

A Random Walk Model for Studying Allocation Patterns in Auction-based Resource Allocation^{*}

Manos Dramitinos, George D. Stamoulis, and Costas Courcoubetis

Network Economics and Services Group (N.E.S.),
Department of Informatics, Athens University of Economics and Business
76 Patision Str. Athens, GR 10434, Greece
{mdramit, gstamoul, courcou}@aueb.gr

Abstract. We consider users bidding in a series of multi-unit sealed-bid auctions, aiming at reserving the same amount of units of the resource auctioned, e.g. transmission slots in a wireless network. Each user attains from each successful allocation of resource units an instant marginal utility that depends on his history of resource allocation. The user's bid at each auction equals this marginal utility. We introduce a random walk model for transient analysis of this series of auctions, we study the properties of the resulting user resource allocation patterns and we provide a numerical and experimental evaluation of this model.

Keywords: Auctions, random walk, resource allocation.

1 Introduction

In this paper, we consider users who are participating in a series of consecutive sealed-bid multi-unit auctions, aiming at reserving the same amount of units of the resource (or, in general, of the good) auctioned. Each user attains from each successful allocation of resource units an instant marginal (i.e. additional) utility that depends on the resource allocation pattern. We assume that the user's bid at each auction equals this marginal utility. Hence, the auction's price fluctuations and the users' utility functions greatly affect both the bids submitted and the resulting resource allocation patterns. This is a problem of practical importance in communication networks. A prominent example is bandwidth allocation in UMTS (and other, e.g. GPRS) networks. Indeed, in [1], we study the problem of resource reservation in UMTS networks in which users request services other than telephony that last for long time intervals. Each of these sessions has a fixed target QoS level, which for simplicity we assume that corresponds to a certain bit-rate. The duration of network time-slots over which resource units are allocated is much shorter. We define in [1] an auction-based mechanism achieving nearly consistent reservation of the resources of a UMTS network by the users that value them the most, in order to satisfy the longer time scale QoS requirements

^{*} This research was partly supported by EU-IST-2003-507607 B-BONE and IST-NoE EuroNGI. The authors wish to thank B-BONE and EuroNGI partners for useful discussions on the subject of this paper.

of their service sessions. Thus, due to our mechanism, these users receive service of very good quality at a charge determined by the market. The non-competitive users receive service of very inferior quality for a very short time period at a very low charge. Therefore, our resource reservation mechanism serves a soft CAC. The mechanism is based on a series of Generalized Vickrey Auctions and a set of predefined user utility functions that we introduce [1]. Bidding is performed automatically on behalf of the users on the basis of each user's selection of one of these utility functions and his declaration of a total willingness to pay, and is dependent of the user's history of resource allocation. This approach has also been adopted by other researchers in the field [2] as well as also in the EU-funded IST project B-BONE [4]. In this paper we present a random walk model for transient analysis of this mechanism. Based on this model, we study the impact of the parameters of one of these utility functions in the resource allocation patterns that users attain. We also provide a numerical and experimental evaluation of the model, setting its parameters according to the input from project B-BONE. Finally, we explain how it can help users select parameters of their utility/bidding functions. The analysis of the properties of auctions in dynamic environments by means of mathematical tools has also received attention in the literature [3].

2 The auction mechanism

The problem of UMTS resource allocation to sessions with QoS requirements is very complicated. Indeed, users demand sessions spanning partly overlapping intervals with different durations, which in general are much larger than the time scale t_a of the network frames in which resources are allocated. The approach that we introduced in [1] is to conduct a sequence of auctions, each concerning reservation of bits within one UTRAN frame. Each auction is a sealed-bid Generalized Vickrey Auction (GVA), with bids of the type (p, q) , where q is the quantity of units (bits) sought in the present frame and p is the price proposed for each such unit. These bids are essentially *atomic*, i.e. a winning bid results in the allocation of all the resource units demanded, except for the *cut-off* bidder.

In a realistic case of a UMTS network, it would not be feasible for users to participate in all these auctions. Thus, since the user cannot place his bids on a per auction basis, we define utility functions pertaining to the various services. These functions are provided by the network operator as bidding functions for the user to choose from; they are scaled by the user's i total willingness to pay $U_{s,i}$ for the service s , which is given by the user himself as part of his service request. Then, the network runs the various auctions by bidding on *behalf* of each user. We assume that the user's i utility $u_{s,i}$ from obtaining the service s is the sum of the marginal utilities attained due to each successful allocation; thus, $u_{s,i}(x_i^{(1)}, \dots, x_i^{(K_{s,i})}) = \sum_{t=1}^{K_{s,i}} v_{s,i}^{(t)}(x_i^{(1)}, \dots, x_i^{(t)})$, where $K_{s,i}$ is the number of auctions where user i participates during his service session. Also, for every user the *bid placed by the network at each auction equals the marginal utility* to be attained if his bid is a winning one in the auction. This choice is motivated by the

incentive compatibility property of the Generalized Vickrey Auction, whereby sincere bidding is a dominant strategy; see [1].

In this paper, we restrict attention to one of these predefined utility functions, namely that suitable for users sensitive to service continuity, such as audio and video streaming. Thus, these users prefer the allocation pattern of Fig. 1(a) to that of Fig. 1(b). In order to express this preference, we define the sub-utility function to be $v_{s,i}^{(t)}(x_i^{(1)}, \dots, x_i^{(t)}) = \mathbf{1}(x_i^{(t)} = m_i) \frac{U_{s,i}}{K_{s,i}} \cdot \alpha^{d_i}$, where: $\alpha \in (0, 1]$ is a discount factor; d_i is the distance between the current and the previous slots during which user i achieved reservations; $\mathbf{1}(\cdot)$ denotes the indicator function, which is justified by the fact that bids are atomic. Therefore, *history* of previous allocations influences $v_{s,i}^{(t)}$ through the value of d_i . Hence, when the user fails to reserve resources at some auction, his marginal utility decreases exponentially, in order to express the user's displeasure from the existence of gaps in his resource allocation pattern. On the contrary, as long as a user is allocated resources, his marginal utility equals $\frac{U_{s,i}}{K_{s,i}}$, which is henceforth denoted as u_0 . The utility function considered is suitable for the UMTS Streaming Class which is destined to serve streaming audio and video sessions; m_i pertains to the Maximum Bit-rate parameter of this class.



Fig. 1. Inconsistent resource allocation patterns.

3 The Random Walk model

We consider a user who is participating in a series of consecutive sealed auctions (see Sect. 2), aiming at reserving the same amount of units of the resource auctioned (e.g. bits in a UTRAN frame). Recalling how bidding is performed, it follows that the auction's price fluctuations and the parameters u_0 , α of the users' utility function greatly affect both the bid submitted and the resulting resource allocation pattern. In this section, we present a random walk model descriptive of the price fluctuation. Employing this model, a user can examine the impact of the value of parameter α of his bidding function on his average gap length, while "ignoring" the actual auction competition (of which he is unaware of), which is now simulated by means of the model. In particular, the fluctuation of the auctions' *cutoff price* over time is modeled as a random walk [5]. That is,

$$p_t = \begin{cases} p_{t-1} + \delta, & \text{with probability } q \\ p_{t-1} - \delta, & \text{with probability } 1 - q \end{cases} \quad \text{for } t = 1, 2, \dots \quad (1)$$

Henceforth, p_t is simply referred to as the auction price. Furthermore, according to the definition of the utility function considered (see Sect. 2) the fluctuation of the user's marginal utility is modeled as follows:

$$u_t = \begin{cases} u_0, & \text{if } u_{t-1} \geq p_{t-1} \\ \alpha \cdot u_{t-1}, & \text{if } u_{t-1} < p_{t-1} \end{cases} \quad \text{for } t = 1, 2, \dots, \quad (2)$$

where $\alpha \in (0, 1]$ is the discount factor. At time 0 the auction price raises for the first time above user's bid (i.e. $p_0 > u_0$) and fluctuates according to the random walk model thereafter.

Next, we discuss the appropriateness of our model. First, note that unlike with the actual auction, under our model, the price is not affected by the bids of the user considered. This is a good approximation of the actual auction because in an actual network many bidders will be participating. Furthermore, it could be argued that a random walk is not descriptive of the actual auction, since it is expected that when the mechanism is employed in an actual network the price will fluctuate around a the long term average value. However, we do not perform stationary analysis of the auction price in this paper. Instead, the model introduced is appropriate for *transient* analysis of the price dynamics of the auction. Consider, for instance, a case where a user is competitive in general, because his marginal utility is higher than the average auction price; however, his bid is instantly topped by the auction price due to some abrupt increase in demand (e.g. due to new users' arrival). Price evolution in this case can be emulated as a random walk process with positive drift. This actually leads to a conservative analysis, since in the actual auction the price would have the tendency to return more rapidly to its long-term average when it has deviated considerably from it. It should also be noted that transient analysis is very important, because due to the exponentially decreasing marginal utility, the user in the case described above becomes gradually less competitive, his service is interrupted and he may as well decide to drop out.

4 Analysis of the general model

In this section, we study the probability of user's "re-entrance", i.e. the probability that the user will succeed in being awarded again resources in future auctions, as well as the resulting resource allocation patterns. First, we compute the number of continuous price decrements k that are required until the user's marginal utility exceeds again the (constantly decreasing) auction price. We define the distance of the auction's price minus the user's exponentially decreasing valuation as $dist_t = p_t - u_t$. The parameter k sought is the minimum integer t that satisfies the inequality $dist_t \geq 0$, which implies that

$$k = \min_t \{p_0 - t \cdot \delta - \alpha^t \cdot u_0 \geq 0\}. \quad (3)$$

Thus, solution of this is the minimum possible "time" to elapse until the user's re-entrance. Assuming that after $k + j + i$ auctions, with j price increases and

$k + i$ price decreases, the user has still not succeeded in winning in the auction, the distance $dist_{k+i+j}$ is then:

$$dist_{k+i+j} = p_{k+i+j} - \alpha^{k+i+j} \cdot u_0 = p_0 + (j - i - k) \cdot \delta - \alpha^{k+i+j} \cdot u_0.$$

We proceed to compute the probability of user's re-entrance. We define:

$$i_j = \min_i \{ dist_{k+i+j} \leq 0 \} \quad (4)$$

Since we have already computed k , using (4) we can compute for each number of price increments j , the minimum number of price decrements $i_j + k$ that are required in order for the user to have a winning bid in the auction. Thus, we can compute the feasible values of the ‘‘gap size’’ of the user's resource allocation pattern, which equals $j + i_j + k$, for some j . For example, assume that for certain u_0 , δ , and α , we have: $k = 1$ and $i_1 = 1$, $i_2 = 2$ and $i_3 = 3$. Since $k = 1$, if the price drops at $t = 1$, then the user's bid will be a winning one again in that auction. However, if the price increases at $t = 1$ and then keeps decreasing (i.e. $j = 1$), then $2 = k + i_1$ price drops suffice for the user's marginal utility to top the auction's price. If this does not happen, because $j > 1$, then the soonest possible time for user's is $t = 5$ if $j = 2$ etc. Figure 2(b) depicts some possible auction's price fluctuations for which the user re-enters the network after 7 time units.




Fig. 2. The price fluctuates and user's marginal utility and bid are reduced.

Next, we compute the probability that user's exponentially decreasing marginal utility exceeds the auction's price for the *first time* at time $t = k + j + i_j$ for some j . In order to compute $Pr_{hit}(t)$, we must first compute the number $A(j, k)$ of ‘‘price fluctuation patterns’’ that lead to the user's re-entrance *for the first time* at time $t = j + i_j + k$. This number is obviously less than the total number of patterns of length $j + i_j + k$, namely $\binom{j+i_j+k}{j}$. Indeed, Fig. 2(a) depicts two possible ways of price fluctuation, with $j = 3$, that also lead to the user's re-entrance. However, for these patterns this does not happen for the *first time* after *exactly* 7 steps, thus these patterns should not be computed in $A(3, 1)$. (Also, for these patterns, the user's marginal utility would not continue to decrease exponentially until $t = 7$, as depicted in Fig. 2(a); it would have become u_0 once the

user achieves to re-enter.) Hence, in order to compute $A(j, k)$, we must exclude all those patterns that result in the user's marginal utility to exceed the auction price at some time t prior to $j + i_j + k$. In particular, we must exclude all the sub-patterns having size $j' + i_{j'} + k$, for all $j' < j$. Each such sub-pattern must be excluded $\binom{j+i_j-j'-i_{j'}}{j-j'}$ times; this is the number of possible allocations from time $j' + i_{j'} + k + 1$ to $j + i_j + k$ that are combined with the original sub-pattern. Hence, the number of acceptable price fluctuation patterns, $A(j, k)$ is:

$$A(0, k) = 1$$

$$A(j, k) = \binom{j + i_j + k}{j} - \sum_{j'=0}^{j-1} A(j', k) \cdot \binom{j + i_j - j' - i_{j'}}{j - j'}, \text{ for } j = 1, 2, \dots \quad (5)$$

As already explained, it is feasible for the user's marginal utility to exceed the auction's price only at certain times, namely at times $j + i_j + k$, where i_j is derived from (4). Hence the probability $\text{Pr}_{\text{hit}}(t)$ that the user's marginal utility (and bid in the auction) exceeds the auction price at some time t for the first time, is:

$$\text{Pr}_{\text{hit}}(t) = \begin{cases} 0, & \text{if } \nexists j \text{ s.t. } j + i_j + k = t \\ A(j, k) \cdot (1 - q)^{k+i_j} \cdot q^j, & \text{otherwise.} \end{cases} \quad (6)$$

The cumulative probability that the user's marginal utility will have exceeded the auction price up to some time t , denoted as $\text{Pr}_{\text{cHit}}(t)$ is $\text{Pr}_{\text{cHit}}(t) = \sum_{t'=1}^t \text{Pr}_{\text{hit}}(t')$. We also denote as $\text{Pr}_{\text{cHit}}^*$ the asymptotic value of $\text{Pr}_{\text{cHit}}(t)$, as $t \rightarrow \infty$, which equals the probability that the user will ever achieve re-entrance.

5 A special case

Assume that $\alpha = 1$, or equivalently that $u_t = u_0$ for $t = 1, 2, \dots$. This implies that $i_j = j$. This assumption results in a model that is much easier to analyze. This model will also provide a bound for $\text{Pr}_{\text{cHit}}^*$ of the general model, which is revisited in Sect. 6.

Proposition 1. *There holds:*

$$\lim_{t \rightarrow \infty} \text{Pr}_{\text{cHit}}(t) = \text{Pr}_{\text{cHit}}^* = \begin{cases} 1, & \text{if } 0 < q \leq \frac{1}{2} \\ \left(\frac{1-q}{q}\right)^k, & \text{if } \frac{1}{2} < q < 1. \end{cases} \quad (7)$$

Proof. The cumulative probability is $\text{Pr}_{\text{cHit}}^* = \sum_{t'=0}^{\infty} \text{Pr}_{\text{hit}}(t')$. Using (6) we obtain:

$$\text{Pr}_{\text{cHit}}^* = \sum_{j=0}^{\infty} A(j, k) \cdot (1 - q)^{(k+j)} \cdot q^j \quad (8)$$

First we consider the case where $0 < q < \frac{1}{2}$. In this case, the price has a negative drift. Hence, it is certain that at some point it will drop below the user's marginal utility $u_t = u_0$; see [5]. Thus, the respective cumulative probability is 1. Next, we consider the case where $\frac{1}{2} < q < 1$. We have:

$$\Pr_{\text{cHit}}^* = \sum_{j=0}^{\infty} A(j, k) \cdot (1-q)^{(k+j)} \cdot q^j = \left(\frac{1-q}{q}\right)^k \cdot \sum_{j=0}^{\infty} A(j, k) \cdot q^{(j+k)} \cdot (1-q)^j \quad (9)$$

The sum in the right hand side of (9) is identical to that in (8) except that q and $1-q$ are swapped. Note that in (8), if $0 \leq q < \frac{1}{2}$ the sum converges to 1. On the contrary, in (9) we have that $0 \leq 1-q < \frac{1}{2}$. Therefore, the sum in (9) converges to 1, which implies that $\Pr_{\text{cHit}}^* = \left(\frac{1-q}{q}\right)^k$, if $\frac{1}{2} < q < 1$. Finally, by continuity for $q \rightarrow \frac{1}{2}$, $\Pr_{\text{cHit}}^* = 1$ too. \square

The same result for $k = 1$ and $q \in (\frac{1}{2}, 1)$ is also established differently in [5]. However, the above proposition was proved for any value of k , because it will be used in the next section. The aforementioned limits are useful in order to understand the user resource allocation patterns that result when users participate in a series of auctions. When the probability of price increase is lower than $\frac{1}{2}$, it is certain that the user will eventually succeed in topping the auction's price and reserve network resources again. This is obviously the case for the actual network (where the cut-off price is determined by competing users) at periods of low to medium demand. Of course, the lower the probability of price increase, the sooner the user will be re-allocated network resources in the auction. On the contrary, when the probability that the price decreases is higher than the probability that the auction's price will increase - that is, whenever the competition is high - it is uncertain if the user will eventually manage to receive service again.

The probability that this happens is $\left(\frac{1-q}{q}\right)^k$, which is a decreasing function of q . The higher the probability of price increase, that is the more intense the competition in the auction, the less probable it is for the user to receive service again. This is justified since if the user's bid is topped at a very competitive auction, it becomes very hard for this user's subsequent bids to become winning again in the series of auctions that follow, where the cut-off prices tend to increase. Exponential reduction of the user's bids (i.e. $\alpha < 1$) may only make matters worse, i.e. result in a smaller value of \Pr_{cHit}^* . Notice also, that for $\alpha = 1$, \Pr_{cHit}^* depends on q as well as k , which due to (4) equals $k = \lceil (p_t - u_0)/\delta \rceil$. It should also be noted that although the result of Proposition 1 is asymptotic for $t \rightarrow \infty$, the time horizon is not actually infinite in the sense that the probability of re-entrance can be well approximated by taking the first few terms of the series in (8). Thus, the result is still applicable for the purpose of analyzing transient phenomena. The same comment also applies for the results of Sect. 6 and 7.

6 Revisiting the general model

Having studied the special case of Sect. 5, we return to the general model presented in Sect. 3 with $\alpha < 1$, the analysis of which is more complicated, as already mentioned. The difficulty of providing a closed-form equation for \Pr_{cHit}^*

(i.e. the asymptotic value of $\Pr_{\text{cHit}}(t)$ for $t \rightarrow \infty$) is due to the fact that neither i_j nor $A(j, k)$ are known in closed form. However, it is possible to provide some bounds for these probabilities. Under certain assumptions these bounds are tight.

Clearly, whenever the probability q that the price goes up in the random walk is less than $\frac{1}{2}$, then the cumulative probability converges to 1, i.e. $\Pr_{\text{cHit}}^* = 1$, $0 \leq q \leq \frac{1}{2}$. This can be proven by applying the same arguments with these provided in the proof of Proposition 1. For $q > \frac{1}{2}$, the computation of \Pr_{cHit}^* is very complicated. Below, we provide an upper and lower bound for this probability.

Proposition 2. *We have:*

$$(1-q)^k + k \cdot \frac{(1-q)^{k+i_1}}{1-q \cdot (1-q)^{i_1}} \leq \Pr_{\text{cHit}}^* < \left(\frac{1-q}{q}\right)^k, \quad \text{if } \frac{1}{2} < q < 1. \quad (10)$$

Proof. First, we prove that $\Pr_{\text{cHit}}^* < \left(\frac{1-q}{q}\right)^k$, if $\frac{1}{2} < q < 1$. This is easily proven since $\left(\frac{1-q}{q}\right)^k$ is the value of \Pr_{cHit}^* if $\alpha = 1$, i.e. if $u_t = u_0$ for $t = 1, 2, \dots$. In the general model, we have $u_t \leq u_0$. Clearly, by monotonicity, under the exponentially decreasing marginal utility (and bid) the user can never be better off compared to the case where $u_t = u_0$ for $t = 1, 2, \dots$. Thus, $\Pr_{\text{cHit}}^* < \left(\frac{1-q}{q}\right)^k$, if $\frac{1}{2} < q < 1$.

Next, we prove that $(1-q)^k + k \cdot \frac{(1-q)^{k+i_1}}{1-q \cdot (1-q)^{i_1}} \leq \Pr_{\text{cHit}}^*$. We have

$$\Pr_{\text{cHit}}^* = (1-q)^k + A(1, k) \cdot q \cdot (1-q)^{k+i_1} + A(2, k) \cdot q^2 \cdot (1-q)^{k+i_2} + A(3, k) \cdot q^3 \cdot (1-q)^{k+i_3} + \dots \quad (11)$$

Due to the exponential decrease of the marginal utility, we have $u_j - u_{j-1} = \alpha \cdot (u_{j-1} - u_{j-2})$, which easily implies that the difference $\Delta i_j = i_j - i_{j-1}$ is non-increasing in j . That is, the number Δi_j of extra price decrements required to cope with one additional price increment is non-increasing in j , although i_j itself is increasing in j . Thus, $\max\{\Delta i_j\} = \Delta i_1 = i_1$. Furthermore, $\Delta i_2 = i_2 - i_1$. Therefore, we obtain that $i_2 \leq 2 \cdot i_1$. Similarly, $\Delta i_3 = i_3 - i_2 \leq i_2 - i_1 \leq i_1$, which implies that $i_3 \leq i_2 + i_1 \leq 3 \cdot i_1$. Continuing this argument, it follows easily that in general:

$$i_j \leq j \cdot i_1 \quad \text{for } j = 2, 3, \dots \quad (12)$$

Applying (5), we have $A(1, k) = k$. It is also easily seen that $A(j, k) \geq A(j-1, k)$ for $j = 2, 3, \dots$. The proof is omitted for brevity reasons. Therefore, $A(j, k) \geq k$ for $j = 1, 2, \dots$. Combining this bound for $A(j, k)$ and that for i_j as given by (12), we obtain from (11):

$$\begin{aligned} \Pr_{\text{cHit}}^* &\geq (1-q)^k + k \cdot (1-q)^{k+i_1} \cdot [1 + q \cdot (1-q)^{i_1} + q^2 \cdot (1-q)^{2 \cdot i_1} + \dots] \Rightarrow \\ \Pr_{\text{cHit}}^* &\geq (1-q)^k + k \cdot \frac{(1-q)^{k+i_1}}{1-q \cdot (1-q)^{i_1}} \quad \{\text{since } q \cdot (1-q)^{i_1} < 1\} \end{aligned}$$

□

We have so far provided an upper and a lower bound for \Pr_{cHit}^* for $\frac{1}{2} < q < 1$. We proceed to investigate the distance between these bounds, which expresses their accuracy as an approximation of \Pr_{cHit}^* . Notice that the value of the distance depends also on i_1 , which can be computed from (4) employing the values of p_0 , u_0 and δ . Recall that we have implicitly assumed in our analysis that the user considered is competitive, which implies that p_0 is not much larger than u_0 . Users generally select large values of α (i.e. close to 1), in order to avoid very fast decaying of their marginal utility. Therefore, in order to simplify our study on the bounds on \Pr_{cHit}^* , we restrict attention to small values of k only, which are the ones of practical importance, and we approximate i_1 with k . (Notice that, for $\alpha = 1$ and $p_0 = u_0$, we have $i_1 = k$.) As depicted in Table 1, the difference of the two bounds is very small if $q \geq 0.6$. This distance increases significantly for $q \in (0.5, 0.6)$ and is unacceptable for $k = 2$ and $k = 3$. Additional numerical results show that the accuracy improves again for large k .

Table 1. The distance of the bounds of \Pr_{cHit}^* as a function of k and q .

	Difference of bounds of \Pr_{cHit}^*								
q :	0.51	0.53	0.55	0.57	0.60	0.65	0.70	0.75	0.80
$k = 1$	0.15	0.12	0.09	0.08	0.06	0.03	0.01	0.01	0.002
$k = 2$	0.55	0.45	0.37	0.31	0.23	0.13	0.08	0.04	0.02
$k = 3$	0.72	0.56	0.43	0.33	0.22	0.11	0.05	0.02	0.01

7 Studying the resulting user resource allocation patterns

Recall that for the user considered, his marginal utility is topped by the auction price at $t = 0$. If the marginal utility exceeds again the auction price *for the first time* at time t , then the gap length is defined as t . The computation of the average length can be done numerically as follows: First, we compute k , and i_j for $j = 1, 2, \dots$. Then, we compute numerically $\Pr_{\text{hit}}(t)$ from (6). In order to compute the average gap length, we must distinguish two cases:

- $0 < q \leq \frac{1}{2}$: In this case, since the probabilities $\Pr_{\text{hit}}(t)$ sum to 1, it suffices to add the products of all possible gap lengths $j + i_j + k$ with the respective hitting probability $\Pr_{\text{hit}}(j + i_j + k)$.
- $\frac{1}{2} < q < 1$: In this case, since the sum of probabilities $\Pr_{\text{hit}}(t)$ do not sum to 1, we add the products of all possible gap lengths $j + i_j + k$ with the respective hitting probability $\Pr_{\text{hit}}(j + i_j + k)$; then we normalize, and by dividing this sum with the sum of aforementioned probabilities, that is with the limit \Pr_{cHit}^* . Therefore, in this case we compute the conditional average gap length, given that the user does re-enter.

Therefore, the average gap length AGL is given by:

$$AGL = \begin{cases} \sum_{j=1}^{\infty} \Pr_{\text{hit}}(j + i_j + k) \cdot (j + i_j + k), & \text{if } 0 < q \leq \frac{1}{2} \\ \frac{1}{\Pr_{\text{cHit}}} \cdot \sum_{j=1}^{\infty} \Pr_{\text{hit}}(j + i_j + k) \cdot (j + i_j + k), & \text{if } \frac{1}{2} < q < 1 . \end{cases}$$




Fig. 3. The average gap length as a function of q for $u_0 = 100$, $\delta = 10$, $\alpha = 1$.

Numerical results are depicted in Fig. 3 for $\alpha = 1$. The values of the average gap length are symmetric with respect to the axis $q = \frac{1}{2}$. Note that for values of q that are very small the user is most likely to top the auction's price immediately (after 1 price decrease at the next auction), hence resulting in a value of the average gap length that is very close to 1. Also, for values of q close to 1 the same values of AVG apply, since in this case the user's bid will either become winning again immediately or this will never happen.

The simulation results reported in [1], indicate that the vast majority of user resource allocation patterns are either perfectly consistent (high-value users) or comprise very few resource allocations (non-competitive users). For users whose value is often close to the auctions' cutoff prices, it has been observed that their respective resource allocation patterns are typically in accordance with that depicted in Fig. 1 (a), having few, large gaps; the less preferable patterns with frequent small-sized gaps (like those depicted in Fig. 1 (b)) are rare as required. These phenomena are captured by our model. Indeed, users whose bids are often close to p_t can be viewed as participating in an auction with $q \simeq \frac{1}{2}$ and thus the gaps of their resource allocation patterns are expected to be large.

8 Validation and usefulness of the model

The analysis of the paper was based on the assumption that it suffices to model the auction's cut-off price by means of a random walk model. In this section, we provide an experimental evaluation of the proximity of the estimate of the model of Sect. 3 regarding the average gap length to that obtained by auction simulations.

The methodology we adopt is the following: We run various auction simulations where a large population of competing users bid for resources, (see [1]). After each auction, we store the users who have gaps in their resource allocation patterns, their respective u_0 and the gap length. Then, we define an equivalent random walk model: We estimate the probability q of price increase in the equivalent random walk by dividing the number of price increments of the actual auction with the total number of price fluctuations. We also estimate the step δ by averaging the actual price differences $|p_t - p_{t-1}|$ throughout the actual auction. After running a large number of simulations of the auction, we group together those users that have gaps¹ with similar u_0 and competed in auctions that can be described with similar δ , q . For each such group, we compute the mean actual gap lengths obtained in the experiments; this is then compared to the average gap length derived by means of the equivalent random walk model using Mathematica. Such comparisons reveal that in most cases our model provides a very good approximation of the actual gap length arising in the simulated auction. A typical example is provided in Table 2: in three simulation runs a total of 10 gaps for users with $u_0 \simeq 272$, $q \simeq 0.61$, $\delta \simeq 16$ were recorded. Their mean gap length was 3.3636 while our random walk model gives for the same u_0 , q and δ an average gap length of 3.4883.

Table 2. Typical experimental results demonstrating the accuracy of our model.

Auction	Auction simulation results ($\alpha = 0.99$)	Estimated R. Walk	
	$(u_0, \text{Gap size})$	q	δ
1	(274, 2), (274, 1), (269, 2)	0.62	15.9
2	(262, 5), (262, 3)	0.59	15.9
3	(276, 7), (276, 3), (276, 2), (278, 7), (278, 2), (278, 3)	0.62	16.1
Equivalent random walk model: $u_0 = 272$, $q = 0.61$, $\delta = 16$			

Next, we explain how the model can be employed in order to support users. Recall that users select a utility function on the basis of which the network bids on behalf of the users; see Sect. 3 and [1]. Note that the user's selection of α is of particular importance for the QoS attained. The higher α the higher the probability of user's re-entry and the lower the average gap length experienced, but also the higher the charge. Therefore, budget-constrained users whose u_0 is often close to the auction cut-off price end up with a small net benefit. Therefore, it might be more profitable for the bids of such users to be losing in some auctions, as long as the resulting gaps are of acceptable sizes. Hence, by selecting an appropriate value of α , a user can affect both the QoS experienced and the expected net benefit to be attained from the auction. The best choice of course depends on the user's sensitivity with respect to the QoS versus the respective charge. Clearly, the user should employ in this selection process a proxy of the expected QoS level. The average gap length computed by means of the random walk model can definitely be used, and as already explained it constitutes an

¹ Note that a user may experience multiple gaps in his resource allocation pattern.

accurate estimate. To better illustrate these ideas we present certain numerical results in Table 3, which shows the sensitivity of the average gap length with respect to $\alpha \in [0.960, 0.999]$ for the equivalent random walk model of Table 2. We assume that we are in the context of a UMTS network where a new auction is run every 10 msec [1]. Clearly, QoS-sensitive users should opt for a value of α in $[0.99, 0.999]$, while those more interested in attaining a somewhat lower charge than receiving almost perfect QoS should select α in $[0.965, 0.990]$. On the contrary, any selection of $\alpha < 0.96$ results in an average gap length of more than 500 msec, and thus it should be avoided.

Table 3. Sensitivity analysis depicting the relation of the value of α involved in the user's utility/bidding function and the respective expected average gap length.

Average Gap Length sensitivity analysis (model parameters: $u_0 = 278$, $q = 0.355$, $\delta = 10.9$)										
α :	0.955	0.960	0.965	0.970	0.975	0.980	0.985	0.990	0.995	0.999
<i>AGL</i> :	53.06	51.85	30.89	27.78	24.08	18.37	12.57	7.08	4.34	3.47

9 Conclusions

In this paper, we have analyzed the mechanism of [1], where users participate in a series of consecutive sealed-bid auctions, aiming at reserving the same amount of units of the resource auctioned. We have introduced a random walk to model the fluctuations of the auctions' cut-off price. Using this model, we have studied the resource allocation patterns of a user whose utility/bidding function remains constant as long as the user's bids are winning while it is exponentially decreasing in case the cut-off price exceeds the user's bid. Finally, we have provided experimental results validating the model and demonstrated how it can be used in order to support the user in the selection of one of the parameters for his bidding function.

References

1. M. Dramitinos, G. D. Stamoulis, C. Courcoubetis.: Auction-based Resource Reservation in 2.5/3G Networks. Kluwer/ACM Mobile Networks and Applications Special Issue: Mobile and Pervasive Commerce, 9:6, pp. 557-566, December 2004.
2. P. Maillé and B. Tuffin.: An Auction-Based Pricing Scheme for Bandwidth Sharing with History-Dependent Utility Functions. In Proc. of the First International Workshop on Incentive Based Computing, France, September 2005.
3. P. Maillé and B. Tuffin.: The Progressive Second Price Mechanism in a Stochastic Environment. Netnomics: 5:2, pp. 119-147, 2003.
4. FP6-IST-507607 Project B-BONE, <http://b-bone.ptinovacao.pt/>.
5. S. Ross.: Stochastic Processes. Wiley Series in Probability and Mathematical Statistics, ISBN 0-471-12062-6, 1996.