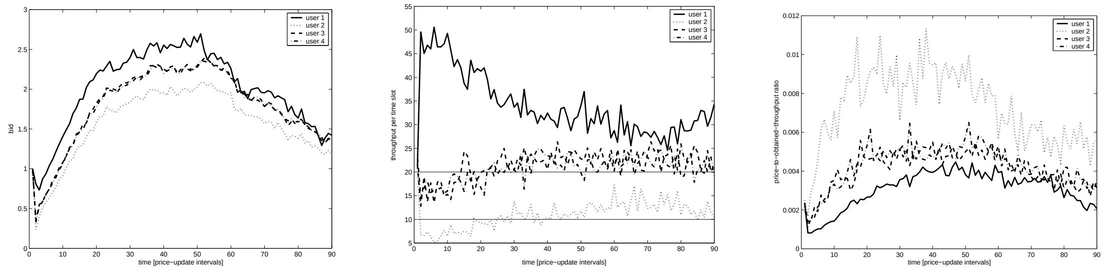

(a) (b) (c)

FIGURE FIGURE 2. (a) The evolution of the bids for the four users with user 1 minimizing the price-performance-related estimated loss (10) and the other users employing (9) with desired rates 10, 20 and 20 respectively. (b) The obtained throughput per time slot for the four users. (c) The evolution of the price-to-performance ratio (the bid divided by the obtained throughput) for the four users.

use of this loss means that a 1-unit price increase is acceptable only if it leads to more than a doubling of the obtained throughput. Only if the throughput becomes more than $2^{q_u} \times 8$ bits is a non-zero bid q_u preferable. Users $2-4$ continue to minimize the expected loss (9) for a desired rate per time slot of

$$\phi_2 = 10 \quad \phi_3 = 20 \quad \phi_4 = 20. \tag{17}$$

In Figures 2 (a), (b), and (c), the bids, obtained throughput and the price-to-obtained-throughput ratio (PTR) q_u/x_u are plotted as a function of time. The results are averages from running a simulation consisting of 1800 auctions 25 times. The average obtained throughput per auction in this case becomes

$$\bar{x}_1 = 34 \quad \bar{x}_2 = 11 \quad \bar{x}_3 = 21 \quad \bar{x}_4 = 21, \tag{18}$$

where we see that users $2-4$ obtain rates corresponding well to their preferences. From Figure 2 (c) we see that user 1 achieves the lowest PTR while the user with the lowest rate requirement has the worst PTR.

All in all, the performance examples show that the bidding strategies seem to function well, but it should be noted that a full analysis of the behavior of the bidding policies is extremely complex and has not been carried out here. The individual bidder, in trying to make a reasonable bid in terms of his/her loss function, bases his/her decision on information which is different for different customers (because the estimates of the other users' best price-capacity products become different for different users depending on the number of wins for that customer). Therefore, the behavior becomes very complex and hard to predict. We need to find better theoretical means for such a deeper analysis.

REFERENCES

1. L. Friedman, *Operations Research*, **4**, 104–112 (1956).
2. E. T. Jaynes, *Probability Theory - The Logic of Science*, Cambridge University Press, 2003.
3. J. A. Nelder, and R. Mead, *Computer Journal*, **7**, 308–313 (1965).
4. M. Johansson, *Resource Allocation under Uncertainty: Applications in Mobile Communications*, Ph.D. thesis, Uppsala University, Signals and Systems (2004).