

Auction-based Resource Allocation in UMTS High Speed Downlink Packet Access (HSDPA)

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
Email: {mdramit, gstamoul, courcou}@aueb.gr

Abstract—Data, audio and video services over 3G networks possess certain requirements in QoS and thus in resources. Accommodating such service requests constitutes a challenge for UMTS networks. To this end, 3GPP Release 5 introduces the High Speed Downlink Packet Access (HSDPA). In this paper, we propose: i) an auction-based HSDPA resource allocation mechanism, which consists of a series of Generalized Vickrey Auctions that are conducted in the short time-scale over which the network resources are allocated and ii) a set of innovative user utility functions. These functions provide a quantification of both the user's willingness to pay for reserving resources in the short time-scale auctions and his attained utility from the pattern of resources allocated during his long time-scale service session. The utility functions are additive and are used as bidding functions in the series of auctions the user participates. Each user selects one of these functions, which is scaled by his total willingness to pay. Then, the network runs the auctions by bidding on behalf of each user. The effectiveness of our mechanism has been assessed experimentally and it appears that most of the users either are served very satisfactorily or are allocated very limited quantities of resources (if at all) at a low total charge.

Index Terms—Auctions, resource reservation, UMTS, HSDPA.

I. INTRODUCTION

AUCTIONS have recently received considerable attention as an economic mechanism for resource reservation and price discovery in networks. This is of particular importance for many practical cases such as the provision of network services with long duration and high demand for network resources. A prominent case is that of UMTS [1], [2]: except for voice, the duration of services that users generally request is significantly longer than a single time slot¹ over which resource units can be reserved; e.g., large file transfers or video streaming of a certain bit rate. These services require Quality of Service (QoS), and thus are very demanding

in resources. Accommodating such services consists a challenge for 3G networks due to the scarcity of radio spectrum. To this end, 3GPP Release 5 defines the High Speed Downlink Packet Access (HSDPA) with a maximum bit rate of 10Mbps, a significant improvement compared to the 2Mbps of the 3GPP Release '99 [3]. Since the mobile industry is extremely competitive and spectrum is a scarce resource, efficient exploitation of spectrum is crucial for commercial providers.

In this paper, we propose an auction-based Resource Allocation mechanism for HSDPA (HERA). This mechanism comprises: (a) a series of repeated Generalized Vickrey Auctions, one per slot, and (b) certain bidding functions for the user to choose from. HERA is an extension of the ATHENA mechanism which we had proposed in [4] for resource reservation in 3GPP Release '99 UMTS networks. In particular, here we extend and adapt the approach of [4] in order to be applicable with HSDPA. We deal with how to satisfy the demand for large time scale services of constant bit rate (CBR), by allocating *nearly consistently* resources to those users that value them the most. We also provide an extension for variable bit rate (VBR) services. We assume that the population of users generally varies over time, as in practical cases, which further complicates the problem. Our objective is to allocate resources efficiently. That is, award the network resources to the users who value them the most while providing proper incentives for rational network usage. Note that demand fluctuations make it impossible for a provider to publish fixed prices for the resources, which would have been the way to attain efficiency under steady demand.

As already mentioned, we propose a series of consecutive auctions as a means for attaining efficiently consistent reservation of resources. Each auction spans a time period of 1 sec and its outcome specifies the allocation of 500 2msec HSDPA frames (see below). Each user aims in being allocated a certain quantity of frames within the 1 sec period so that his target

¹For UMTS, the time slot is taken to be the 10msec UTRAN frame.

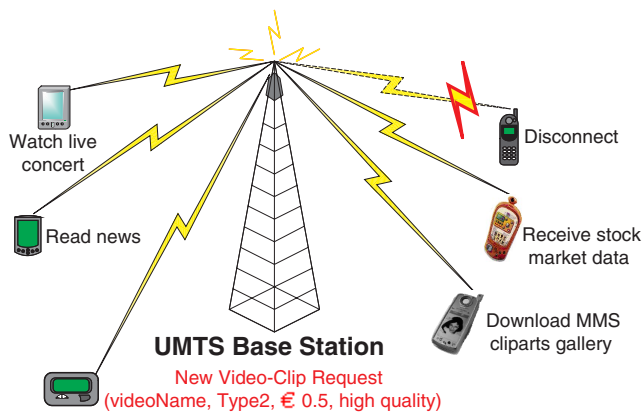


Fig. 1. A user of a UMTS cell submits a video service request consisting of the video name, his selection of a utility function and his total willingness to pay for that service of a certain quality (bit rate).

bit rate is attained. Since constant resource allocation among auctions may not be feasible for all users, it is very important to construct *meaningful user utility functions* that actually express users' preferences regarding the degradation of the service quality. These utility functions are additive, although not memoryless. Hence, they can be used as bidding functions in our mechanism, thus providing a quantification of users' preferences in cases of inconsistent resource allocation patterns. Each user selects one of these functions, which is scaled by his maximum total willingness to pay, which is to be 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, according to his respective selection of bidding function. This way, a user demanding a service just declares once this selection as well as his willingness to pay for the service desired, as shown in Fig. 1. The definition of these functions and the fact that users have to choose one of them constitutes an integral part of our mechanism for nearly consistent resource reservation. The effectiveness of our approach has been assessed experimentally by analyzing the resulting resource allocation patterns. It is seen that users with competitive bids are allocated patterns that are in very good accordance to their respective preferences, while users with non-competitive bids are allocated very limited quantities of resources (if at all) at low charges.

As already mentioned, HERA is an extension of the ATHENA mechanism which we had proposed in [4] for resource reservation in 3GPP Release '99 UMTS networks. In this paper, we extend and adapt ATHENA in order to be applicable with HSDPA. The new features introduced and the modifications concern the time-scale of the auction, the capacity of the downlink shared

channel, the definition of the resource units auctioned, and the technological means for mapping service requests to resource reservations. Moreover, in HERA we also provide a new user utility function pertaining to variable bit rate (VBR) services. Both the simulation model and the software employed for the evaluation of the mechanism have been enhanced accordingly. We have run numerous new experiments involving a variety of mixes of competing CBR and VBR users' service sessions under various network congestion levels, i.e. low, medium and high demand.

An alternative to performing auctions sequentially (one per slot) would be to run combinatorial auctions, thus allowing users to bid for resources spanning several consecutive slots. Such an approach should not be adopted due to the following reasons: i) the computational overhead is excessive, ii) guaranteeing resources for the entire duration of the sessions currently served would render impossible for the network to serve users with higher value that would arrive late, thus resulting in high connectivity delay, poor revenue and inefficiency.

The remainder of this paper is organized as follows: In Section II, we provide a brief background on auctions. In Section III, we present UMTS HSDPA and its main features. In Section IV, we define the proposed auction-based mechanism for the HSDPA. In Section V, we define user utility functions capable to express certain types of user preferences with respect to patterns of inconsistent resource allocation. In Section VI, we present certain extensions of our utility functions. In Section VII we provide an experimental assessment of our auction. Finally, in Section VIII, we provide some concluding remarks.

II. BACKGROUND ON AUCTIONS

In this section, we present fundamental definitions, theorems and results from auction theory; see [5], [6].

An auction is a mechanism based on a pair of rules, namely the *allocation rule* that defines which good is allocated to whom and the *payment rule* that defines the charge of the auction winner(s). A participant of an auction is called *bidder* while the entity conducting the auction *auctioneer*.

There are many ways to classify the various auction mechanisms. Thus, auctions are referred to as *simple* or *single-unit* if only one good is auctioned and *multi-unit* if multiple units of a good (e.g. integral units of a link's bandwidth) are to be traded. Moreover, depending on whether bids are made in public or submitted as sealed envelopes, the auction is referred to as *open* or *sealed* respectively. Auctions that maximize seller's revenue are referred to as *optimal* and those that maximize

social welfare are referred to as *efficient*. If the auction is conducted in rounds, then it is called *progressive*. Depending on the payment rule, the auction can be either *uniform* if the same amount of money is paid by all for each unit of the good awarded and *discriminatory (pay-your-bid)* if each user pays his bid for each object he wins.

A. Simple auctions

A *bid* in the context of simple auctions is the amount of money offered by a bidder for the item auctioned. The best-known mechanism is by far the *English* auction, where the seller starts with a minimum price that is actually incremented until there is only one person claiming the item, whom the item is awarded to. The *Dutch* auction is exactly the opposite mechanism. The price is initially high and is gradually decremented until a bidder claims the object. The item is awarded to him for a charge equal to the current price. The sealed bid auctions (*1st price* and *2nd price* or *Vickrey*) consist of two phases: a) the first one, where bidders submit sealed envelopes with their bids and b) the second one, where these envelopes are opened. The item is then awarded to the bidder who submitted the highest bid. The winner pays his bid at the *1st price* auction and the highest losing bid, - i.e. the *2nd highest bid* - at the *Vickrey* auction. It has been proved that under the *Vickrey* auction it is best for each bidder to honestly bid his true value for the item being awarded. This property is referred to as *incentive compatibility*.

B. Multi-unit auctions

A *bid* in the context of multi-unit auctions is defined to be the pair (p, q) of the per unit expressed willingness to pay p for a quantity q of units. All simple auctions can be generalized to multi-unit auctions. Incentive compatibility holds only for the generalizations of the *Vickrey* auction. The rules of the *Generalized Vickrey Auction (GVA)* prescribe that: i) each user reports his valuation for a subset or for all points of his demand function for units of the good auctioned, ii) units are allocated to the highest bids until demand exhausts supply, iii) each user is charged according to the social opportunity cost that his presence entails. Hence, each user pays for the units he is awarded the losing bids that would have become winning if all his own bids were set to zero. Hence, winners pay less than their respective bids.

III. HSDPA

HSDPA was finalized in the 3GPP Release 5 specification and is considered the most significant radio-related update since the release of the first 3GPP Release

'99 specification. It is based on a distributed architecture where processing is closer to the air interface at the base station (Node B). The HSDPA operation is performed on the High Speed Downlink Shared Channel (HS-DSCH). This is superior to the Downlink Shared Channel (DSCH) of the original 3GPP Release 1999 in terms of performance, efficient usage of radio spectrum and target bit rates. HSDPA extends bit rates up to 10Mbps, a significant upgrade compared to the 2Mbps of the previous release². This results from a set of unique additional features of HS-DSCH, namely: i) the very short frame length (only 2msec) ii) the fixed spreading factor of 16 with maximum 15 codes and iii) the support of 16QAM modulation in addition to QPSK of DSCH. The number of HS-DSCH codes in each cell is configured by the Radio Network Controller (RNC). The HSDPA system architecture is depicted in Fig. 2.

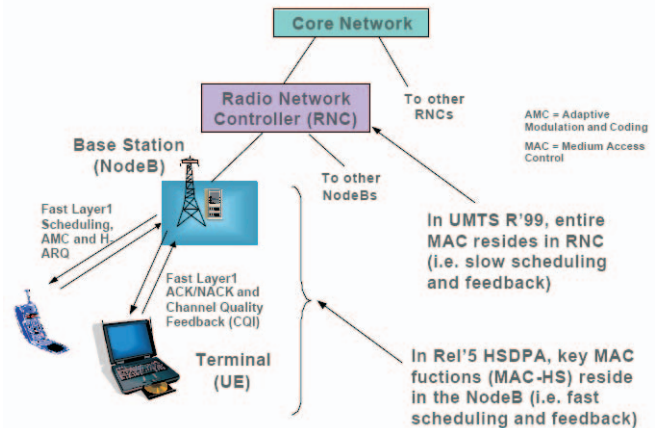


Fig. 2. The HSDPA system architecture and its innovative features compared to UMTS Release 99.

Each user is allocated a bit rate of the HS-DSCH which depends on the number of codes and frames that have been assigned to him. The HS-DSCH frame can be configured to contain either 1 or 15 codes. Therefore, several packet users can be code- and time-multiplexed. However, it is possible that the sharing of the HS-DSCH among users is done mainly in the time domain by the HSDPA scheduler; this is considered in [8] as the most preferable use of the channel. Hence, due to the very short time duration of the HS-DSCH frame, it is desirable that one user is awarded exclusively the HS-DSCH within the 2msec time scale.

²The capacity of the WCDMA sector can be typically improved via base station beamforming that allows to reduce interference levels by transmitting the signal via a narrow beam to the desired user [7].

IV. HSDPA AUCTION-BASED RESOURCE ALLOCATION

The HSDPA scheduler is responsible for awarding access to the HS-DSCH. We propose that the assignment of resources, i.e. the 2-msec HS-DSCH frames, be performed by means of auctions. We conduct a sequence of “mini-auctions”, each concerning reservation of resources within 1 sec; hence 500 2msec frames are auctioned. Each mini-auction is a sealed-bid Generalized Vickrey Auction (GVA) [5], with atomic bids (i.e., bids that are either fully satisfied or rejected) of the type (p, q) , where q is the quantity of resource units sought in the present auction and p is the price proposed for each such unit. Henceforth, unless otherwise specified, we restrict attention to services requiring constant bit rates. For UMTS, since resources are allocated in bits, if the service involves traffic of a specific bit rate m , then we have $q = m \cdot t_a$, where t_a is the time period per mini-auction. (For each such service session, one bid is placed per mini-auction; see below.) In a GVA, users have the incentive to bid truthfully their willingness to pay. This is a very attractive property, motivating our selection of running a GVA in each mini-auction.

The time period t_a , was set to be 1 sec, so that it is both (a) small enough to capture market fluctuations due to the entry or departure of users from the network, and (b) large enough to accommodate various bit rate requests. The minimum bit rate that can be awarded to a user within 1 sec - by awarding just one frame per second - is very small (approximately 20 Kbps or 1,33 Kbps, if the HS-DSCH frame is configured to contain 1 or 15 codes respectively), thus allowing for a very fine granularity in the set of achievable bit rates. This way, the exclusive use of the High Speed Downlink Shared CHannel (suggested in [8]) is also possible, as long as the RNC configuration defines that the number of HS-DSCH codes per frame is set to 1.

The mapping of users’ bids to actual network resource reservation is carried out arbitrarily without having to resort to elaborate techniques for the allocation of the various frames within the 1 sec period. For example, a user awarded a bit rate of 5Mbps within the auctioned 1 sec, actually uses the HS-DSCH exclusively for 250 frames of 2 msec. The frames used by the same session within two consecutive periods may be at most 2 secs apart, which may affect jitter adversely. This problem can be taken care of, if user terminals are equipped with a playout buffer to smooth the data flow and bound jitter. The necessary buffer capacity is of the order of magnitude of MBytes, i.e. capable of storing data corresponding to a time duration of few seconds. This

requirement is reasonable with the present technology for devices. Another solution would be to spread the frames awarded within the 1 sec period uniformly in time so that jitter is maintained low. The formal definition of such an algorithm is left for future research.

In a realistic case of a UMTS network, it would not be feasible for users to participate in all these mini-auctions. Hence, we define utility functions, pertaining to the various services. These functions are provided by the network operator as bidding functions for the user i to choose from; they are scaled by the user’s total willingness to pay $U_{s,i}$ for the service s , which is to be given by the user himself (as part of his service request). This approach is similar to that of 3GPP regarding predefined Quality of Service (QoS) profiles. Then, the network runs the various mini-auctions by bidding on *behalf* of each user, according to his respective selection of bidding function.

V. USER UTILITY FUNCTIONS

Below, we briefly present the utility functions and their properties. We assume that the user’s i value $u_{s,i}$ for obtaining the service s is the sum of the marginal “sub-utilities” $v_{s,i}^{(t)}$ attained due to each successful allocation (i.e. the allocated bit rate x_i equals the desired m_i); thus, $u_{s,i}(x_i^{(1)}, \dots, x_i^{(K_{s,i})}) = \sum_{t=1}^{K_{s,i}} v_{s,i}^{(t)}(x_i^{(1)}, \dots, x_i^{(t)})$. (Indices i and s denote the user and service respectively; $K_{s,i}$ denotes the number of mini-auctions for user i and service session s ; superscript t ranges from 1 to $K_{s,i}$, indicating the “current” mini-auction.) Next, we define meaningful utility functions, pertaining to the various services. These functions reflect the fact that, when there are gaps in the resource allocation pattern, not only the amount of slots but also the *way* these are allocated - i.e. the resource allocation pattern - makes considerable difference to the degree of user satisfaction. Thus, by selecting one of the predefined user utility functions, each user declares his preferred form of allocation pattern for the cases where perfectly consistent resource reservation is not possible. We have defined the following utility functions:

- Type 1: *Indifferent to the allocation pattern*. This type pertains to volume-oriented users, such as those downloading news. The utility attained depends only on the quantity of resources allocated, as opposed to the allocation pattern. Hence the user’s total value $U_{s,i}$ is equally apportioned among the various slots. That is,

$$v_{s,i}^{(t)}(x_i^{(1)}, \dots, x_i^{(t)}) = \mathbf{1}(x_i^{(t)} = m_i) \cdot \frac{U_{s,i}}{K_{s,i}},$$

where $U_{s,i}$ is the user’s declared total willingness-to-pay and $\mathbf{1}(\cdot)$ is the indicator function. There-

fore, the lack of information transfer due to the gaps in the resource allocation pattern, results in proportional loss of user satisfaction. As already mentioned, an appropriate example of users belonging to this type is those accessing news with a certain target bit rate. Also, users downloading information by means of FTP can be considered as users of this type, provided that the FTP session is capable of resuming download when time-outs occur. This utility function is suitable for the UMTS Background Class services with m_i pertaining to the Maximum Bit-rate parameter of this class [9]. Note that since users are “volume-oriented”, a rate less than m_i would also be useful to them. That is, if a user has the last winning bid and he can be awarded just a fraction of the target bit rate, it is both meaningful and efficient to do so. This can be attained if the assumption of atomic bids is relaxed. On the contrary, it is obviously inefficient to assign to the higher value users a bit rate less than their maximum.

- Type 2: *Sensitive to the service continuity*. This type pertains (among other cases) to users that prefer watching consistently half of a football match rather than watching multiple shorter periods. Thus, they prefer the allocation pattern of Fig. 3(a) to that of Fig. 3(b). A sub-utility expressing this preference is

$$v_{s,i}^{(t)}(x_i^{(1)}, \dots, x_i^{(t)}) = \mathbf{1}(x_i^{(t)} = m_i) \frac{U_{s,i}}{K_{s,i}} \cdot \alpha^{d_i^{(t)}},$$

where $\alpha \in (0, 1)$ is a discount factor, $d_i^{(t)}$ is the distance between the current and the previous slots during which user i achieved reservations. History of previous allocations influences $v_{s,i}^{(t)}$ through the value of d_i , which is kept track of. This utility function is suitable for the UMTS Streaming Class services [9].

- Type 3: *Sensitive to the regularity of the resource allocation pattern*. Such users prefer the allocation pattern of Fig. 3(b) to that of Fig. 3(a). The corresponding sub-utility is

$$v_{s,i}^{(t)}(x_i^{(1)}, \dots, x_i^{(t)}) = \mathbf{1}(x_i^{(t)} = m_i) \frac{U_{s,i}}{K_{s,i}} \cdot \alpha^{\max\{0, \Delta d_i\}},$$

where Δd_i is the difference of $d_i^{(t)}$ (as defined above) and the length of the previous gap. Note that the factor $\alpha^{\max\{0, \Delta d_i\}}$ equals 1 (resp. is less than 1) if Δd_i is non-positive, i.e. when the received quality of service improves (resp. deteriorates).

For all three types, if the user is constantly allocated resources (and thus the best possible quality is achieved), then the utility obtained is $U_{s,i}$.

The aforementioned utility functions are not the only possible ones reflecting the user satisfaction w.r.t. the




Fig. 3. Inconsistent resource allocation patterns.

allocation pattern attained. What is important, is that the form of these utilities reflects rightly the preferences of each type of user.

VI. EXTENSIONS OF THE USER UTILITY FUNCTIONS

We have extended the definitions of these utility functions so as to cover more interesting cases.

For instance, user i may be willing to watch just “good quality” video - of a rate r_{good} - whenever watching the video with the preferred “high quality” r_{high} is not feasible. Watching the video with consistently either “good quality” or “high quality” results in different degrees of user satisfaction; hence, the total willingness to pay respectively equals $U_{\text{good},i}$ and $U_{\text{high},i} = U_{\text{good},i} + \Delta U_i$ where ΔU_i expresses the user’s *extra* satisfaction for video of “high quality”. Due to the possible fluctuations of the rate attained, a proper utility function for this type of users is

$$v_{s,i}^{(t)}(x_i^{(1)}, \dots, x_i^{(t)}) = \mathbf{1}(x_i^{(t)} \geq r_{\text{good}}) \frac{U_{\text{good},i}}{K_{s,i}} \cdot \alpha^{d_{i,1}^{(t)}} + \mathbf{1}(x_i^{(t)} = r_{\text{high}}) \frac{\Delta U_i}{K_{s,i}} \cdot \alpha^{d_{i,2}^{(t)}},$$

where $d_{i,1}^{(t)}$ (respectively $d_{i,2}^{(t)}$) is the distance between the current and the previous slots during which user i achieved a reservation of at least r_{good} (resp. exactly r_{high}). Finally, note that the number of atomic bids to be given on behalf of each user at each mini-auction equals the number of alternative quality levels. Thus, for the aforementioned case of two such levels, two *summable* bids should be given per user: one expressing his willingness to pay for the basic rate r_{good} and the other expressing his extra willingness to pay for the extra rate $r_{\text{high}} - r_{\text{good}}$. Of course, in the simple case of a single quality level that was discussed in the previous paragraphs, only one bid is to be submitted per user in each mini-auction.

This generalization is very useful in the context of 3G networks where *hierarchical encoding* of streaming video is common [10]. Under this approach, streaming video is coded in hierarchical streams of predefined quality. The “bottom” stream is the video’s base layer,

which contains all the basic video frames; when this layer is received at the user terminals, it enables users to watch the video at a very low quality. The additional streams provide the enhancement layers that improve the video's quality. This way, terminals with different QoS capabilities can be supported without wasting bandwidth, since only the required layers are transmitted over the network. Moreover, when the Radio Access Network (RAN) is congested, some enhancement layers can be dropped but the users continue receiving video service (though the quality is much poorer). This is reflected in the aforementioned utility function that defines a different user willingness to pay for receiving service under various QoS levels.

Next, we introduce a new utility function for Variable Bit Rate services, in cases where a predefined traffic pattern must be obtained by participating in the auction. This function expresses the utility attained by computing the deviation of the allocated pattern from the desired: $v_{s,i}^{(t)}(x_i^{(1)}, \dots, x_i^{(t)}) = \frac{x_i^{(t)}}{\bar{m}_i} \cdot \frac{U_{s,i}}{K_{s,i}} \cdot \alpha^{m_{\text{gap},i}/\bar{m}_i}$, where: $m_i^{(t)}$ is the target bit-rate at the various time slots; $\bar{m}_i = \frac{1}{K_{s,i}} \sum_{t=1}^{K_{s,i}} m_i^{(t)}$ is the mean rate of user's i target resource allocation pattern; and $\sum m_{\text{gap},i}$ is the sum of the user's target bit-rates at the time-slots of the "on-going gap", that is the consecutive zero-allocation slots since the last slot $\tau < t$ where due to an allocation, or to the fact $m_i^{(\tau)} = 0$, we have $x_i^{(\tau)} = m_i^{(\tau)}$.

We have argued in [4] that a user whose preferences are accurately expressed by one of the predefined functions, has the incentive to *truthfully* declare both this function and his maximum willingness-to-pay $U_{s,i}$. Furthermore, the charge reflects the actual market demand and proper incentives for rational network usage are provided to the users: Receiving service is expensive in periods of high network load, while it is inexpensive when the network is under-utilized.

VII. EXPERIMENTAL ASSESSMENT

We have extended the simulation software of ATHENA [4] so that the HERA auction is simulated. We have run numerous simulation experiments according to a detailed simulation model, specifying: i) the distributions of user arrivals, departures, and service requests, and ii) the mix of users in terms of the number of users per type and the distribution of their total willingness to pay. The total quantity of resource units available at each mini-auction also fluctuates in the simulation model thus imitating the varying allocation of resources to telephone calls and SMS/MMS in a real network. Finally, in each experiment, the value of the discount factor α in user utility functions was in the range $[0.95, 0.995]$.

The quantities demanded by users are in accordance with the demand of services as specified in IST project B-BONE [10]. Hence high quality audio streaming sessions demand 64-128 Kbps, video streaming sessions demand 384 Kbps, telemetry monitoring requires 28 Kbps and non-RT VBR sessions are simulated as periodic on-off sources having a peak rate of 100Kbps. A sample resource allocation pattern of an on-off source is provided in Fig. 4.




Fig. 4. Screenshot depicts the aggregate demand for resources (top-most bars), the total available resource units (line), and the allocation pattern (lowest bars) of an on-off source under low competition in the auction.

A sample resource allocation pattern of a Type 2 user (watching video streaming) is depicted in Fig. 5, while his bids in the auctions he participated are provided in Fig. 6.

A sample resource allocation pattern of a Type 3 user is depicted in Fig. 7, while his bids in the auctions he participated are provided in Fig. 8.

It is seen that the usage of the aforementioned utility functions as bidding functions in the auction, results in meaningful resource allocation patterns for the users. Users with competitive bids are served very satisfactorily while they are allocated patterns that are in very good accordance to their respective preferences. On the other hand, users with non-competitive bids are allocated very limited quantities of resources (if at all) at a low total charge.

"Intermediate" allocation patterns in general arise rarely. (Such patterns are those that are not nearly consistent and yet comprise a non-negligible percentage




Fig. 5. Screenshot depicts the aggregate demand for resources (topmost bars), the total available resource units (line), and the allocation pattern (lowest bars) of a Type 2 session under high competition in the auction.

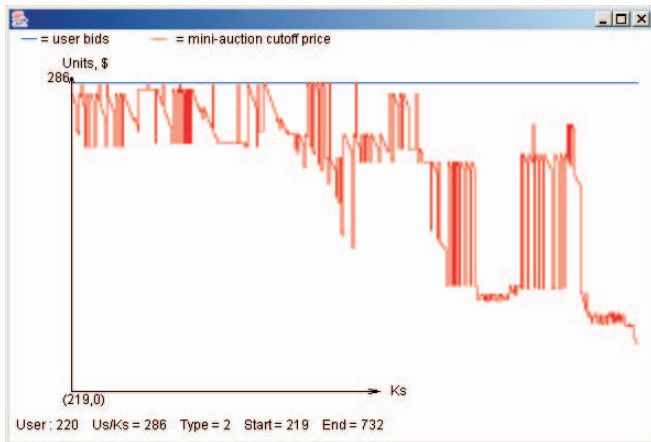


Fig. 6. Screenshot depicts the user's bids (topmost line) and the auction's cutoff price (fluctuating line) under high competition in the auction.

of the total quantity of resources that would arise for the same user in the case of consistent allocation.) A typical distribution of allocation patterns under high demand is depicted in Fig. 9. The total percentage of users being allocated intermediate patterns is very low, due to the effectiveness of the aforementioned utility functions. Indeed, intermediate patterns mostly arise in cases of low market demand, and their associated charge is typically very low. Also, the higher the level of market demand, the larger the percentage of users allocated almost no resources, as depicted by the leftmost bar of Fig. 9. This is also depicted in Fig. 10 where the classification of the user resource allocation patterns is provided in cases of very high competition in the auction.

Hence, our mechanism serves as soft Call Admission Control mechanism, which takes into account both the needs and the economic means of the various contending




Fig. 7. Screenshot depicts the aggregate demand for resources (topmost bars), the total available resource units (line), and the allocation pattern (lowest bars) of a Type 3 session under high competition in the auction.

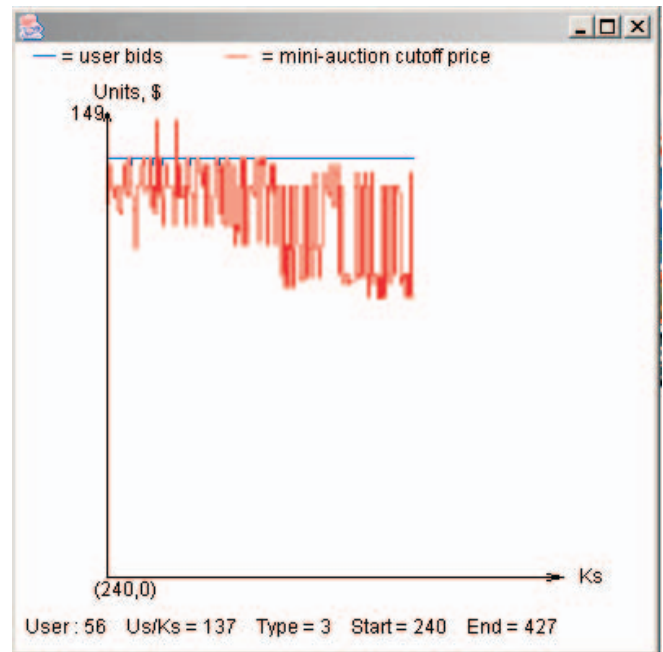


Fig. 8. Screenshot depicts the user's bids (topmost line) and the auction's cutoff price (fluctuating line) under high competition in the auction.

users. Overall, our approach can serve as the resource allocation logic of the HSDPA scheduler in UMTS networks.

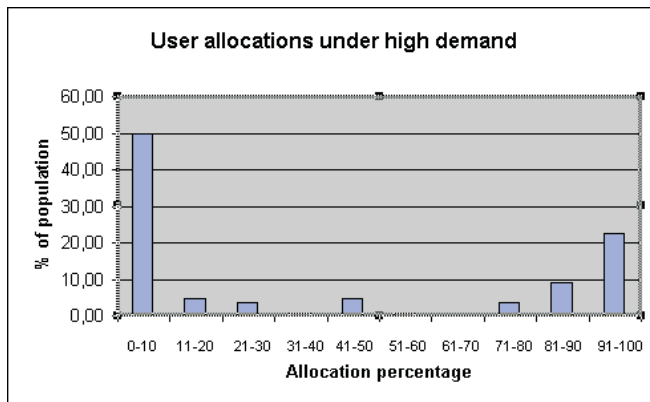


Fig. 9. Distribution of resource allocation patterns with respect to the percentage of the total targeted resources that was attained under high demand.




Fig. 10. Distribution of resource allocation patterns with respect to the percentage of the total targeted resources that was attained under very high demand.

VIII. CONCLUSIONS

In this paper we presented HERA, an effective auction-based mechanism for the UMTS High Speed Downlink Packet Access (HSDPA). The mechanism can deal successfully with long time-scale data, audio and video services in practical cases of networks with large numbers of competing users. It is based on a series of Generalized Vickrey Auctions and a set of user utility functions expressing the user's preferences with respect to the resource allocation patterns. The additive form of these functions enable their usage as bidding functions in the series of the auctions each user participates. Hence, we have developed suitable bidding functions that enable users to construct meaningful resource allocation patterns on the long time scale of user service by

bidding for many short time scale network resources. We have assessed experimentally the effectiveness of our mechanism and it appears that most of the users either are served very satisfactorily or essentially are not served at all.

IX. ACKNOWLEDGEMENTS

This work was partly funded by EU project FP6-IST-507607 BBONE [10].

The authors wish to thank BBONE and EURONGI partners for useful discussions on the subject of this paper.

REFERENCES

- [1] H. Holma and A. Toskala, "WCDMA for UMTS: Radio Access for Third Generation Mobile Communications", John Wiley, ISBN 0-471-72051-8.
- [2] 3GPP, URL: <http://www.3gpp.org/>.
- [3] H. Holma and A. Toskala, "Third Generation WCDMA radio evolution", *Wireless Communications and Mobile Computing*, 3, 978-992, 2003.
- [4] M. Dramitinos, G. D. Stamoulis, C. Courcoubetis, "Auction-based Resource Reservation in 2.5/3G Networks", *Kluwer/ACM Mobile Networks and Applications [MONE] special issue*, accepted 2003, to appear, Vol. 10, No. 2, April 2005. Also available at <http://nes.aueb.gr/>.
- [5] C. Courcoubetis and R. Weber, "Pricing Communication Networks : Economics, Technology and Modelling", John Wiley & Sons, 2003, ISBN 0-470-85130-9.
- [6] V. Krishna, "Auction Theory", Academic Press, April 2002, ISBN 0-12426297-X.
- [7] J. Ramiro-Moreno, K.I. Pedersen, P.E. Mogensen, "Capacity gain of beamforming techniques in a WCDMA system under channelization code constraints", *IEEE Transactions on Wireless Communications*, 3(4), 1199-1208, July 2004.
- [8] M. Frodigh, S. Parkvall, C. Roobol, P. Johansson and P. Larsson, "Future-Generation Wireless Networks", *IEEE Personal Communications*, 8(5), 10-17, October 2001.
- [9] 3GPP, "Technical Specification Group Services and System Aspects, QoS Concept" (3G TR 23.107 version 5.6.0), URL: <http://www.3gpp.org/>.
- [10] FP6-IST-507607 Project B-BONE, "Broadcasting and multicasting over enhanced UMTS mobile broadband networks", URL:<http://b-bone.ptinovacao.pt/>.