EFFICIENT AGENT-BASED NEGOTIATION FOR TELECOMMUNICATIONS SERVICES *

George D. Stamoulis, Dimitrios Kalopsikakis, Anna Kyrikoglou,
and Costas Courcoubetis
Institute of Computer Science (ICS)
Foundation for Research and Technology - Hellas (FORTH),
and
Department of Computer Science, University of Crete,
P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece
Tel: +30 81 391 728, Fax: +30 81 391 601
e-mail: {gstamoul, kalopsik, kirik, courcou}@ics.forth.gr

Abstract

We investigate the usage and functionality of an intelligent software agent negotiating for telecommunications services on behalf of the user. The objective of this agent is to select the best of the service combinations offered by multiple service retailers, taking into account both user preferences as encoded in the user profile and the corresponding charges given by the retailers. The agent analyzed is based on a simple yet informative model for on-line calculation of the utility of the user for the various service combinations, and on an algorithm for fast and efficient negotiation, which is not based on any assumptions regarding the structure of retailer tariffs. We assess our approach both analytically and experimentally, and we argue that the negotiation algorithm is either optimal or close to this for a variety of interesting cases. Although our approach is illustrated by dealing with a particular News-on-Demand service, it is broadly applicable. We also argue that it pertains to both Internet and ATM network technologies.

1. Introduction

The picture in the telecommunications market is changing dramatically; competition among service retailers (which in turn sell services offered by other providers) is increasing due to both technological and regulatory changes. Moreover, the wide diversity and the level of customization of services have now made it possible for users to buy affordable services matching well their requirements and preferences. Although it is beneficial for a user to optimize this selection, based also on his own preferences, this may prove a very complicated task in case there are several "degrees of freedom" involved in this selection (such as media selection, QoS levels etc.). On the other hand, this task is appropriate for an intelligent software agent that will negotiate with service retailers on behalf of the user.

In this paper, we develop such an agent and analyze its performance. In particular, this agent negotiates for a News-on-Demand service under three alternative optimization criteria set by the user; that is, utility maximization, or charge minimization, or net benefit maximization, which proves to be the most interesting case. The agent makes use of a simple yet informative model for user utility and a negotiation algorithm that appears to be optimal or near-optimal in a wide variety of cases. The algorithm is not based on any assumptions regarding the tariff structure. Both the utility model and the negotiation algorithm are simple and computationally efficient. Thus, the agent can be implemented as a *mobile* agent that migrates to retailers' hosts in order to negotiate locally. Such an implementation has been adopted by the ACTS project MONTAGE (AC325) [9], funded by the European Commission, which aims to exploit agent technology to support efficient service provision to fixed and mobile users in competitive telecommunications environments. In the course of MONTAGE project, our approach is applied in assisting mobile users to negotiate with retailers for a News-on-Demand service when located in visited domains. Finally, it should be noted that our approach pertains to both Internet and ATM technologies.

Agent-based software components have recently become very popular in several tasks arising in telecommunications, such as network management, resource allocation, support of mobility etc; see [5] for an overview of the main related issues. Extensive work is performed within CLIMATE, which is the cluster of agent-related ACTS projects; see [1] and references therein. Finally, agents purchasing goods or negotiating on behalf of users have been proposed for several applications; e.g. for Web browsing, travel assistance and

* This work was partly supported by the European Commission under ACTS project MONTAGE (AC325).

shopping [12], for negotiating during the provision of a dynamically priced ABR ATM service [2], etc. Such agents are often based on AI techniques. Our work is conceptually simpler, since it only involves optimization of some scalar function carried out by means of a non-exhaustive negotiation algorithm. In fact, the theoretical and experimental results suggest that this algorithm terminates in a few steps and has in general good performance. It is also worth noting that our approach for user utility approximation and agent-based negotiation can be applied to other kinds of services that can be represented as a vector of options the exact values of which have to be specified optimally for the user.

The remainder of the paper is organized as follows: In Section 2, we describe the News-on-Demand service. In Section 3, we develop the model for user utility, while, in Section 4, we present the negotiation algorithm. Furthermore, we assess the performance of this algorithm analytically (in Section 5), and experimentally (in Section 6). In Section 7, we argue that our approach is applicable to a wide variety of network technologies (Internet, ATM, etc.). In Section 8, we compare our approach with related work. In Section 9, we present some concluding remarks.

2. The News-on-Demand Service

The approach developed in this paper is applicable to a multitude of services requiring the selection of media and QoS levels, or offering other options for which the related choices should match user preferences. Such preferences are encoded in the *user profile*, which is set by asking appropriate questions to the user when subscribing to the service. To make our discussion more concrete, we assume that a user is offered a News-on-Demand service (in the form of news clips) provided by multiple retailers. Each time the user wants to have access to this service, he has to make selections regarding the "ingredients" of the service combination, namely the constituent media and their respective QoS levels. The selection is made by means of a User Agent (UA) that negotiates with retailers, querying them either one-by-one or in parallel; in the latter case, multiple replicas of the UA have to be created, one per retailer with whom a negotiation will take place. The UA (or its replicas) can be implemented as a mobile agent, but our approach also applies to the non-mobile case too. The final selection of service combination is made after taking into account the outcome of the negotiations with all retailers.

First, we discuss user preferences regarding the content. It is assumed that the user profile comprises the content sections of interest to the user, and the associated percentages expressing their relative significance to the user. For simplicity, we always consider that the content sections chosen by the user are available by the retailer at all valid QoS levels.

There are many different ways of presenting each content section, with each of them corresponding to a selection of a set of media and their QoS levels. These media are taken to be *text, still pictures, audio* and *video*. (We assume that *video* is always accompanied by its own audio, which is not taken as a separate medium; *audio* is assumed to have different content.) A fully-specified set of media and the corresponding QoS levels (for each content section) is referred to as service combination. Furthermore, note that for each content section, there is a level of preference associated with each medium, expressing how important it is for the user. Four possible levels of preference are considered: *required, desired, excluded, don't care*. The *required* and *excluded* levels suggest non-negotiable requirements, limiting the set of acceptable service combinations, while the *desired* and *don't care* suggest flexible requirements.

Furthermore, we discuss the selection of QoS levels for the different media. Instead of considering the specific values of the parameters determining the QoS levels offered, we resort to a "high-level" description. In particular, for all the media we refer to the QoS levels as: *excellent*, *very good*, *good* and *low* quality. (Note that text is an exception, because we only need to consider a single QoS level.) Moreover, we assume that each retailer applies a straightforward rule mapping these QoS levels to the appropriate encoding parameters (such as resolution, sampling rate, etc.). These are subsequently translated to traffic parameters (e.g. mean rate or bandwidth, peak rate, loss probability) for the connections to be used. (Further discussion on this issue is beyond the scope of this paper.) The choice (by the retailer) of the cheapest appropriate connections is based on these parameters, as well as on the anticipated traffic characteristics.

For each content section, the user profile includes the levels of preference for the various media, as well as the requirements for the minimum and the maximum acceptable QoS levels. The latter may either be *excellent*, or lower if the user terminal imposes technological restrictions. An example of a user profile is shown in Table 2-1. Notice that for content sections with significance of 0% we do not have to specify preferences regarding the media composition.

Obviously, the most pleasing service combination is the most demanding one (to be referred to as the *maximal*

service combination S_{max}). That is, the one that contains for each content section all the media except the *excluded* ones; moreover, the QoS level for each medium is the maximum, as specified in the profile. For reasons to become apparent in the next section, the user profile also includes the maximum amount of money W_{max} the user is willing to pay for the maximal service combination. The user profile also includes the selection criterion of the user, which is further discussed in Section 3.

Content Selection			
Content Section	*Significance Percentage (k_c)*		
Politics	70%		
Sports	30%		
Weather	0%		
Media specification for Politics			
Medium	*Level of Pref*	*Max QoS*	*Min QoS*
Video	Required	Excellent	Good
Still Pict	Desired	Excellent	Very Good
Audio	Excluded	--------	----------
Text	Don't Care	One QoS Level	
Media specification for Sports			
Medium	*Level of Pref*	*Max QoS*	*Min QoS*
Video	Required	Excellent	Low
Still Pict	Desired	Very Good	Good
Audio	Don't Care	Good	Low
Text	Required	One QoS Level	
W_{max} = $ 10			
Selection Criterion: Net Benefit Maximization			

Table 2-1: An example of user profile.

3. Modeling User Utility

Introduction - The ideal utility function

In order to assess the satisfaction of the user from the various service combinations, we employ the notion of user utility function, which is fundamental in microeconomic theory [13]. Moreover, we consider three optimization criteria for the selection of service combination: *Utility Maximization* (i.e., selection of the most pleasing service combination); *Charge Minimization* (i.e., selection of the cheapest service combination); and *Net Benefit Maximization*, where the objective of the user is to select the service combination S and the respective retailer in order to maximize:

$$Net_benefit(S) = utility(S) - charge(S),$$

$$\text{subject to } Net_benefit(S) \geq 0.$$

This criterion expresses a trade-off between user satisfaction and charge, while also assures that with the selected service combination the user is better-off than if not being offered the service at all.

Ideally, in order to obtain the utility function accurately, the user should associate a value to each service combination, thus expressing his preference, or provide a list of all possible service combinations ordered according to his preferences. Such a list is hard to create for several practical reasons. The number of possible combinations may be prohibitively large. Also, if the UA is implemented as a mobile agent, then the cost and the delay for the transport of the utility function can be considerable. Extensibility is yet another problem that arises when more features are added to the set of already existing ones.

An approximation of the utility function

Below, we develop a simple weighted-sum utility model. This model combines certain parameters, the values of which are determined by means of specific questions to the user (asked when setting the profile). These terms are then added, which implies that (for simplicity), we are ignoring complementarities or synergies between media. Our approach has some similarities with that taken independently in [4]; see Section 8 for a comparison.

Let us first consider a single content section c; this is presented by means of the media in the set M_c, which is a subset of the set of non-excluded media for this section. The utility value u_c associated with this content section is computed as a weighted sum of the individual utilities of the constituent media; for each medium m, this utility value depends on the respective QoS levels q_m and is denoted as $u_m(q_m)$. The weight ρ_m^c of medium m for this summation depends on the medium's relative significance (to the user) for the presentation of the content section c. Therefore, the utility associated with this content section is

$$u_c = \sum_{m \in M_c} \rho_m^c u_m(q_m) \ .$$

To compute the total utility of a service combination S, we sum these values over all content sections; the weight of a content section in this summation equals its significance percentage k_c given by the user in the profile. Thus, we obtain

$$u(S) = \sum_c k_c u_c = \sum_c k_c \sum_{m \in M_c} \rho_m^c u_m(q_m), \qquad (1)$$

where the summation extends over all content sections c with $k_c \neq 0$.

It only remains to specify appropriate values for $u_m(q_m)$ and the weights ρ_m^c. In principle, these can be assessed for

each particular user by asking related questions when setting the user profile. Instead of this, we propose the following simpler approach. Suppose that all users had the same level of preference (say *required*) for all media used in the presentation of a content section *c*. Then, we could set the same values for ρ_m^c for all users. These would correspond to the relative utility r_m^c per medium for the "average" user, and can be obtained by means of user studies and statistics [8]. Based on the information therein and everyday experience, one can form for each content section a table such as the one below, which provides a reasonable set of values for relative utilities for the sports news. (Further investigation of this issue is beyond the scope of the paper.) Without loss of generality we can assume that the r_m^c are normalized, i.e. $\Sigma_{m \in M_c} r_m^c = 1$.

Medium	Relative utility
Video	70 %
Audio	15 %
Still Picture	5 %
Text	10 %

We can combine the weights of the "average" user with the user's levels of preference as expressed in his profile. Thus, we set $\rho_m^c = r_m^c$ for a *required* medium, $\rho_m^c = (2/3)*r_m^c$ for a *desired* medium, $\rho_m^c = (1/3)*r_m^c$ for a *don't care* medium, and finally we re-normalize the ρ_m^c.

Furthermore, we study the impact of QoS level on the utility per medium. Below, we study the dependence of utility on bandwidth, which is taken as a proxy for QoS. In particular, it is assumed that there are two classes of utility function depending on whether the service is elastic (best effort) or guaranteed. For elastic services, the shape of the utility function is approximated as increasing and concave, while for guaranteed services it is similar to a step function that has a convex and a concave segment (Figure 1); see [14]. The concavity of the utility function for elastic services (e.g., Web like services) is justified by the fact that the increment of the user satisfaction is decreasing as the QoS (bandwidth) is increasing. The step form of the utility function for guaranteed services (e.g., video-conferencing) is justified by the fact that there is a threshold below which the user has very limited satisfaction (convex segment) and above which his satisfaction rapidly reaches an asymptotic value (concave segment). Note that in both cases of utility function, the "acceptable" part is increasing and concave. The utility of the maximum QoS level is set to 1 for all the media. The utility of lower levels is decreasing as QoS deteriorates.

For simplicity, we take the same values of the parameters $u_m(q_m)$ for all users and for all media. These are defined by selecting appropriate points from the curves of Figure 1, so that the minimum level has half the utility of the maximum level.

Since we are interested in computing the net benefit of the user, we have to convert utility into *monetary* units. This is accomplished by taking into account the willingness-to-pay W_{max} for the maximal service combination S_{max} as specified in the user profile. Indeed, expressing utility in monetary units, we have:

$$utility(S) = utility(S_{max}) * \frac{u(S)}{u(S_{max})} = W_{max} * \frac{u(S)}{u(S_{max})}$$

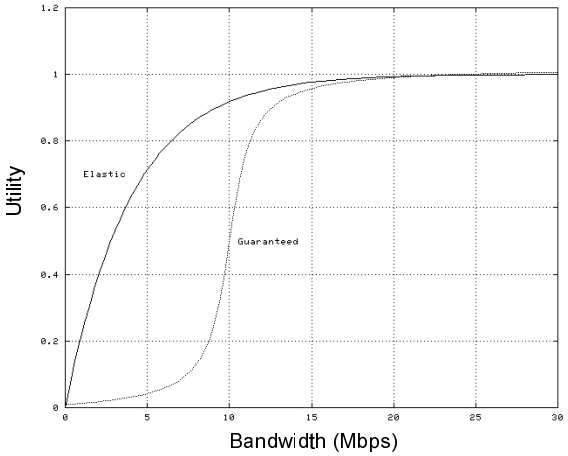

Figure 1: Utility for elastic and guaranteed services as a function of bandwidth (the values of which should be treated as an example).

It is easily seen that $u(S_{max}) = 1$, due to normalization of the parameters in (1). Thus, we finally obtain

$$utility(S) = W_{max} * u(S).$$

The above conversion of utility into money also justifies the constraint that the selected service combination should have a non-negative net benefit.

4. Agent Intelligence for Negotiation

In this section, we propose an algorithm for agent-based selection of service combination. Note that in order to perform the negotiations leading to this selection, the UA employs the user profile. This is also copied to by all UA-replicas in case parallel negotiation with multiple retailers will take place. In such a case, the final selection is made by the parent UA, by comparing the selections of the

replicas. For simplicity, we deal hereafter only with the case of a single retailer.

Utility Maximization and Charge Minimization

In the case of *Utility Maximization* or *Charge Minimization*, it is simple for the UA to make the optimal selection. Consider first the case of *Utility Maximization*. The UA should get an offer only for the *maximal* service combination S_{max} (defined in Section 2), and accept it if the charge is less than W_{max}. Similarly, for the case of *Charge Minimization*, if we assume that tariffs are monotonic, the cheapest service combination is the least demanding one. That is, the one that contains only the media *required* by the user; moreover, the QoS level for each such medium would be the minimum acceptable. Thus, the UA has only to get an offer for this service combination, and accept it if the charge is less than W_{max}.

The algorithm for Net Benefit Maximization

The selection is more complicated and interesting when *Net Benefit Maximization* is the selection criterion. A straightforward optimal algorithm for net benefit maximization would be an exhaustive one. In particular, the UA would ask the retailer for the charge of all possible service combinations, calculate the utility of the user for each one (according to the formula presented in Section 3) and finally compute the net benefits. Then, the UA would return the combination that maximizes the net benefit if the associated value is non-negative. However, the overhead for the UA to find the optimal selection may not be acceptable, because the number of service combinations may be very large.

The algorithm that we propose finds a "good" service combination (often the optimal one) in a few steps. In fact, as discussed in Section 5, the algorithm has certain provable properties regarding the relation of its selection to the optimal one.

The UA asks the retailer one-by-one questions regarding the charge of fully specified service combinations. The algorithm starts from the maximal service combination S_{max} and moves to "neighbor" combinations, step-by-step, as long as the computed net benefit increases. The algorithm stops if the computed net benefit does not improve locally. At each step of the algorithm, the utility can only decrease. In particular, the only permissible action is either to decrease the QoS of a certain medium by one level or to exclude a certain medium from the combination. (Of course, only *desired* and *don't care* media can be excluded.) Note that the neighbors are considered in decreasing order of utility reduction.

Additional speed-up of the algorithm can be attained as follows: First, we can ignore neighbors that do not cause a significant increment of the net benefit. Moreover, we can compute (from past experience) a target value of net benefit that is considered as satisfactory: the algorithm can stop when this value is attained and return the current service combination instead of searching for a better one.

Regarding the computational overhead of the algorithm, note that the utility of a neighbor can be easily calculated *on-line*, by simply adjusting the utility of the parameter (QoS or medium) modified. Although the utility of multiple neighbor combinations may have to be evaluated prior to asking for an offer, this is done by the UA *itself*, without requiring additional negotiation with the retailer.

The approach presented above is robust, in the sense that it is not based on any assumptions regarding the structure of tariffs. Also note that if the total charge equals the sum of charges per medium, then the maximum net benefit can be attained by selecting the QoS level for each medium separately so as to maximize the net benefit per medium, and then find the best combination of media. This approach is restrictive, since it requires that the UA knows how the retailer charges. It will be seen in Sections 5 and 6 that our algorithm proves to be efficient in the case of additive tariffs.

5. Theoretical Assessment of the Performance of the Algorithm

In this section, we state certain properties of the algorithm for *Net Benefit Maximization* as regards the relation between the service combination it selects and the optimal one. We refrain from providing the proofs, which will be included in a next version of the paper.

First, we define an ordering relation between service combinations as follows: We say that service combination x *dominates* y (denoted as $x \geq y$) if, for each content section, x comprises all media of y (for the same content section) at QoS levels that are no worse than the corresponding QoS levels of y.

We consider the case where tariffs have both of the following properties: If $x \geq y$, then

- a medium (comprised in neither x nor y) costs more to add to y rather than to add it to x.
- a medium that is in the same level at x and y costs more to upgrade its QoS at y rather than at x.

These properties are often valid in real market environments, where the same extra or improved feature costs less if the user has already paid more.

It can be established that if all media are either *required* or *excluded*, then the service combination selected by the algorithm *dominates* the optimal one. That is, the algorithm gives a service combination with higher utility and higher charge than the optimal.

One case where the above properties of tariffs applies is when tariffs are *additive* over media and content sections; i.e.,

$$charge(S) = \sum_{c} \sum_{m \in M_c} f_m^c(q_m)$$

where $f_m^c(q_m)$ is the charge for medium m of content section c when the QoS level is q_m. However, in an interesting special case, a stronger result applies. It can be established that if the function $f_m^c(q_m)$ is *increasing* and *concave*, i.e. if QoS upgrading costs more at the lower levels (which is also a reasonable assumption), then the selection of the UA negotiation algorithm is the *optimal* one. In fact, this result applies to all cases where the net benefit associated with each medium is a bitonic function of QoS; that is, it consists of two strictly monotonic segments one of which is increasing and the other is decreasing.

6. Experimental Assessment of the Performance of the Algorithm

To evaluate the efficiency and accuracy of the selection of the UA negotiation algorithm versus the optimal selection, we developed (and implemented in PERL) an environment for simulation tests. (Obviously, it is necessary to run such tests only when the service selection criterion is *Net Benefit Maximization*.) We have tested the performance of the algorithm for different retailer tariffs, all of which were of the following type:

$$Charge = Fixed\ Charge + Transport\ Charge\ ,$$

where the Transport Charge is computed in an additive manner (over the constituent media), and the contribution per medium depends on its QoS level (see also Section 5). We also provided the option to apply a non-fixed discount over the total transport charge computed; thus, we could perform tests with non-additive tariffs. In our experiments, we have tried several formulae for the transport charge per medium, combined with various functions for the discount. In fact some of the parameters were at cases chosen randomly so as to make the tests less predictable. For each case of tariff, we have compared the UA selection to the optimal one (derived by means of an exhaustive search) for all valid user profiles of the following instance of the News-on-Demand service: the user only asks for one content section, always excluding text; the number of valid service profiles is 2196.

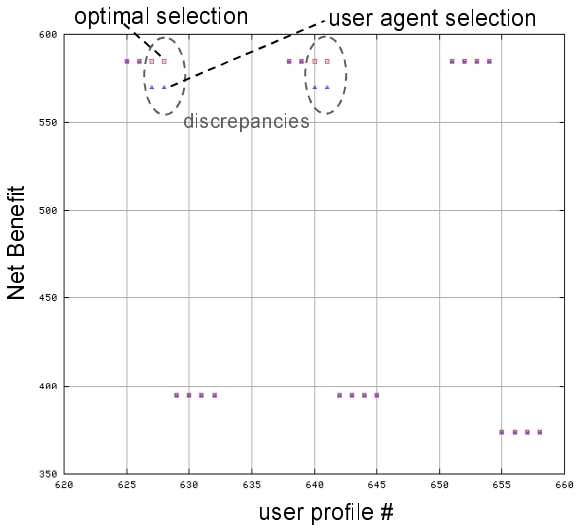

Figure 2: Comparing the net benefit of the optimal and the UA algorithm selections.

In general there is excellent agreement between the UA selection and the optimal one, even in cases where the net benefit is neither additive over media nor bitonic over QoS. Such a case is depicted in Figure 2, where we compare the net benefit of the optimal and UA algorithm selections for a set of 35 service profiles. This set was selected so that it contains 4 cases where the two selections are not the same; nevertheless, the difference in the net benefit value is very small. In fact, for this particular case of tariff, a total of just 16 slight mismatches was observed over the set of 2196 service profiles. Of course, when the net benefit is additive over media and bitonic over QoS, then perfect agreement was observed in the experiments, as expected due to the theoretical results of Section 5.

7. Application of the Proposed Approach to Different Network Technologies

Application of our approach presupposes a concrete presentation of the service requirements (in terms of media combination and QoS levels) and an appropriate mapping thereof to network level parameters. If the underlying network infrastructure is plain Internet, then this mapping only concerns bandwidth. Things are more interesting

when the underlying technology supports different QoS levels.

Furthermore, our approach is directly applicable with the complete protocol stack of ATM. Indeed, the QoS levels of different media can be mapped to network QoS parameters that can be met by either guaranteed or elastic transport services. For example, *video* with a certain QoS level could be satisfied using either CBR, or VBR (with a particular traffic contract) or even ABR (e.g., for the case of *low* QoS, provided that congestion in the ABR class is not considerable).

We now consider the case of Internet as the underlying infrastructure. In the current IP technology (IPv4) all the provided services are "best-effort". Consequently, no guaranteed services can be provided and no QoS differentiation can be supported. The difference in QoS only stems from the difference in encoding quality, and the optimal selection is non-trivial to find only in case of volume-based Internet charging (see [10]); in such a case our approach can offer a "good" selection of the service combination. In fact, some extensions could adapt the proposed model to better perform in an IP infrastructure. The UA could be enriched with a functionality of *estimating* the current state of the network, as this influences the perceived quality and thus, the utility of different transmission rates. This estimation could be done by sending some testing packets to the retailer, as in [4].

Several approaches have been recently proposed in order to enrich Internet with QoS provision. RSVP [11] accompanied with *Internet Integrated Services* [6] can offer the appropriate mechanisms for guaranteed integrated services. Also, IPv6 [7] supports flow discrimination and priorities, some of which are reserved for real-time traffic. This separate treatment of flows of information would meet the demand for different QoS levels of the provided service. Thus, similarly, with the case of ATM, our approach would also be applicable in the aforementioned cases.

Finally, we consider the case of *differentiated services* [3]. In this technology, various categories of services are created, consisting mainly of different conditioning functions and per-hop-behaviors. This distinction does not imply strict guaranteed services but can accommodate different user expectations. The differentiated classes could be used to accommodate QoS levels, thus implying that our approach is again applicable.

8. Comparison of the Proposed Approach with Related Work

We found that there are certain common ideas of our approach for agent-based negotiation with the WebTP project [4]. WebTP is a receiver-driven Web Transport Protocol that optimizes web transfers to ensure maximum user satisfaction, which depends on its arrival time and the acceptable delay of the various objects. In particular, both approaches encode the user preferences in the user utility function, that is a weighted sum of the utilities of the individual media (objects). In our case, if the selection criterion is *Net Benefit Maximization*, then there is a need for further negotiation with the retailer during the selection procedure in order to get the charge offers for different combinations. Furthermore, the user agent employs the algorithm of Section 4 to make a good selection. On the other hand, WebTP project examines the problem of optimizing the transfer order and uses the evaluation of all the possible orderings. Finally, our approach is applicable to both best-effort transport services and services with guaranteed QoS, as opposed to WebTP, which applies to a best-effort transport service.

9. Concluding Remarks

In the emerging open telecommunications market, users will have the opportunity to buy affordable services matching well their requirements and preferences; this however may involve complicated decisions regarding features of customized services and selection of service retailer (or provider) and/or tariff. This task is appropriate for an intelligent software agent that will negotiate with service retailers on behalf of the user. In this paper, we developed and analyzed the performance of such an agent that negotiates for a News-on-Demand service in order to optimize the selection criterion of the user. The agent makes use of a simple yet informative model for user utility and a negotiation algorithm that appears to be optimal or near-optimal in a wide variety of cases. Our approach is computationally efficient, thus being amenable to mobile agents, which migrate to retailers' hosts in order to negotiate locally. Our ideas for user utility approximation and agent negotiation are applicable to a wide variety of network technologies (including Internet, ATM, etc.), as well as to other kinds of services that can be represented as a set of options to be specified in the negotiations phase. We have also presented theoretical and experimental results on the assessment of the performance of the negotiation algorithm. These results suggest that this algorithm has in general good performance.

An interesting direction for future research is to combine our approach with learning, in order to customize the user

utility parameters. Indeed, recall that some of the parameters were based on the user profile, while others were taken as common for all users. In reality, however, each user has his own valuation of the various QoS levels as well as his own relative utility per medium in a content section etc. These will in general not coincide with the corresponding "average-user" values used. Thus, the defined parameters can be employed as initial values in the model, which will be updated by means of the learning capabilities of the UA.

An interesting as well as essential direction for future work is to test our approach in a real environment by estimating the user satisfaction from the agent's selections. Such tests will be conducted in the user trials of ACTS project MONTAGE.

Acknowledgements: The authors would like to thank the members of the consortium of the ACTS project AC325 MONTAGE, and Dr. Vasilios Siris, for useful discussions on the subject of the paper.

References

[1] http://www.fokus.gmd.de/cc/ima/climate

[2] C.Courcoubetis, C.Manolakis, G.D.Stamoulis, and F.P.Kelly "An Intelligent Agent for Optimizing QoS-for-Money in Priced ABR Connections", Technical Report 218, Institute of Computer Science - Foundation of Research and Technology - Hellas, May 1998 – available at: http://www.ics.forth.gr/netgroup/publications/1998.tr218.abr_intelligent_agent.ps.gz

[3] http://www.ietf.org/html.charters/diffserv-charter.html

[4] Rajarshi Gupta, Mike Chen, Steven McCanne and Jean Walrand, "WebTP: A Receiver-Driven Web Transport Protocol", available at:

http://www.path.berkeley.edu/~guptar/webtp/index.html

[5] IEEE Communications Magazine, Special Issue on Mobile Software Agents for Telecom-munications, 36(7), July 1998.

[6] http://www.ietf.org/html.charters/intserv-charter.html

[7] http://www.cs-ipv6.ac.uk/ipv6/documents/papers

[8] Keith Jack, *Video Demystified: A Handbook for the Digital Engineer*, Second Edition, High Text Publications Inc, 1998.

[9] AC325 project MONTAGE: Mobile Intelligent Agents in Accounting, Charging, and Personal Mobility Support; see http://montage.ccrle.nec.de

[10] J.K. MacKie-Mason and H.R. Varian, "Pricing the Internet", *Public Access to the Internet*, B. Kahin and J. Keller, Prentice-hall, Englewood Cliffs, New Jersey, 1994.

[11] http://www.ietf.org/html.charters/rsvp-charter.html

[12] Proceedings of UM97: the Sixth International Conference on User Modeling, Chia Laguna, Sardinia, Italy, June 2-5, 1997.

[13] Hal R. Varian, *Microeconomic Analysis*, W. W. Norton, New York, 1992.

[14] Scott Shenker, "Fundamental Design Issues for the Future Internet", *IEEE J. Select. Areas Commun.*, vol. 13, no. 7, pp. 1176 –1188, 1995.