Quality of Service Management for the Virtual Home Environment

Alvin Yew, Christos Bohoris, Antonio Liotta and George Pavlou

Centre for Communication Systems Research, School of Electronics, Computing & Mathematics
University of Surrey, Guildford, Surrey, GU2 7XH, UK
E-mail:{K.Yew, C.Bohoris, A.Liotta, G.Pavlou} @eim.surrey.ac.uk

The Virtual Home Environment (VHE) concept poses important requirements on Quality of Service (QoS) management, which are not fully addressed by current OSA and Parlay frameworks. OSA/Parlay expose a limited amount of network information, on the one hand, whilst not allowing for the required level of QoS control, on the other hand. This paper elaborates on the QoS functionality required by the VHE and discusses the related limitations of and necessary extensions to current OSA/Parlay specifications. A case study is used to illustrate the issues tackled in the paper.

Keywords: VHE, QoS Management, MExE, Parlay, OSA.

1. Introduction

The Virtual Home Environment (VHE) is a concept introduced by the 3rd Generation Partnership Project (3GPP) for UMTS aiming to enable its users' personal service environment to be portable across network boundaries and between terminals. The VHE's target is to ensure that users are consistently presented with the same personalised features, user interface customisation, and service preferences. The complexity of the role of the VHE is apparent given that the task can involve any type of underlying network, any type of access terminal, wherever the user may be located. An important implication in the VHE is that the Quality of Service (QoS) provided by the network affects the provision of such VHE services to the user. Users may choose to map VHE services to the QoS demanded from the network, thus creating a mapped services-to-'preferred network QoS' user profile. Among one of the methods proposed by 3GPP to facilitate the VHE concept is the Mobile Execution Environment (MExE) [1]. MExE allows service providers to develop terminal-centric applications that suit different types of terminals according to their capabilities and characteristics, regardless of the terminal's platform. These MExE executables that reside on the terminal can then interact with a MExE server to provide terminal-independent services in the VHE. Although the main emphasis of MExE is to ensure service interoperability between different types of MExE terminals, 3GPP has also stated the requirement for a QoS manager to be present in these terminals [2]. According to 3GPP, the MExE QoS manager is responsible for managing the QoS streams for MExE executables and notifying the negotiated and delivered QoS to the end user and MExE executable. However, this approach only satisfies the VHE QoS management requirements of 3rd Generation (3G) MExE-compliant terminals, as only they will have the MExE QoS manager. Furthermore, the MExE QoS manager achieves QoS management through an interface that is compatible with 3G core network protocols only. It is therefore important to developing a general model of a QoS manager for the VHE that is not utterly dependent on 3G technology, with the aim of providing the benefits of VHE capable QoS-aware services to users of any network.

If the VHE concept were viewed from a non-3G mobile network point of view then providing dedicated servers that allow access to interfaces exposing network capabilities and QoS management functionality would seem to be the most appropriate means of providing services in the VHE. This would allow non-MExE terminals, such as current laptops, to access VHE services through non-3G mobile networks (e.g. fixed IP networks). However, 3GPP's current set of such interfaces, i.e. the

Open Service Access (OSA), do not incorporate any QoS management functionality [3]. As 3GPP acknowledges that VHE should be applicable to all types of future networks, other consortia such as the Parlay group have been working closely to 3GPP and have defined interfaces comparable to that of OSA to work in fixed networks [4]. In view of these recent developments, issues on QoS provisioning and management through standardised network interfaces such as OSA/Parlay should be extremely important in the future. By taking advantage of this 'liberalisation' of network capabilities, we can develop a general QoS manager by transferring the responsibilities of the QoS Manager from the terminal (i.e. 3GPP's approach) onto a dedicated server as shown in figure 1. The server-based QoS manager will then perform its duties through manipulation of the Parlay/OSA APIs, and act as a mediator for server-centric services requesting QoS-related support and functionality. When planning the interactions between the QoS manager, service, and the network, we need to consider two key design requirements. First, the information exchanged via the QoS manager needs to be comprehensive enough to ease the realisation of 'QoS aware' services, i.e. services that can exhibit a sufficient level of adaptation with respect to the underlying network layer. Second, the QoS manager needs a sufficient level of access to the underlying networks to be able to dynamically negotiate the effective use of network resources. As neither OSA nor Parlay currently provides such dynamic QoS management features, there is a need to propose extensions and modifications to these interfaces. The issues of locating a QoS manager on a dedicated server instead of on a terminal with respect to VHE operation must be examined as well.

The next two sections present a short background on the VHE concept, and the OSA/Parlay architectures. We then elaborate on the role and functionality of the QoS manager, an essential component of a VHE system. Finally, we discuss extensions to current specifications of OSA/Parlay necessary to the VHE for the support of adaptable, QoS-aware services and applications. We conclude in Section 6 with a simple case study which serves to illustrate the concepts introduced in the paper.
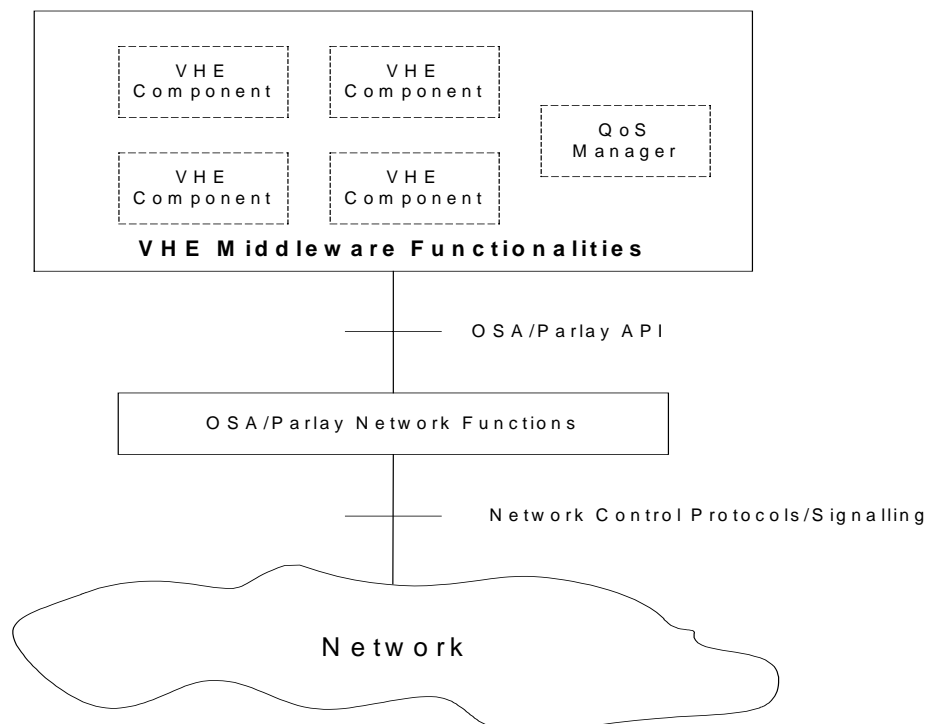

Figure 1   - A Server-based QoS Manager for the VHE

2. The VHE Concept

There are many definitions of the term VHE from various sources, each describing its interpretation of VHE enabling features. This section examines the similarities and differences in the expected VHE enabled features. Some definitions of the VHE are presented in table 1 below. (Note that the last two definitions are from European projects researching the VHE concept).

	Definitions
3GPP	"The VHE ensures that users are consistently presented with the same personalised features, User Interface customisation and services in whatever network and whatever terminal (within the capabilities of the terminal and network), where ever the user may be located" [5]
GSM MoU	"Virtual Home Environment (VHE) is a system concept for service portability in the Third Generation across network borders" [6]
ITU / IMT2000	"VHE is a capability whereby a User is offered the same service experience in a visited network as in his Home system." [7]
UMTS Forum	"VHE means that the user will have the same interface and service environment regardless of location (personalised user interface independent from the current serving network)." [8]
Eurescom P920-GI	"The Virtual Home Environment is an environment which presents the user with a common look and feel interface and service experience regardless of location, network and terminal type. The VHE is based on standardised service capabilities and personalised features that are consistently presented so that the user always "feels" that he is on his home network even when roaming across network boundaries" [9]
IST VESPER	"VHE main feature is that the customised environment will be following the user while he/she is roaming within different networks and using different terminals" [10]

Table 1 VHE Definitions

The common denominator among those definitions is provisioning of customisable and personalised services, portable across heterogeneous environments. However, there are some issues with regards to other aspects of the VHE, as discussed below.

2.1. Home Environment

3GPP defines the 'Home Environment' (HE) as the environment that is responsible for the overall provision of services to users. In ITU's VHE definition, the term 'Home system' is used in comparison with the term 'visited network'. Through this comparison, it appears that the HE is within the domain of the user's home network provider. The term 'home network' must be clearly defined if the VHE concept includes non-3G mobile networks. For example, if a user subscribes to different mobile and IP network providers, then both networks can be logically referred to as the home network. In this case, the distinction is that the former is the home mobile network, while the latter is the home IP network. If the user personal service environment (PSE) is to be portable between different types of networks, it is possible that the HE is controlled by a business entity separate from the home network provider.

2.2. Service Performance Preferences

In ITU's definition of VHE, the user is offered the same service experience regardless of network used. The expected performance of the service can be interpreted as part of the user's service experience. For server-centric services (e.g. video-conferencing) service performance is partly dependent on the QoS offered by the network. In this respect, there are three options available to

ensure that the user gets the same service experience: a) let the service adapt to the network conditions; b) let the network adapt to the service requirements; and c) a mixed approach including both service and network adaptation. In the example of video conferencing, option (a) could involve adjusting the frame rate of the video to suit the network conditions. For option (b), a guaranteed QoS would be required from the network before the video conferencing begins.

2.3. Networks and Terminals

As all of the definitions in table 1 state that the user's service experience should be similar regardless of the network used, the primary understanding of the VHE concept seems to be that of service interoperability and portability between networks. However, it is important to note that 3GPP have added the condition of terminal-independent service functionality to its definition of the VHE (although service provision itself is limited to within the capabilities of the terminal and network).

This presents another consideration for service providers when developing services or applications. Service providers must therefore ensure that they can adapt their applications and services to both terminal and networks. These two factors are not mutually exclusive as the terminal's capabilities may also determine the maximum bit rate to run the application or service. The bandwidth required from the network to support the service should therefore not exceed the maximum bandwidth supported by the terminal for reasons of costs and efficiency. In the case of a video-on-demand service, the bandwidth reserved in the network should complement the maximum receivable bit rate of the terminal. For example, a 56 kb/s modem user using a dial-up Internet connection shouldn't reserve 2 Mb/s bandwidth from the network. The terminal must therefore be able to announce its capabilities to an interested entity (e.g. network, service etc.) in order to prevent this from happening. A suitable solution to this problem using HTTP extensions is given in [11].

3. Parlay and the Open Service Access

As the Parlay APIs and the OSA will probably be the technologies to hide network heterogeneity to the VHE, it is important to understand their differences in architecture and implementation. This section presents a short introduction on the two technologies.

The OSA was developed by 3GPP to allow service providers to develop server-centric services requiring UMTS network capabilities in an unrestrictive manner. Figure 2 illustrates a high level view of the open service access while figure 3 shows the Parlay architecture. The key component of both Parlay and the OSA is the framework server and its Application Programming Interface (API). These allow applications to discover the network functionalities that are exposed by the network provider. The network provider is not obliged to expose the entire functionality of the network through the framework API, and therefore service discovery is a key feature in both Parlay and OSA.


Figure 2   - The Open Service Access

QoS Management in the Virtual Home Environment

The APIs also exposes the service capabilities, such as call control and messaging, of the network. For the case of the OSA, the API calls are handled by different service capability servers (SCS). Parlay uses a different approach whereby the different categories of service capabilities are separated into different respective APIs.

As was mentioned before, a limitation in the OSA API is that there are no methods for QoS control and management that can be used by external applications. QoS management has been partly addressed by the Parlay APIs through the introduction of the connectivity manager interface, i.e. interface 6 in figure 3 [12]. This interface allows an enterprise to enquire about and configure aspects of QoS provisioning in its virtual private network provided by the network provider. However, the connectivity manager interface was designed with the intention of providing QoS management on an enterprise scale and not for individual users. The intended user of the connectivity manager interface is shown to be the enterprise operator administration tool in figure 3. Therefore, Parlay does not provide any APIs for a server-centric service application to perform QoS management on behalf of its individual users. Another difference is that Parlay allows for some enhancements by third party vendors through the use of interfaces 3 and 5 in figure 3 whereas OSA does not. This makes Parlay a better platform than OSA to implement additional network service functionality.
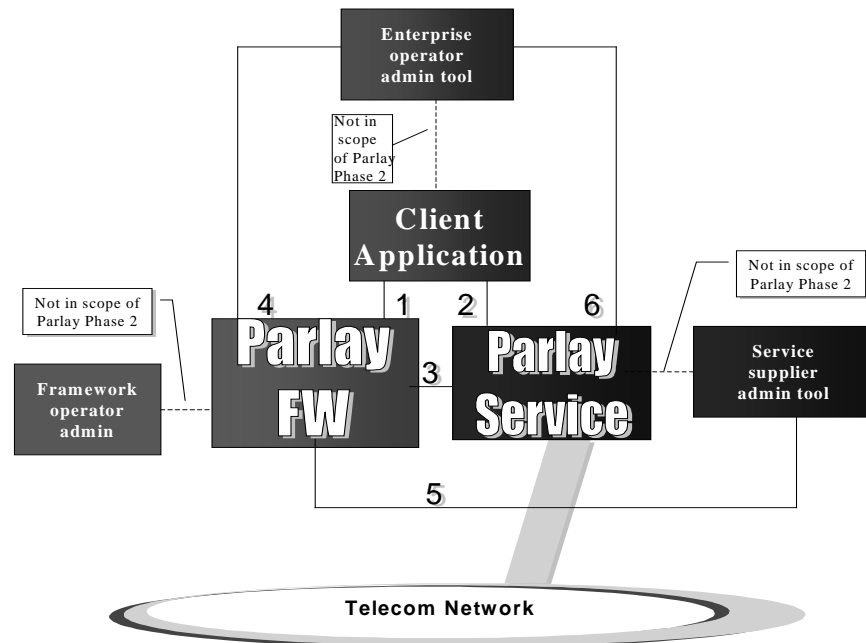

Figure 3   - The Parlay Architecture [12]

4. Role of the QoS Manager

In order to better understand the role of the proposed server-side QoS manager, the logical view of the MExE QoS manager should be presented and explained first. The right side of Figure 4 shows the logical view of the MExE QoS manager as specified by 3GPP while the left side shows the logical view of the proposed server-side QoS manager. It is apparent from the diagram that the MExE QoS manager acts as a mediator between the MExE executable and the network. By providing an interface to the MExE executable (i.e. the MExE QoS API), the MExE QoS manager permits MExE developers to allow their applications to request QoS related functionality from the network in a high level manner without resorting to the complicated 3G core network protocols. These network protocols are represented by the interface named 'Network Control API' between the MExE QoS manager and the network.
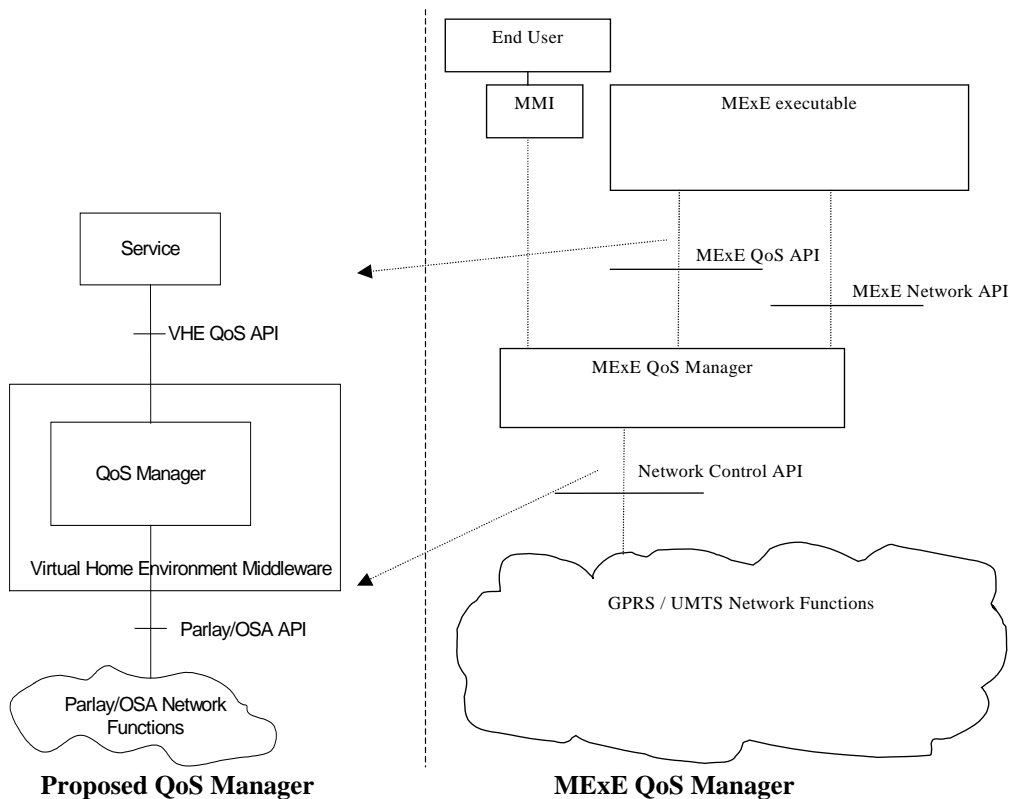

Figure 4   – Logical views of the MExE, and the proposed server-side QoS Managers

4.1. VHE QoS API

In figure 4, the arrows between the two logical views of the QoS managers provide a mapping of the functionalities of the 3GPP-defined interfaces to its equivalent in the proposed server-side QoS manager. The only difference between an interface and its equivalent in the other logical view is the manner in which they are accessed or used. The Network Control API operates through the use of wireless signalling channels whereas the Parlay/OSA API operates through the use of distributed object technologies. For the case of the VHE QoS and the MExE QoS APIs, the former can be used remotely whereas the latter is always used locally (i.e. the MExE executable and the MExE QoS manager are always located on the same terminal).

It is intuitive to wonder about the need for a VHE QoS API and the server-side QoS manager when the interfaces defined in Parlay/OSA are sufficiently high-level for service developers. There are two reasons that justify the need for the QoS manager and VHE QoS API. The first reason stems for the business models of both Parlay and OSA. For the service to access these interfaces directly, a prior agreement must be made between the network and the service provider. This is due to the high security required by network providers as to prevent an unauthorised user manipulating resources and capabilities of their networks. This becomes a problem if a user of the service is a subscriber to more than one type of network (e.g. mobile and fixed IP), which is likely if the VHE concept applies to all types of networks. In this scenario, economy of scale plays a major role as the service provider would not like to sign a service agreement with a network operator where only a few of its users are patrons. It is likely that only large-scale service providers will find it financially sound to adopt this approach. An alternative would be to deal with a HE provider that has service agreements with numerous network operators. This would then allow smaller scale service providers to have the benefits of accessing the network's capability indirectly through the HE provider. This alternative approach does

not deviate from the original VHE concept as the HE is still in charge of the overall provisioning of services to the user.

The other reason behind the justification of the VHE QoS API is that it provides a way for the service to state their QoS preferences and requirements in a network independent manner. It therefore alleviates the service of the burden of searching through the various types of QoS classes offered by the network to find one that suits its requirements by offering the QoS Manager to perform this task. To summarise, the QoS manager translates the QoS requirements of the service received through the VHE QoS API and tries to map these requirements to the QoS offered by the network.

4.2. QoS manager capabilities

The core responsibilities of the QoS Manager are the management of QoS streams on behalf of the user, the negotiation of the requested QoS with the network provider, and the notification of the delivered QoS to the service. The list below summarises the required capabilities of the server-side QoS manager.

1. QoS management on behalf of the user in both 'best effort' networks (if the network does not provide guaranteed QoS) and networks supporting guaranteed QoS.

2. Mapping of the service QoS parameters onto network-specific QoS parameters.

3. Dynamic QoS management including dynamic QoS monitoring and control. For instance the QoS manager needs to be able to handle the service QoS requests, provide guarantees to the service, and monitor the network against its agreement with the user. (Note that dynamic QoS management is not supported by OSA or Parlay).

4. Dynamic QoS negotiation. For instance the QoS manager needs to attempt to re-negotiate the QoS if it falls below the lowest guaranteed QoS traffic class. (Note that dynamic modification of a connection in the network is not supported by OSA or Parlay).

5. Coherent bandwidth allocation. For instance, it is important to ensure that neither the overall guaranteed throughput nor the mean throughput of the service(s) is greater than that of the user terminal's capabilities. If this is not the case, the QoS manager needs to issue a request to the service to reduce its throughput. (Note that this capability is not supported in the MExE QoS manager).
6. Trigger service adaptation. The QoS manager needs to be able to send notifications to the service whenever it cannot sustain service throughput or cannot guarantee a specific QoS level. This may, in turn, trigger service adaptation.

7. Priority management. In the case where the user accesses two services concurrently, and the sum of the total requested bandwidth is greater than the capabilities of the user's terminal, the QoS manager will refer to the user's preference to decide which service has higher priority in terms of bandwidth reservation. (Note that this capability is not supported in the MExE QoS manager).

Some of the capabilities mentioned in the list above require functionality from Parlay and OSA that are not currently supported. The next section will present the extensions and modifications needed to the Parlay and OSA in order for the QoS Manager to function appropriately.

5. Additional QoS support required from Parlay/OSA

This section presents the need for additional functionality in the Parlay/OSA APIs for 'QoS' aware services to function more effectively and efficiently via the server-side QoS manager. The proposed additions will also enable providers of such services to develop their applications with the confidence of interoperability in heterogeneous networks, and encourage creativity in the adaptability of the service in dynamic network conditions.

5.1. 'Change in QoS' notification and monitoring

The QoS manager has to be notified of any changes in the achievable QoS, enabling the application to adapt its service accordingly. For example, this change in the QoS may result from a dynamic change of QoS policy in an IP network, or it may be the result of a handover or bad reception in a mobile network environment. A change of QoS notification must therefore inform the application of the details about the new QoS level provided, which in turn may cause the application to decide whether a change of network is necessary to support the service. For example, in the case of a mobile network environment a different network provider offering better reception in the area may be used. For the case of a content-based service, the notification may serve to trigger an adaptation in the service provisioning (e.g. lowering the video quality etc.). As such, the QoS manager should have the option of stating the level of monitoring it requires from the network. Examples are:

-        Not to monitor for changes in quality of service;
-        To receive a notification when the change has occurred.

5.2. Dynamic setting of QoS preferences for service session

The QoS manager should be able to negotiate QoS with the underlying network connectivity provider on the user's behalf before starting a new service session. If the request is rejected, an alternative proposal should be sent from the network back to the QoS manager.

5.3. Dynamic modification of a connection's QoS

It is possible that the QoS provided by the network during the start of a service may be lower than that desired by the user. This can be due to reasons such as bad reception in the case of UMTS, or network congestion for the case of fixed IP networks. Therefore, in the case where a higher level of QoS is made available to the user during the service, the service provider should be able to initiate a sequence of upgrading the QoS on behalf of its users. A second example is to allow the user to dynamically change his or her QoS requirements.

5.4. Providing an independent means of exposing the QoS options available in the network

This functionality arises through the need to incorporate the VHE concept into heterogeneous networks and already exists in Parlay's connectivity manager interface. Parlay allows its users to discover the different classes of QoS provided by the network (e.g. Gold, Silver etc.). Users of the Parlay interface have two options – to select one of the classes with the default QoS attributes, or to modify attributes within the selected class. However, the attributes in the QoS class templates currently defined by Parlay are only specific to fixed networks. These templates should be enhanced to include those attributes shown in table 2 for UMTS networks [13]. Attributes that are not supported by the network should be indicated in the template accordingly. It would be convenient for

service providers if there were standards in place that require all network providers to define customisable classes of QoS that support the four types of applications stated in table 2, i.e. conversational, streaming, interactive, and background type of applications. The QoS manager could then be presented with an easy and uniform way of discovering the QoS capabilities of the network that are related to the type of service that they are providing.

Traffic class	Conversational class	Streaming class	Interactive class	Background class
Maximum bitrate	X	X	X	X
Delivery order	X	X	X	X
Maximum Service Data Unit (SDU) size	X	X	X	X
SDU format information	X	X		
SDU error ratio	X	X	X	X
Residual bit error ratio	X	X	X	X
Delivery of erroneous SDUs	X	X	X	X
Transfer delay	X	X		
Guaranteed bit rate	X	X		
Traffic handling priority			X	
Allocation / Retention priority	X	X	X	X

Table 2 UMTS bearer attributes defined for each bearer traffic class [13]

6. Case Study

The case study described here serves as an illustration of the concepts introduced in this paper. Let's now consider a typical sequence of events starting with a VHE service that upon initiation contacts the VHE server (Figure 5). During this phase, the service will request from the VHE QoS Manager the required QoS to be set on the connection between the service host server and the user terminal. The parameters of this operation include end-node addresses, the preferred and minimum acceptable values of QoS parameters, as well as the service type (conversational, interactive, etc.). For such a request, the QoS manager has to first map these QoS parameters from a service view to a lower level network view as supported by the underlying network technologies. The QoS parameters have to be checked against the networking capabilities of the terminal and the resulting data will be translated by the QoS Manager into a specific traffic class offered by the network. Finally, the resulting QoS is applied to the connection and the service is notified. We must note here that parameters for performance monitoring of the connection (used bandwidth, delay, etc.) may be applied along with the parameters denoting a traffic class. Performance monitoring plays a significant role of supporting service continuity in the VHE, by ensuring that services adapt gracefully as the provided QoS degrades to a "best-effort" connection. Additionally in all cases, monitored performance parameters keep the QoS Manager informed of the usage levels of the user's connections, allowing it to make more effective QoS decisions.
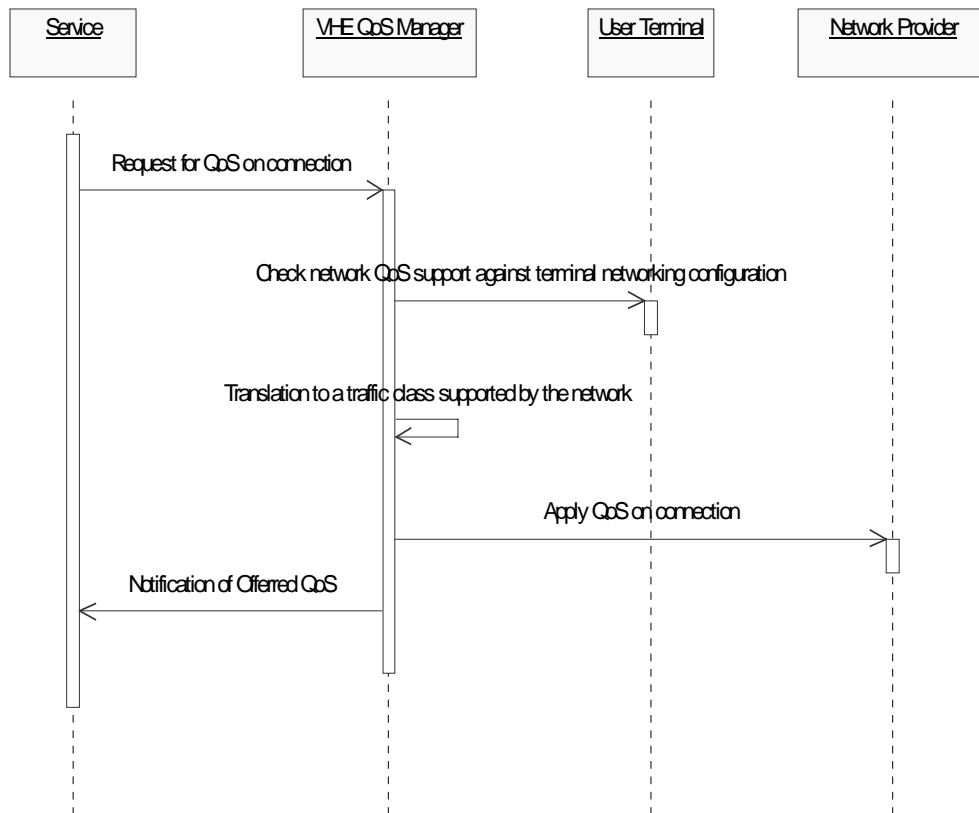
Figure 5   –   Typical scenario of request for QoS on a connection

We will now consider a variation of the above scenario where the service request cannot be fulfilled as smoothly as in the typical case described. Let's suppose that the VHE user initialises a second application instance on his terminal, which requires a new connection for its associated service to run (Figure 6). The service will again contact the VHE server and will request from the QoS Manager the required QoS just as before. After the initial mapping of the parameters from a service view to a lower level network view, the QoS manager concludes that one minimum acceptable value of a QoS parameter cannot be currently applied. As an example in our scenario we can assume that this parameter is the guaranteed bit rate, the requested value of which cannot be offered to the user because it goes over the terminal's networking capabilities. In this case the QoS manager based on performance monitoring data, the QoS of any other running services, the terminal networking capabilities and the user profile (users may have defined service priorities, etc.) will attempt to re-distribute the QoS on behalf of the user. If this process fails the user has to be notified that the newly requested service has been rejected by the network. If this process is successful, again translations of these parameters to traffic classes offered by the network takes place, the QoS is applied to the connections and the services are notified.
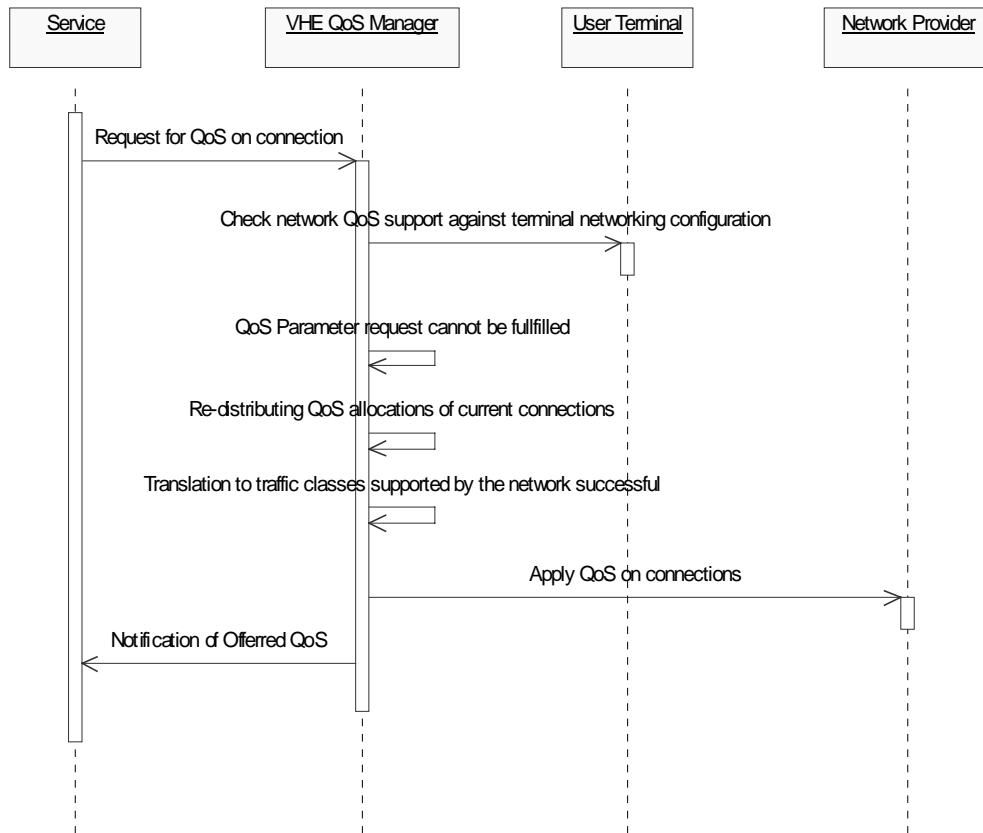
Figure 6   – Scenario involving re-distribution of QoS parameters.

7. Conclusion

The VHE concept introduces, among others, the idea of facilitating the job of service developers by providing them with a 'layer' that sits between the service (or the application) and the network. OSA/Parlay offer a unified set of APIs which hide network heterogeneity, exposing to the service designer particular views of the underlying network resources on the one hand, whilst allowing for a certain level of control on the network, on the other hand. In this sense OSA/Parlay can be seen as a first step towards network-independent service development. However, the VHE concept goes beyond that; it aims at supporting services which are 'adaptable' and 'QoS-aware'. This aim brings about the need for QoS management in the VHE that can serve both the user's and the VHE service's interests.

In this paper, we have discussed in particular the VHE requirements in terms of dynamic QoS management. Service adaptation to continuously changing network conditions requires constant and efficient QoS monitoring. By providing the VHE with the sufficient level of information on the quality of the underlying connection, it will be possible to trigger adaptation at service (or application) level. On the other hand, when the constraint is on the application –e.g. when the application cannot adapt or requires increased network quality– the VHE needs to be able to operate on the underlying network to address the request. Ideally this should be done through OSA/Parlay but current specifications do not provide the needed level of support. In this paper we have elaborated on QoS functionality required by the VHE, indicating the necessary enhancements to OSA/Parlay towards adaptable, network- and QoS-aware, service provisioning.

Besides providing support for services, QoS management in the VHE also involves satisfying demands on the user level. By taking advantage of user preferences, the VHE QoS manager is able to balance the QoS required by different concurrent services according to the terminal's capabilities. This functionality serves to ensure that the user is able to have the same service experience at all times even when using multiple concurrent services. By ensuring that the overall QoS provided by the network to the user conforms to the terminal used, an additional aspect of service to terminal inter-working was introduced to the VHE concept.

We have reported the initial results of our on-going work on QoS management for the VHE which is being carried out in the context of the IST VESPER project [14]. VESPER is tackling the issues involved in the VHE from a wider angle, including the introduction of a suitable VHE architecture and the realisation of service continuity, service portability, service personalisation, and session mobility. Further work will serve for the implementation of such architecture and the experimentation with the VHE concept.

8. Acknowledgements

9. References

[1] 3GPP Technical Specification 22.057 v5.0.0: "Mobile Station Application Execution Environment (MExE); Stage 1 (Release 5)", October 2000.
[2] 3GPP Technical Specification 23.057 v4.1.0: "Mobile Station Application Execution Environment (MExE); Functional Description; Stage 2 (Release 4)", March 2001.
[3] 3GPP Technical Specification 22.127v4.0.0: "Stage 1 Service Requirements for the Open Service Access (OSA)", January 2001.
[4] The Parlay Group. Home page at   http://www.parlay.org.
[5] 3GPP Technical Specification 22.121 v4.0.0: "The Virtual Home Environment (Release 4)", October 2000.
[6] GSM MoU: "3G Service Requirements and Concepts", PRG TG.21, V3.1.0, April 1998.
[7] ITU-T: "Report of the meeting held in Geneva from 5-22 May 1998 Part II". , Draft Recommendation Q.1711 (May 1998).
[8] UMTS Forum Report 1: "A Regulatory Framework for UMTS", June 1997.
[9] EURESCOM Project P920-GI: "VHE concept description, scenarios, and protocols", Deliverable 1, June 2000.
[10] IST VESPER project (IST-1999-10825): "VHE Requirements", Public Deliverable D21, November 2000.
[11] M. Nilsson, J. Hjelm, H. Ohto: "Composite Capabilities/Preferences Profiles: Requirements and Architecture", W3C Working Draft dated 21 July 2000, available at   http://www.w3.org/TR/CCPP-ra/.
[12] The Parlay Group: "Connectivity Manager Service Interfaces", Parlay APIs 2.1 Issue 1.1, June 2000.
[13] 3GPP Technical Specification 23.107 v4.0.0: "QoS Concept and Architecture (Release 4)", December 2000.
[14] Information Society Technologies (IST) Virtual Home Environment for Service Personalization and Roaming Users (VESPER) project. Home page available at http://vesper.intranet.gr.