Applying a Policy-based Framework to Manage Quality of Service Requirements in the Virtual Home Environment

Alvin Yew, Antonio Liotta, George Pavlou
Centre for Communication Systems Research
University of Surrey
Guildford, Surrey GU2 7XH, United Kingdom

***Abstract*-The deployment of the Virtual Home Environment (VHE) concept in 3G mobile systems in the near future will place many demands on managing a user's personalised service environment. Key to the VHE concept is the performance of such services, resulting in a need to manage Quality of Service (QoS) demands and allocation on behalf of users. This paper examines the issues involved in managing QoS demands from the various roles and entities in the VHE and proposes a policy-based framework to satisfy their needs.**

I. INTRODUCTION

The Virtual Home Environment (VHE) is a concept introduced by the 3rd Generation Partnership Project (3GPP) for 3rd generation (3G) mobile communication systems that aims to enable its users' personal service environment to be portable across network boundaries and between terminals. The core of the VHE concept is to ensure that users are consistently presented with the same personalised features, user interface customisation, and service preferences at all times [1]. The complexity of implementing the VHE is apparent given that the VHE encompasses any type of underlying network, access terminal, and user location, thus requiring coordination between many entities in achieving its aims. An important implication in the VHE concept is that the Quality of Service (QoS) provided by the network may affect the provisioning of VHE services to the user. A bandwidth-intensive video-on-demand service is an appropriate example of a service that is dependent on the available QoS in the network for its quality and performance during service delivery. The QoS provided in the network could therefore determine the 'look and feel' of VHE services from the user's perspective as the quality and performance of VHE services constitutes as an integral part of the user's service experience in the VHE. As such, there is an important need to manage QoS in the network on behalf of VHE users whenever they access their services so as to fulfil that aspect of the VHE concept.

Managing QoS on behalf of a user traditionally involves negotiating and establishing a service level agreement (SLA) between the network provider and the user. In most cases, the SLA is of a static nature, relying on the fact that the user's network usage for a service is predictable and therefore the required QoS can be reserved from the network resources in advance. A static SLA between the network provider and the user that satisfies the user's requirements most of the time can therefore be seen as an appropriate trade-off between a lack of flexibility as a result of reserving the user's required QoS in advance, and increased SLA management overheads. However, the trade-off in user flexibility for static SLAs is no longer advantageous when considered in the VHE and service mobility contexts. If we extend the VHE concept beyond 3G mobile communication systems, then users will be able to access their VHE services through a variety of access networks (e.g. Wireless LAN, 56 Kbps dial-up modem etc.). The user's network usage in any single network then becomes more difficult to predict for network resources to be reserved in advance. Furthermore, the QoS available in some access networks may vary considerably during service usage (e.g. handover in mobile networks etc.). Therefore, we can deduce that based on the issues raised above, QoS negotiation between the user and the network provider must be done on the fly, leading to the requirement of dynamic SLAs within the VHE. With dynamic SLAs, intelligent decisions must be made about the user's QoS requirements on the fly. From a non-'VHE aware' point of view, the decision could be made from invariable-like factors such as projected amount and the shape of traffic generated by the VHE service. However, considering that a 'VHE capable' service supports a number of user terminal types and is readily adaptable to various access networks, the factors on choosing the correct level and managing the allocated QoS for the VHE service are more variable. Therefore, there is a need to find a suitable method of weighing up all the factors to decide how to manage QoS on behalf of a user in the VHE.

II. DETERMINING THE USER'S QOS REQUIREMENTS

It is important to derive the user's QoS requirements from the network before any kind of QoS management in the VHE can be done. There are many factors that can be used in determining the user's QoS requirements. These can be all or a combination of the following: the terminal characteristics,

0-7803-7400-2/02/$17.00 © 2002 IEEE

the VHE service QoS requirements, the user's personal preferences, and the access network policies.

A. Terminal Characteristics

The terminal characteristics may be important in determining the VHE service's QoS requirements. For example, a content-based VHE service may adapt its service to the screen size of the user's terminal (for example, use smaller pictures). The manner in which the VHE service obtains this information could be though the Content Capabilities/Profile Preferences (CC/PP) framework, and the CC/PP exchange protocol, both of which are currently work in progress by the World Wide Web Consortium (W3C) [2][3]. After examining the capabilities of the user's terminal, the VHE service can then decide on the appropriate service model that best suits the terminal before delivering the service. A typical sequence of QoS reservation from the network through the initial user request incorporating the CC/PP protocol is shown in Fig. 1. The different service models catered for different types of terminals can have different QoS requirements. This therefore implies that the user's terminal capabilities can play a part in deriving the QoS required from the network.

Another important consideration regarding the issue of terminal capabilities with regards to QoS is the amount of bandwidth that the terminal supports. The amount of bandwidth reserved in the network should never be more than that which the terminal can handle.

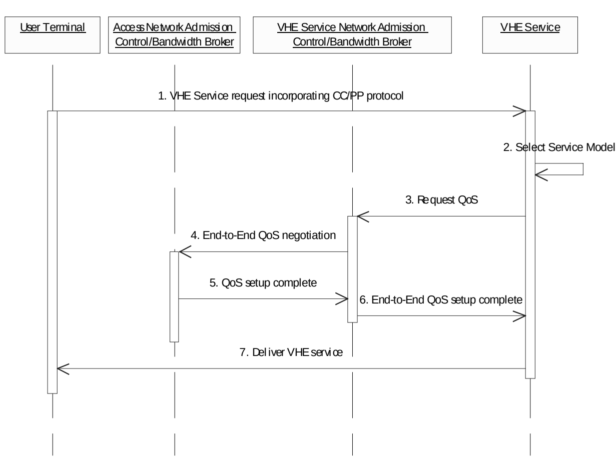

Fig. 1. Resource reservation procedure based on CC/PP exchange protocol

B. VHE service QoS requirements

Services can be classified into two different categories – adaptable, and non-adaptable. An adaptable service changes its service model to achieve certain predefined goals in its service delivery. These goals could vary from user satisfaction through its perceived performance, to service accessibility catering for a wide range of terminal types. A good example of an adaptable service is a real player video service that adapts the video media bit rate to suit the bandwidth in the user's access network. A non-adaptable service uses the same service model to deliver the service to its users irrespective of any external circumstances. An example of this could be a file transfer service that depends on the File Transfer Protocol (FTP), although the latter uses the Transmission Control Protocol (TCP) that provides a degree of adaptation through its congestion control mechanisms. A common effect experienced by both adaptable and non-adaptable services is that the performance of a service is proportional to the level of QoS delivered by the network.

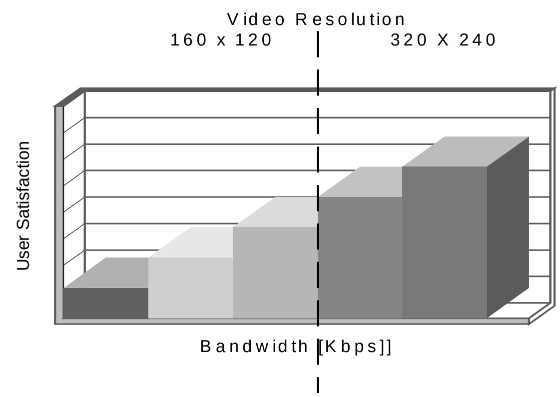

Fig. 2. Service model selection for video service

Deriving QoS requirements from the network for adaptable services involves more complexity than that of non-adaptable services. This is because that manner in which the service chooses its service model will have to be considered. For example, the choice of service model could be related to terminal characteristics. Fig. 2 shows a layered video service with two service models represented by two separate video files with video resolutions of 160×120 and 320×240 pixels respectively. The decision on the choice of the video file to deliver to the user is based on the user's terminal maximum screen resolution. Within each of the two service models, the video can adapt to the available bandwidth in the network giving various levels of user satisfaction (e.g. adjusting the frame rate of the video). If the manner in which an adaptable service chooses its service model is known in advance, dynamic QoS reservation in the network can be performed more efficiently. For example, in the case of the layered video service shown in Fig. 2, it can be deduced in advance that there is no point in requesting a bandwidth that is between two bandwidth thresholds, as this would not result in greater user satisfaction. A consideration regarding a VHE service's service model selection is that it can be based on other suitable criteria (e.g. type of access network) instead of the user's terminal characteristics.

C. User Preferences

Choosing the preferred QoS is one of the essential methods that a user can use to customise a service. By customising a service and stating preferences, the user can have an input in the 'look and feel' of his service experience. This subsection presents some additional user customisation options that deal with QoS allocation in the network.

It is possible to deduce the optimum QoS on behalf of a user for a particular service if the user's terminal capabilities and the VHE service QoS requirements were the only two factors considered. However, the user may not always desire to have the best QoS available in the network for his VHE service. Reasons for this include the service subscription and the cost of the network usage. For the latter, research in allowing users to dynamically select the QoS required based on costs considerations include a framework for integrating pricing with resource reservations [4]. For the former, the user may have a contract with the service provider that differentiates the service offered at the application and/or the network level within the service domain. For a video-on-demand service, this could be achieved by assigning different video servers with different traffic policies to different classes of subscriber [5]. A customer may then have the cheapest subscription to the video services where the video is of the smallest available resolution regardless of his terminal screen resolution. This factor is dependent on factors within the service's business model and results in the user having a predefined service profile for the particular service.

D. Access Network Policy

In some cases, the access network's policy on resource and QoS allocation may instil an upper bound level of QoS available to the user that is lower than that of the terminal capabilities. For example, a General Packet Radio Service (GPRS) terminal should, according to the 3GPP specifications, be able to receive up to 384 Kbits/sec of data from the network (downlink). However, due to the 'in-maturity' of the GPRS network equipment used, the marketing policy, and the capacity planning involved, the amount of bandwidth that most current GPRS network operators assign to the downlink data channel is approximately half of the theoretical maximum GPRS capability. Another possibility of the access network policy being a major factor is that the VHE provider, which manages and coordinates the user's VHE, may have a pre-established SLA with various access network providers on the maximum amount of QoS that its users are allowed to request and use.

By considering all the factors presented in this section regarding the choice of QoS for a VHE service, it is possible to find an ideal level of QoS to request and reserve from the network. Fig. 3 shows the factors considered in the four-step decision process to achieve the desired QoS parameters that will enable the VHE service to deliver the same personalised 'look and feel' regardless of network or terminal as required in the VHE concept. As each step of the decision process is completed, the range of the QoS narrows. In some cases, one or more factors in Fig. 3 need not apply in the decision process. For example, if the user did not personalise the service, then the user service profile factor becomes irrelevant. For such a scenario, there will be a range of QoS levels that are applicable for the service at the end of the decision process. Choosing a suitable QoS level may then depend on factors discussed in the next section.

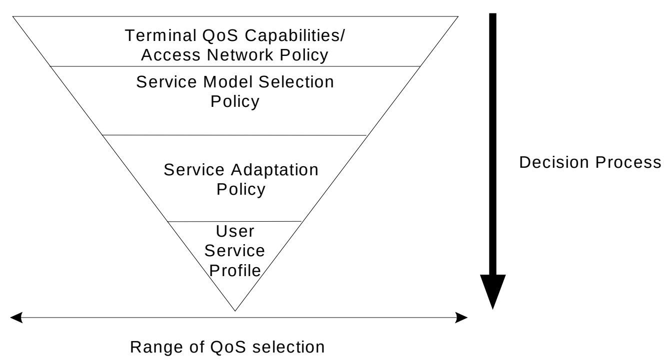

Fig. 3. Decision process of selecting QoS level for a VHE service

III. Managing multiple active VHE services

While the previous section presented the factors involved in choosing the QoS parameters for a VHE service on behalf of a user, it is also important to consider how the decision process outlined in Fig. 3. may be affected if the user accesses more than one VHE service concurrently. An obvious consideration when a user is accessing two or more VHE services concurrently is that the total QoS requirements of all active VHE services should never be more than that of the user's terminal capabilities at all times. Management of the user's VHE services' required network QoS would then be performed under the confines of this restriction.

With the restricting factor of the terminal capabilities determining the upper bound performance of multiple active VHE services, managing the connectivity and QoS of these services to conform to the user's personalised VHE becomes crucial to maintain user satisfaction and cost efficiency.

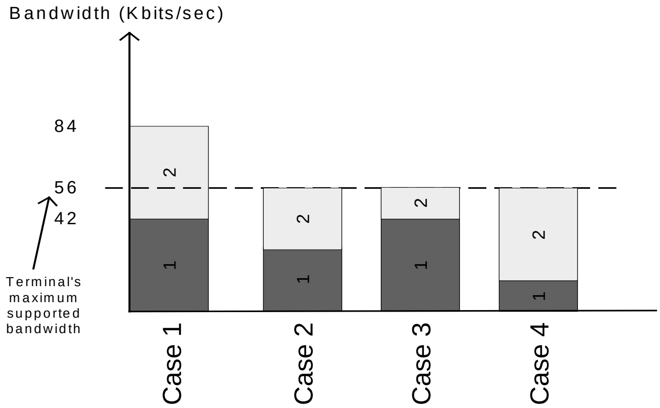

Fig. 4. Case study of bandwidth allocation of 2 active adaptable services

Fig. 4 shows four possible suggested ways in which to make QoS decisions when two VHE services (labelled 1 and 2) with similar service adaptation policies and QoS requirements are active. In case 1, the user activates the second VHE service (labelled 2) while still using service 1 with a 56 Kbits/sec dialup modem. The two services, unaware that another service is active, both request a bandwidth of 42 Kbits/sec for suitable operation. The over-subscription of bandwidth occurs because without coordinated QoS management, the two services each believe that they are still requesting a bandwidth within the terminal capabilities. Case 2 illustrates the effect of coordinated QoS management. In this case, the overall decision process is based on balancing the user satisfaction derived on the quality of both services, resulting in an equal share of the terminal's maximum supported bandwidth as proposed in [6]. In case 3, the allocation of bandwidth between the two services is based on a first-come-first-served basis. When the user activates the second service, a decision is made based on the terminal capabilities to allocate the remainder of the maximum terminal supported bandwidth to the second service. This method works when the remaining bandwidth is high enough to satisfy the second service adaptation policy's minimum bandwidth requirements. In case 4, an advanced priority based decision process is used to divide the bandwidth between the two services. The user is allowed to have an input into which service that s/he values more and therefore should be given preference when performing QoS reservation and allocation. In this case, the user indicates that s/he favours service 2 over service 1, resulting in service 2 being assigned a larger share of the bandwidth over service 1. This requires modifying the existing allocated bandwidth of service 1 and making service 1 adapt to the lower available bandwidth.

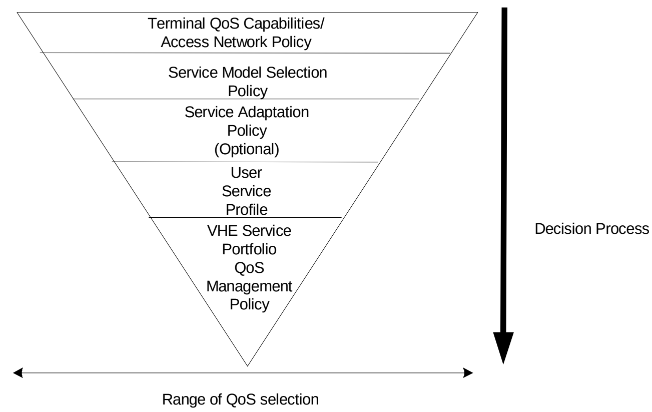

Fig. 5. Overall decision process involved in coordinated QoS management

The case study above is a simple example to highlight the choices that could be made when trying to manage two or more active VHE services, and the deficiencies of a non-coordinated resource allocation procedure as shown in Fig. 1. Other advance decisions in deciding QoS allocation could be based on minimising the cost against the total perceived user satisfaction, or maximising the total perceived user satisfaction gained. After explaining the need for coordination for QoS allocation between multiple active VHE services, a further step in the decision process should be appended to the end of the decision process described in Fig. 3 to derive the optimum QoS level required by each service when making decisions on QoS allocation and reservation. This step is to consult the 'VHE service portfolio QoS management policy' which determines how to balance QoS allocation within terminal capabilities in the case of multiple active services. The addition of this further step in the decision process is mandatory and the overall decision process is shown in Fig. 5. The range of customisation options for the VHE service portfolio QoS management policy is dependent on the VHE provider. These options can then be accessed and personalised by the VHE user.

IV. APPLYING A POLICY-BASED FRAMEWORK

In this section we suggest a policy-based framework adapted from the Internet Engineering Task Force's (IETF) work on policy-based management to manage the various QoS requirements of the services on behalf of the user in his VHE service portfolio. We use IETF's terminology on policy-based management to describe our proposed framework [7].

In the framework, there are four separate logical entities involved in delivering the VHE service to the user – VHE provider, Access Network provider, VHE service provider, and the VHE user. The VHE provider is in charge of coordinating and managing the user's home environment to support the VHE concept. The Access network provider provides the network that is used as an access point for the user to access his home environment. The VHE service provider owns and provides a service that is part of the user's VHE service portfolio. The VHE user uses VHE services through the VHE provider via the access network (provided by the access network provider).

In the suggested framework, the policy decision point (PDP) is the VHE provider as it is the logical entity that makes decisions on the operational management of the user's home environment. The VHE provider considers all the factors presented on the previous two sections when making its decision on choosing the policy to use for managing the user's VHE services' QoS requirements. As such, the VHE provider will need to know the relevant service policies from all the VHE service providers in the user's portfolio, the VHE user's own personal preferences, and his/her service portfolio policy in order to make an informed decision on the correct choice of policy. Therefore, all of these policies must be housed in a centralised policy repository made accessible to the VHE provider. Limited access to policies in the centralised policy repository for the VHE service providers and the users allows them to modify only the policies they own (e.g. user preferences for the VHE user etc.).

The choice of a policy enforcement point (PEP) is based on the location of the node that makes the actual QoS reservation. If we consider that heterogeneous access networks are involved in delivering VHE services to the VHE user, then QoS reservation should began at the access network since the range of QoS levels available to the user at the access network would be less than that supported at the VHE service's network. This would result in a more efficient resource reservation process, as rejected reservations would be fed back to the PEP earlier.

With the recent developments by network providers in extending the capabilities of the underlying network to external authorised entities via 'soft network' solutions such as the Parlay APIs and the Open Service Access (OSA), management of QoS aiding the establishment of dynamic SLAs can be implemented in a more distributed rather than traditional centralised manner [8][9].

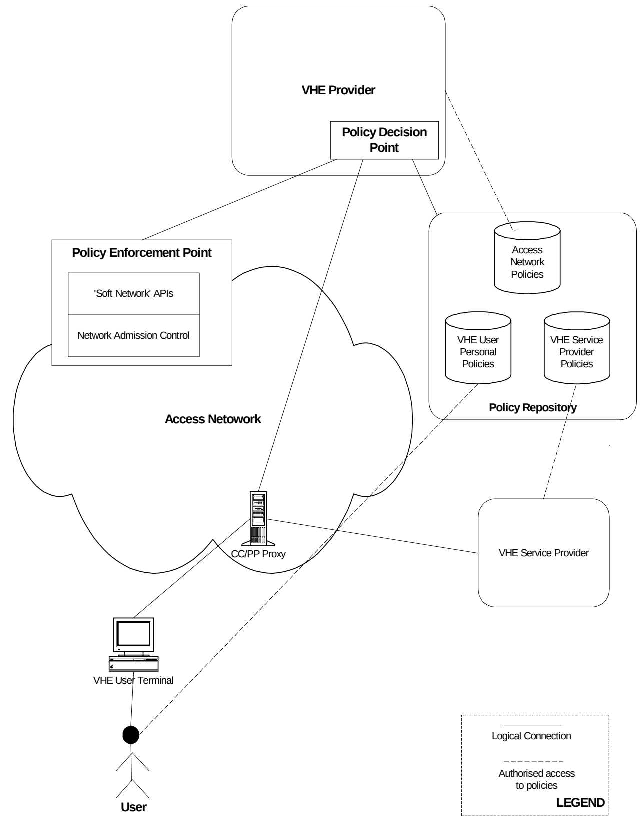

Fig. 6. Architecture of the proposed policy-based framework

By combining distributed object strategy together with such 'soft network' solutions, the VHE provider can manage and configure remotely a QoS reservation object or mobile agent at an edge node at the access network near its resource admission control centre to make resource reservations through the Parlay/OSA APIs on behalf of a user. This object or mobile agent can then act as a PEP in the policy-based framework. We have previously explored the required modifications to the Parlay APIs and the OSA to support dynamic QoS management for the VHE in [10]. This includes the ability to change the QoS parameters of a user's existing connection, and to receive performance-monitoring results of the user's connection so as to make informed decisions when managing the QoS in the user's connection. For example, the PEP will need to attempt to re-negotiate the level of QoS with the access network provider on behalf of the user if it falls below the minimum required to sustain the VHE service. Fig. 6 shows the suggested policy-based framework to manage QoS policies in the VHE. The CC/PP proxy in the access network exposes the user's terminal capabilities to the PDP and the VHE service provider so as to allow policy decisions and service adaptability based on terminal capabilities respectively.

V. Conclusion

Deciding QoS requirements and allocation in the VHE requires flexible management to satisfy the user's demands. Flexible management is required because of the nature of the VHE concept itself, which is that users can have numerous customisation options and a flexible choice of terminals, and access networks used. Current management of user services do not include the full range of flexibility expected in the VHE into account when deciding on service provision. A reason for this is that there is no high level management for the user's entire portfolio of services, a task that we feel that will probably be assigned to the VHE provider in the future. In order to manage such complex management requirements on behalf of the user, we propose to adopt a policy-based framework that will be flexible to different service models from VHE service providers, and allow the user to have some input in customising and personalising his service environment in the VHE. It is true that by giving the user more flexibility determining the 'look and feel' of his/her services, more management overheads will be incurred. However, this can be offset if the management framework is efficient, can provide the user the satisfaction s/he requires when using the services, and is flexible enough to allow new types of terminal and network portable services to be deployed easily into the VHE. The latter will ultimately be the determining factor if the VHE is to be a success in the future, and may facilitate the future development of the 'killer application' that 3rd generation mobile network operators yearn for.

We are currently working towards a detailed design/implementation of a policy-based QoS management framework for the VHE in the context of the IST VESPER project [11]. VESPER is tackling the issues involved in the VHE from a wider angle, including the introduction of a suitable VHE architecture and the realization of service continuity, service portability, service personalization, and session mobility. The integration of our policy-based system with the VESPER platform will serve to verify the proposed approach with a larger number of typical VHE scenarios.

Acknowledgment

This work has been performed in the framework of the project IST VESPER, which is funded by the European Community. The Authors would like to acknowledge the contributions of their colleagues from Intracom Hellenic Telecommunications and Electronics Industry S.A., National Technical University of Athens, Institut National de Recherche en Informatique et Automatique, IKV++ GmbH Informations und Kommunikationstechnologie, GMD - Forschungszentrum Informationstechnik GmbH, Fondazione Ugo Bordoni, Universita' Di Catania, Portugal Telecom Inovação, University of Surrey, Technical Research Centre of Finland and SIEMENS AG Osterreich.

References

[1] 3rd Generation Partnership Project, "The Virtual Home Environment (Release 4)", *3GPP Technical Specification*, TS 22.121 v4.0.0, October 2000.

[2] F. Reynolds, J. Hjelm, S. Dawkins, and S. Singhal, "Composite capability/preference profiles (CC/PP): a user side framework for content negotiation", *W3C working draft*, July 1999, available at http://www.w3.org/TR/NOTE-CCPP.

[3] H. Ohto, J. Hjelm, "CC/PP exchange protocol based on HTTP extension framework", *W3C working draft*, June 1999, available at http://www.w3.org/TR/NOTE-CCPPexchange.

[4] X. Wang, and H. Schulzrinne, "An integrated resource negotiation, pricing, and QoS adaptation framework for multimedia applications", *IEEE J. on Selected Areas in Communications*, Vol. 18, No. 12, pp. 2514-2529, Dec 2000.

[5] F.Hao, E.W Zegura, and M.H. Ammar, "Supporting server selection in differentiated service networks", *Proc. Of IEEE INFOCOM 2001*, Vol. 2, pp. 659-668, 2001.

[6] G. Bianchi, A. T. Campbell, and R. R.-F. Liao, "On utility-fair adaptive services in wireless networks", *6th Int. Workshop on Quality Services 1998 (IEEE/IFIP IWQOS'98)*, May 1998, pp. 256-267.

[7] A. Westerinen, J. Schnizlein, J. Strassner, M. Scherling et al., "Terminology for policy-based management", *IETF Internet Draft*, Policy Framework Working Group, July 2001, available at http://www.ietf.org/internet-drafts/draft-ietf-policy-terminology-04.txt.

[8] The Parlay Group: "Connectivity Manager Service Interfaces", *Parlay APIs 2.1 Issue 1.1*, June 2000.

[9] 3rd Generation Partnership Project, "Stage 1 Service Requirements for the Open Service Access (OSA)", *3GPP Technical Specification*, TS 22.127v4.0.0, January 2001.

[10] A. Yew, C. Bohoris, A. Liotta, and G. Pavlou, "Quality of service management for the virtual home environment", *Proc. of IFIP/IEEE DSOM 2001*, in press.

[11] Virtual Home Environment for Service Personalization and Roaming Users (VESPER) IST project, http://vesper.intranet.gr.