

Auction-based Resource Reservation in 2.5/3G Networks

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
Tel: +30 210 8203693, Fax: +30 210 8203686
Email: {mdramit, gstamoul, courcou}@aueb.gr

Abstract

We consider UMTS networks in which users request services lasting for long time intervals, e.g. video clips lasting for several minutes. The duration of network time-slots over which resource units are allocated is much shorter. This renders consistent reservation of the short time scale resources (that would lead to constant service quality) for longer time scales a hard task. In this paper, we define an auction-based mechanism for 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 requirements of the UMTS services valued by the users. The mechanism is based on a series of Generalized Vickrey Auctions and a set of new user utility functions. Our mechanism is provably incentive compatible under certain assumptions, i.e. users have the incentive to bid truthfully according to their utilities. This provides users with the incentives for rational usage of resources according to their actual needs. The utility functions we define express appropriately the preferences of the users with respect to the resource allocation pattern in the cases where perfectly consistent allocation cannot be attained. We also provide a mapping of these functions to the UMTS service classes. The effectiveness of our resource reservation mechanism is demonstrated by means of

experiments. In practical cases, the mechanism should be implemented in the network base station, with bidding being performed automatically on behalf of the users on the basis of their willingness to pay and their selection of one of the predefined utility functions.

Keywords: Auction, efficiency, resource reservation, UMTS, utility.

1 Introduction

Multi-unit auctions have recently received considerable attention as an economic mechanism for resource reservation in networks. The case where users compete for reserving consistently (i.e., without fluctuations) resources for large time scales remains an open research topic. This is of particular interest for many practical cases such as the provision of network services with relatively high duration. A prominent case is that of UMTS [7], [11]: except for voice, users request other services lasting for longer; e.g., news downloading or video streaming of a certain rate. Note that the duration of such sessions varies, depending on the application and the content: the duration of jokes' videos can be one minute (or even less), while that of a football match is 90 minutes. In any case, however, the duration of such sessions is much longer than that of network slots, over which resource units¹

¹The unit of resource allocation and the definition of a time-slot depends on the network technology. In

can be reserved.

In this paper, we deal with how to satisfy the demand for large time scale services, by allocating *nearly consistently* (i.e., without significant fluctuations) resources to those users who value them the most. We assume that the population of users generally varies over time, which further complicates the problem. Our objective is to allocate resources efficiently (i.e. maximize social welfare, or at least attain a high value thereof), while providing proper incentives for rational usage of resources. Efficiency should be taken seriously under consideration by the commercial UMTS network providers. Apart from UMTS, the open problem of consistent resource allocation also applies to GPRS technology including its enhanced version EDGE [9]. Since the mobile industry is extremely competitive, inefficient resource allocation (i.e. awarding some of the scarce network resources to those that do not value them the most) would reduce satisfaction of high-value users. These users may consider migrating to another provider, which in fact is now easier than the past due to number portability. Hence, the revenue of inefficient providers will deteriorate. Moreover, 3G spectrum is a scarce “fixed-quantity” resource, while users’ demand for bandwidth is continuously increasing. Thus, efficient exploitation of spectrum is crucial.

In this paper, we argue that nearly consistent reservation of resources can be attained by means of a series of repeated Generalized Vickrey Auctions, one per slot. Although the repetition of auctions has been analyzed by economists, most of the work focuses on comparisons of sequential and simultaneous auctions and the problems that auction repetition imposes on incen-

UMTS, a 10msec UTRAN frame constitutes a slot, while resources are allocated in bits; see Appendix. In GPRS and EDGE, the unit of resource allocation is the radio block.

tive compatibility (see [2] and [5] for related results). Hence, it remains an open problem how to reserve resources consistently throughout the service duration by means of auctions. Thus, we propose and evaluate a new such mechanism that is applicable for 2.5/3G networks, i.e. networks based on technologies such as UMTS, GPRS and EDGE. For simplicity reasons, we henceforth restrict our discussion to UMTS, but we employ the general terms “slot” and “unit” (of resource allocations) so as to keep presentation of the material as general as possible; see Appendix. Our mechanism is provably incentive compatible under certain assumptions, i.e. users have the incentive to bid truthfully according to their utilities. This provides them with the incentive for rational usage of resources, according to their actual needs. Furthermore, we define utility functions that express appropriately the preferences of the users with respect to the resource allocation pattern in the cases where perfectly consistent allocation cannot be attained. These constitute an integral part of our mechanism for nearly consistent resource reservation, because such reservation is attained when bidding is performed in accordance to these utility functions. In fact, this is also demonstrated experimentally by analyzing the resulting resource allocation patterns. It appears that, for users with competitive bids, the resulting patterns are in very good accordance to their respective preferences, thus showing the effectiveness of our mechanism. Moreover, users with non-competitive bids are allocated very limited quantities of resources (if at all) at a low charge. We also provide a mapping of our utility functions to the UMTS service classes.

An alternative to performing auctions sequentially would be to run a combinatorial auction (or multiple such auctions), thus allowing users to bid for resources spanning several consecutive

slots. Such an approach should not be adopted due to its excessive computational overhead and to the fact that it would be impossible for the network to serve users with higher value who arrive later than those currently served. Hence, such an approach would not be efficient.

The remainder of this paper is organized as follows: In Section 2, we present our auction-based mechanism for resource reservation. In Section 3, we define user utility functions capable to express certain types of user preferences with respect to patterns of inconsistent resource allocation. In Section 4, we discuss the auction’s incentive compatibility properties under certain assumptions. In Section 5, we experimentally assess the effectiveness of our mechanism. Finally, in Section 6, we provide some concluding remarks.




Figure 1: A UMTS cell and its serving users.

2 ATHENA: A new Resource Reservation Mechanism²

Most of the complexity of the UMTS consistent resource reservation problem lies in the fact that users demand sessions spanning partly overlapping intervals with different durations, which in general are much larger than the time scale t_a of a slot. The approach that we propose is

²ATHENA: Auction-based THird gEneration Networks resource reservAtion

to conduct a sequence of “mini-auctions”, each concerning reservation of resources within one slot. Each mini-auction is a sealed-bid Generalized Vickrey Auction (GVA) [6], with atomic bids (i.e., bids that are either fully satisfied or rejected) of the type (p, q) , where p is the expressed willingness to pay for a quantity q of resource units in the present slot. For UMTS, if the service involves traffic of a specific rate m , then we have $q = m \cdot t_a$. As explained in [6], in a GVA, bidders with elastic utility functions 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. Indeed, most often efficiency comes together with incentive compatibility, since it is hard to allocate the goods auctioned to the bidders that value them the most if their bids do not reveal their true values for these goods.

According to the rule of GVA, each user is charged with the social opportunity cost that his presence entails. Note that this charge is less than the user’s bid. We illustrate this below, by means of a simple example: Assume that three units of a good are auctioned to four bidders, namely A, B, C and D . The valuations of bidder A for the three units are $(u_1^A, u_2^A, u_3^A) = (9, 6, 4)$, while those of B are $(u_1^B, u_2^B, u_3^B) = (8, 5, 1)$. Bidders C and D are only interested in obtaining one unit, with $u_1^C = 2$ and $u_1^D = 3$. As already mentioned, bidders have the incentive to bid truthfully. Bids are ranked in decreasing order and A is awarded two units, while B is awarded one unit. The charge for each bidder equals the social opportunity cost his presence entails: i.e., each pays the highest losing bids after excluding his own bids, paying as many such bids as the units he is awarded. Thus, A is charged $5 + 3 = 8$ and B is charged 4.

In Section 4, we argue that it is best for each

user to bid *truthfully* in each mini-auction according to his actual valuation for resources, under certain assumptions. However, in a realistic case of a UMTS network, it is not feasible for users to participate in all these mini-auctions, neither manually nor automatically by means of an agent running in their respective terminal. Thus, since the user cannot give his utility on a per mini-auction basis, we define utility functions, pertaining to the various services. These functions are provided by the network operator for the user to choose from and scaled by the user's total willingness to pay, which is to be given by the user himself (as part of his service request). This is similar to the 3GPP's approach [12] regarding predefined QoS profiles. Then, the network runs all mini-auctions by bidding optimally for each of the users, according to his respective selection of utility function. Thus, all computation is performed on the network base stations, while there is no need for extra communication with the user terminals. The details of the physical layer and the auction are hidden from the users: A user demanding a service selects among the predefined utility functions the one that better expresses his preferences and declares a willingness to pay U_s . A session that lasts for time t_s is then created. Each user aims in achieving constantly the desired rate m by bidding in a large number K_s of mini-auctions, where $K_s = \frac{t_s}{t_a}$. (Recall, however, that the network is bidding on each user's behalf.) For example, if the user wishes to watch his favorite music video clip lasting for 4 minutes, all that he declares is the video name, a total willingness to pay U_s , the desired quality level, and the utility function type, as shown in Figure 1. Requests are sent from the user's terminal to the UMTS base station over the Random Access CHannel (RACH). The parameters t_s , K_s and $m = 2\text{Mbps}$ are computed automatically

by the network and are transparent to the user. Note that a successful choice of U_s is very important for the resources to be allocated to the user. Such a choice should be based on the user's urgency to receive the service, on the extent to which the content to be transported is interesting to him, on a rough estimate of the duration of the service, as well as on the user's past experience from participation in the ATHENA mechanism. It is plausible, that future terminals may be equipped with software providing some related assistance (or recommendation) to the user.

To summarize, our auction-based mechanism ATHENA comprises three layers: i) the Service Layer, where users request services, ii) the Auction Layer, where user-supplied information is mapped to auction-specific parameters, and resources are reserved by means of auctions, and iii) the Network Layer, where the reserved resources are actually allocated.




Figure 2: The layered structure of ATHENA.

Finally, it should be emphasized that by adopting the ATHENA mechanism a UMTS network provider attempts on a slot-by-slot basis to offer nearly consistent resource reservation to bidders with competitive bids. However, at the start of the service, no guarantee on the resources to be allocated throughout its duration can be offered. Nevertheless, by the appropriate choice of a utility function, it is guaranteed that the charge never exceeds the user's value for the specific service under the quality actually attained.

3 User Utility Functions

Next, we define the utility functions offered by the ATHENA mechanism to the users to choose from. We assume that the user's utility (value) u_s for receiving a service is the sum of the sub-utilities $u_i(x_i; h_{i-1})$ obtained in each slot i , each of which depends on the service rate x_i and the history h_{i-1} of resource reservation for this service session up to the present slot; that is,

$$u_s = \sum_{i=1}^{K_s} u_i(x_i; h_{i-1}). \quad (1)$$

The dependence on the history h_{i-1} implies that when there are gaps in the resource allocation pattern, not only the quantity of resources but also the *way* these are allocated can possibly make 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 allocation is not possible. For simplicity, we first consider three cases, in all of which the user is of the guaranteed type, i.e. he is only interested in attaining a certain rate m in each slot. Thus, the network places on his behalf only one bid for a quantity $q = m \cdot t_a$ of resource units. Due to the incentive compatibility property (see Section 4), we adopt the following *bidding* rule: in each slot i , the network bids on behalf of each user truthfully in accordance to his selected utility function, thus offering an amount of money $u_i(m; h_{i-1})$.

Type 1: *Users indifferent to the allocation pattern.* This is the simplest type, and pertains to "volume-oriented" users, i.e. users whose utility solely depends on the quantity of resources allocated, as opposed to the allocation pattern. Hence the user's total value U_s is equally apportioned among the various slots. That is,

$$u_i(x_i; h_{i-1}) = \mathbf{1}(x_i = m) \cdot \frac{U_s}{K_s}. \quad (2)$$

Therefore, the lack of information transfer due to the lack of resource allocation in certain slots, results in *proportional* loss of user satisfaction. The most prominent example of users belonging to this type is those accessing news. 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.

Type 2: *Users sensitive to service continuity.* This type pertains (among other cases) to users that prefer watching consistently half of a football match rather than watching several shorter periods. In general, it is applicable for services where the degree of user satisfaction depends heavily on both the volume of information transmitted and the delay observed. Thus, users of this type prefer the allocation pattern of Figure 3(a) to that of Figure 3(b), although the total quantity of resources of the two patterns is the same. In order to express this preference, we define the sub-utility function as

$$u_i(x_i; h_{i-1}) = \mathbf{1}(x_i = m) \frac{U_s}{K_s} \cdot \alpha^d. \quad (3)$$

where d is the distance between the current and the previous slots during which this user achieved reservations. Recall that h_{i-1} is the history of resource reservation for this service session up to the present slot, and influences $u(x_i; h_{i-1})$ through the value of d , which is kept track of by the ATHENA module. This utility is suitable for the UMTS Streaming Class services [12].

Type 3: *Users sensitive to the smoothness of the allocation pattern.* This type pertains to users of services such as stock-market information and games. Such users prefer allocation pattern Figure 3(b) to that of Figure 3(a). In order to express this preference, we define the sub-utility function as

$$u_i(x_i; h_{i-1}) = \mathbf{1}(x_i = m) \frac{U_s}{K_s} \cdot \alpha^{\max\{0, \Delta d\}}. \quad (4)$$

where Δd is the difference of the present and the previous values of the distance defined above. Note that $\alpha^{\max\{0, \Delta d\}}$ equals 1 if $\Delta d \leq 0$ and thus the received quality of service improves or remains constant, and is less than 1 (since $0 < \alpha < 1$) if the distance increases and hence the quality deteriorates. Hence, this utility function is suitable for users of the UMTS Interactive Class services [12]. This function can be easily extended to cover the case where a user is fully satisfied when he receives a certain periodic pattern, as opposed to perfectly consistent allocation.

Both Type 2 and 3 utility functions have certain features in common: $\frac{U_s}{K_s}$ expresses the additional user satisfaction per slot in cases of reservations in consecutive slots, while α and its powers declare the dis-satisfaction resulting from the gaps in the allocation pattern. In both cases, if the user is consistently allocated resources, thus receiving the best possible quality, then the utility obtained is U_s . As for incentives, they only concern the user's selection of one of the predefined utility functions and the declaration of his total willingness to pay U_s . The incentive compatibility property (see Section 4) shows that a user whose actual preferences are accurately expressed by one of the predefined functions has the incentive to truthfully declare this particular function as well as U_s .




Figure 3: Inconsistent allocation patterns.

The aforementioned utility functions are not the only ones reflecting the user satisfaction w.r.t. the allocation pattern attained. What is impor-

tant, is the fact that the values of these utilities reflect correctly the preferences of each type of user. In our example, it can be easily seen that for a user of Type 2 the utility value obtained from equation (3) exceeds that from equation (4) despite the fact that the same total quantity of resources is allocated in the two cases. For a user of Type 3, the reverse inequality applies. Experimental assessment reveals the effectiveness of the utility functions presented above with respect to the resulting resource allocation patterns; see Section 5.

So far we have defined utility functions suitable for guaranteed users, i.e. users whose service rate has a single target-value. Next, we extend the definitions of these utility functions to cases of multiple such target-values. For brevity, we only deal with extension of the utility function of Type 2; Type 3 can be treated similarly.

Type 4: Extension of Type 2 for two alternative service rates. We consider a user who is willing to watch video either with the “high quality” rate r_{high} , or with just the “good quality” rate r_{good} , whenever the former 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 V_{good} and $V_{\text{high}} = V_{\text{good}} + \Delta V$, where $\Delta V > 0$. Figure 4 depicts a typical allocation pattern arising when consistent resource allocation with rate r_{high} is not possible. This pattern can be viewed as the superposition of the r_{good} and $r_{\text{high}} - r_{\text{good}}$ allocation sub-patterns. Since the user is still sensitive to service continuity, but there are now two possible service rates, we can extend the user utility function of equation (3) as follows:

$$u(x_i; h_{i-1}) = \mathbf{1}(x_i \geq r_{\text{good}}) \frac{V_{\text{good}}}{K_s} \cdot \alpha^{d_1} + \mathbf{1}(x_i = r_{\text{high}}) \frac{\Delta V}{K_s} \cdot \alpha^{d_2}, \quad (5)$$

where d_1 and d_2 are defined with respect to the length of the gaps incurred in the r_{good} and $r_{\text{high}} - r_{\text{good}}$ sub-patterns, respectively. The type of service considered above corresponds to the UMTS Conversational Class [12], with r_{good} and r_{high} corresponding to the Guaranteed and the Maximum Bitrates of this class. However the lower rate r_{good} is not guaranteed by ATHENA.




Figure 4: A user accepting two alternative rates of service

The number of atomic bids to be given on behalf of each user in each mini-auction equals the number of alternative quality levels. Thus, for each user of Type 4, two *summable* bids should be placed: 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 $r_{\text{high}} - r_{\text{good}}$.

Finally, in the complete version of the paper [1], we show that a super-additivity property applies when superimposing resource allocation patterns. That is, the utility of the resulting pattern exceeds the sum of the utilities of the constituent patterns, when all utilities are calculated by means of the functions presented in this section. This property is in agreement with the super-additivity of the actual user utilities, due to the complementarities arising when additional resources are allocated.

4 Incentive Compatibility

A user wishing to receive service has to participate in K_s successive Generalized Vickrey Auctions (GVAs) of the ATHENA mechanism. We assume that the various bidders have private values for the resources sought, and that the total utility of each of them is the sum of the sub-utilities attained from each reservation; see equation (1). We restrict attention to users of the guaranteed type, i.e. users only interested in attaining a certain service rate m in each slot.

When these sub-utilities are independent, as in the case of users of Type 1 (see Section 3), then the value associated with a slot is independent of the overall allocation pattern. Then, by the incentive compatibility property of a GVA, it is best for bidders to bid truthfully an amount of money equal to the corresponding sub-utility $u_i(m; h_{i-1})$; i.e., only thus their net benefit can be maximized. On the contrary, when these sub-utilities are dependent, the reservation of resources in one slot brings both instant value (sub-utility) to the additive user utility function as well as extra value to resources to be reserved in the future by improving the overall allocation pattern. [This improvement amounts to a reduction of the values of d and Δd in equations (3) and (4).] It is easily proved that bidding less than the corresponding sub-utility is less beneficial than bidding truthfully, because this only increases the risk of losing without affecting the payment. It could be argued however that a bid higher than the sub-utility should be submitted, due to the extra value that a present resource reservation (which becomes more likely by over-bidding) would bring to future ones. However, this is not the case for “extremely uncertainty averse” (i.e. extremely conservative) bidders, who (by definition) in cases of choice/behavior under uncertainty maximize the payoff of the worst possible outcome. This “max-min” be-

havior was proposed by [10] for situations of complete uncertainty, and was adopted by several other researchers (e.g. see [2] and [8]). In our case, extremely uncertainty averse bidders prefer to sacrifice the potential for increased future value that is attained by means of overbidding for the security of avoiding a negative net benefit in the present slot, again due to overbidding.

To summarize, *incentive compatibility* holds for private values, additive user utility functions and either independence of the values of the utility function per slot or extreme uncertainty aversion. A detailed proof of the aforementioned properties can be found in the complete version of the paper [1]. The proof employs techniques on sequential games.

5 Experimental Assessment

In this section we assess experimentally the resource allocation patterns arising in a UMTS network employing the ATHENA mechanism. For simplicity, we do not consider users of Type 4, and restrict attention to the patterns of users of Types 2 and 3 (see Section 3); recall also that the resource allocation pattern does not really matter for users of Type 1. We have implemented special software that simulates the ATHENA mechanism. We have run numerous simulation experiments according to a detailed simulation model, specifying the distributions of user arrivals, departures, and service requests, and the mix of users in terms of the number of users per type and the distribution of their willingness to pay. For example, a specific scenario ran comprised 500 mini-auctions, 50 “low-value” users of Type 1, 25 “high-value” and 25 “low-value” users of Type 2 and 50 “medium-value” users of Type 3. For each user, the willingness to pay is randomly selected according to a uniform distribution over an interval which is determined

by whether the user is of low, medium or high value.




Figure 5: Screenshot (a) depicts the total available resource units (highest bars), the aggregate demand for resources (line), and the allocation pattern of a user of Type 2 (lowest bars) whose bids are in general competitive. Screenshot (b) depicts this user’s bids (darker line) and the mini-auctions’ cut-off prices throughout his service session.

The entry and exit times of each user in the network are also selected randomly, according to a uniform distribution over the interval $[1, K_s]$. The total quantity of resource units available at each mini-auction also fluctuates (due to the varying allocation of resources to phone calls

and SMS/MMS) in the simulation model, and is randomly selected according to a uniform distribution. Finally, note that the values of the discount factor α in user utility functions [see equations (3) and (4)] belong to the range $[0.91, 0.99]$.




Figure 6: Screenshot (a) depicts the total available resource units (highest bars), the aggregate demand for resources (line), and the allocation pattern (lowest bars) of a user of Type 3 whose bids are in general competitive. Screenshot (b) depicts this user’s bids (darker line) and the mini-auctions’ cut-off prices throughout his service session.

Experiments confirm that, as intuitively expected, when a user’s bids are always higher (resp. lower) than each mini-auction’s lowest

winning bid (“cut-off” price), then the allocation pattern was perfectly consistent (resp. the user was allocated no resources at all). Thus, ATHENA serves as a “soft” call admission control mechanism actually preventing users with low willingness to pay from receiving service by the network. The more interesting patterns arise for users whose bids are often close to the cut-off prices of the various mini-auctions, thus being vulnerable to entries of new users and to the fluctuations of the total quantity of resources available. Experiments showed the effectiveness of our mechanism for such users. Indeed, most of them either are allocated very limited total quantities of resources at a low total charge, or succeed to receive meaningful service with a nearly consistent resource allocation pattern at a fair total charge. In particular, successful users of Type 2 receive almost consistent allocation of resources for large time intervals (see screenshots in Figure 5) while those of Type 3 receive fragmented allocation patterns, similar to those depicted in Figure 6. These outcomes are due to the expressions of the utility functions employed, in accordance to which are the gaps observed. In several cases, when a gap starts in the resource allocation pattern, then the bids of the user are reduced due to the discount factor α , and if he is not competitive enough then essentially he quits service. However, this is preferable than receiving mediocre service with large gaps in the resource allocation pattern.

6 Concluding Remarks

In this paper we presented ATHENA, an effective auction-based mechanism for nearly consistent reservation of the resources of a UMTS network by the users that value them the most. The mechanism deals with long time scale data, audio and video services, and is based on a series of Generalized Vickrey Auctions and a set of

new user utility functions regarding the resource allocation patterns. We also provide a mapping of these functions to the UMTS service classes.

In [1] we have extended the mechanism so as to also guarantee resources for the entire duration of a service for a few high-value users, which should pay a related premium. An interesting topic would be to combine our resource reservation mechanism with DiffServ. This issue is left for further research.

References

- [1] M. Dramitinos, G. D. Stamoulis, and C. Courcoubetis, "Auction-based Resource Reservation in 2.5/3G Networks", URL: <http://nes.aueb.gr>
- [2] J. Dow and S. R. da Costa Werlang, "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio", *Econometrica*, 60:1, 197-204, January 1992.
- [5] D. Friedman, "Evolutionary Games in Economics", *Econometrica*, 59:3, 637-666, May 1991.
- [6] T. Groves and J. Ledyard, "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem", *Econometrica*, 45:4, 85-96, May 1997.
- [7] H. Holma and A. Toskala, "WCDMA for UMTS: Radio Access for Third Generation Mobile Communications", John Wiley, ISBN 0-471-72051-8, September 2000.
- [8] J. Rawls, "A Theory of Justice", Cambridge: Harvard University Press, 1971.
- [9] S. Soursos, C. Courcoubetis, and G.C. Polyzos, "Differentiated Services in the GPRS Wireless Access Environment", Proc. of IWDC 2001, Evolutionary Trends of the Internet, September 17-20, 2001, Italy.

[10] A. Wald, *Statistical Decision Functions*, New York: John Wiley, 1950.

[11] The 3rd Generation Partnership Project (3GPP), URL: <http://www.3gpp.org/>

[12] 3GPP, Technical Specification Group Services and System Aspects, QoS Concept (3G TR 23.107 version 5.6.0).

Appendix: UMTS and GPRS Issues

In this Appendix, we define the "units" of allocation, thus the "slots" of the Network Layer of our auction-based resource reservation mechanism. More information can be found at [7], [9], [11] and [12]. Figure 7 illustrates the 10msec WCDMA frame in UMTS. The entity that allocates the network resources within each frame among user connections is the *resource manager*.




Figure 7: Allocation of resources in UMTS [7].

The ATHENA mechanism is also applicable to GPRS and its enhanced version EDGE. In GPRS, user data services compete for reserving a number of *time slots*, with the unit of allocation being the *radio block*. Users attempt to reserve these resources in order to accommodate their large time scale data traffic; this is transferred in bursts, by means of creating Temporary Block Flows within the PDP context, as described in [9].