

Supporting Adaptation-aware Services through the Virtual Home Environment

Antonio LIOTTA, Alvin YEW, Chris BOHORIS, George PAVLOU

Center for Communication Systems Research
University of Surrey, Guildford, GU2 7XH, UK
Tel. +44 1483 689808 Fax +44 1483 686011
{A.Liotta, K.Yew, C.Bohoris, G.Pavlou}@eim.surrey.ac.uk

Abstract – This article focuses on the service management aspects of Virtual Home Environment (VHE) -like middleware for context-aware, adaptable services. We present a functional architecture which is in line with the on-going VHE standardization efforts. We also describe our prototype implementation developed in the context of the IST VESPER project with particular emphasis on the realization of functionality for context-handling and adaptation triggering. Our system is illustrated through a case study, which allows us to highlight the advantages of a middleware-based solution that decouples the adaptation logic from the application and network layers. An important part of our study looks into how our system may be realized outside a laboratory environment, considering two alternative scenarios. First, we consider the case in which the VHE service management logic resides within the Network Operator boundaries. Then we elaborate on an alternative scenario in which service management is achieved outside of the Network Operator, relying on an increased level of network openness not yet available today.

Keywords: Service Management, Service Adaptation, Adaptable Context-aware Services, Virtual Home Environment.

1. INTRODUCTION

The rapid evolution of fixed and mobile networks witnessed during the last decade has created new opportunities to deliver increasingly advanced services to the mobile user. The 2nd Generation (2G) Global System for Mobile Communications (GSM) has been widely deployed with enormous success, despite its very limited support for packet data services. Following the popularity of the Short Messaging Service (SMS), the industry has responded to the increasing demand for Internet-like services with the introduction of the Wireless Application Protocol (WAP). However, the type of services that can be provided through a WAP terminal is greatly limited mainly by the circuit-switched network access (which does not support well asynchronous application events) and by the poor data rate. The General Packet Radio Service (GPRS) is removing some of these limitations by offering packed data services, while the bandwidth is gradually increasing towards the Universal Mobile Telecommunications System (UMTS) or 3rd Generation (3G). 3G networks combine a choice of Quality of Service (QoS) levels with increased bandwidth, enabling bandwidth-intensive applications and paving the way towards QoS-aware applications for the mobile user. It is envisioned that the issues related to heterogeneous wireless networking will be further addressed by 4th Generation (4G) systems, which may provide outdoor connectivity in the order of 100 Mbit/sec and improved end-to-end IP services.

The features of 3G and 4G networks offer a potential vehicle to address the demand for future advanced services and increase the revenues of the various players in the telecommunications arena – i.e. network operators, service providers, content providers or value added service providers, and application developers. Research in developing and managing services that can meet the expectations of the mobile user by taking full advantage of the latest network technologies is very important. This has led to a recent surge of activities towards advanced service provisioning concepts and service creation frameworks. Open Application Programming Interfaces (APIs), such as Parlay [Parlay01] and the Open Service Access (OSA) [3GPP01a], enable hosting applications outside specific networks while giving them controlled access to the network resources. In this context, the Virtual Home Environment (VHE) [3GPP01b] concept plays an important role and will be implemented in 3G mobile communications through OSA and other 3G specifications. The VHE concept allows users to be consistently presented with the same personalized features, user interface personalization, and services regardless of network, terminal, and location.

While a variety of VHE services are likely to populate the future telecommunications arena, the realization of such services still poses numerous problems. A first, vital step involves the creation of a standardized, open framework which grants to the applications a controlled access to network capabilities through well-defined, object oriented APIs. Standardization in this area is currently driven by several consortia comprising network operators, equipment vendors, software suppliers, and terminal manufacturers.

Another important problem regarding the realization of VHE services is the need for a viable architectural solution which allows effective service development, deployment, and management. While service development and deployment will be facilitated by the standardization of open network and service APIs, the management of VHE-like services presents challenges that have received only scant attention so far. The key problem in managing such services is to provide effective adaptation mechanisms. As the user roams to a foreign network, changes his/her preferences, switches to a different terminal, etc, the management system should transparently provide an ‘adapted’ service, preventing service interruption.

The current approach to providing adaptable services or applications is based upon the classic layered architectural model where adaptation is provided at the various layers (data link, network, transport or application layers) in isolation [Haas01]. Because very limited information is shared across protocol layers, adaptation strategies are *uni-dimensional* – i.e. they only consider one parameter at a time – often leading to unsatisfactory results. For example, in a video-streaming application such as RealPlayer, adaptation may be simply achieved at application level by monitoring the percentage of buffer utilization at the client and appropriately controlling the bit rate. This approach may be effective in best-effort IP networks but does not suit QoS-enabled networks where end-to-end performance monitoring is necessary. Such an application provides only a simple form of ‘adaptation to network conditions’; it does not cater for terminal adaptation (i.e. a new version of the application is required if the user wishes to change terminal); and it only allows rudimentary network control mechanisms by the application.

In contrast, the level of adaptation required by VHE-like services is generally *multi-dimensional* because of two reasons. First, adaptation relies on a number of input parameters, i.e. the *service context*, which includes the *user context* (i.e. location, time, terminal capability, and preferences) and the *network context* (entailing the network QoS parameters). Second, adaptation problems generally have more than a single solution. Possible adaptation strategies include: 1) having the service adapting to the context (e.g. controlling the video stream bit rate depending on network QoS and terminal display characteristics); 2) having the service influencing the context (e.g. allocate network resources according to the requested video stream bit rate) to suit service requirements and/or user preferences; or 3) a combination of the previous two approaches in order to achieve maximum user satisfaction. The multi-dimensionality of the adaptation problem becomes even more obvious if we consider more advanced services such as location-dependent or time-dependent services, where the context presented to the user depends not only on his/her terminal, network and preferences, but also on his/her geographical location or time of the day.

This article focuses on the service management aspects of a VHE-like middleware for context-aware, adaptable services. We present a functional architecture which is in line with the on-going VHE standardization efforts. We also describe our prototype implementation developed in the context of the IST VESPER project [IST00a] with particular emphasis on the realization of functionality for context-handling and adaptation triggering. Our system is illustrated through a proof-of-concept case study which allows us to highlight the advantages of a middleware-based solution that decouples the adaptation logic from the application and network layers. An important part of our study looks into how our system may be realized outside of a laboratory environment, considering two alternative scenarios. First, we consider the case in which the VHE service management logic resides within the Network Operator boundary. Then we elaborate on the alternative scenario in which service management is achieved outside of the network operator, relying on an increased level of network openness not yet available today.

2 THE VHE CONCEPT

The VHE concept has been described by different consortia, as depicted in **Table 1**. The idea behind those definitions is that a VHE system should facilitate the development and deployment of *interoperable* services, which are also *portable* and *accessible* across heterogeneous network boundaries. 3GPP has also included the requirement of portability across heterogeneous terminals. Hence, the VHE concept pursues the idea of *service universality* which allows a user to transparently access services at anytime, from anywhere and from any terminal.

	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” [3GPP01b]
GSM MoU	“Virtual Home Environment (VHE) is a system concept for service portability in the Third Generation across network borders” [GSM98]
ITU / IMT2000	“VHE is a capability whereby a User is offered the same service experience in a visited network as in his Home system.” [ITU98]
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).”[UMTS97]
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” [EURESCOM00]
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” [IST00b]

Table 1. VHE definitions.

Beyond its requirements, the VHE is developing into a layered architecture which separates the network layer from the service layer through a well-defined set of interfaces that are currently subject to standardization. On the one hand, network heterogeneity will be hidden through OSA and Parlay interfaces. These allow applications to discover the network functionality exposed by the network provider and to have a controlled access to network resources. On the other hand, the VHE offers useful building blocks to the service such as the OSA Service Capability Servers (SCSs) or other specialized Parlay APIs. In such a way, application developers can rely on pre-existing network-independent functionalities such as those for call control or location management.

In addition to OSA, other VHE-enabling technologies are being developed. These include: 1) the UMTS SIM Application Toolkit (USAT) [3GPP01d], which allows applications residing in the SIM card of a mobile terminal to interact with services located on a SAT server; 2) the Mobile Station Execution Environment (MExE) [3GPP00], which allows the development of applications that can run on different types of terminal; and 3) the Customized Applications for Mobile networks Enhanced Logic (CAMEL) [3GPP01c], providing the mechanisms to support Intelligent Network (IN) type of services in 3G mobile networks.

Standardized APIs and VHE-enabling tools provide a common framework for service developers, service providers, and network operators and are meant to facilitate service development and deployment. However, the type of services that can actually be supported depends critically on the degree of adaptability exhibited by the VHE system. In this article, we argue that sophisticated, context-aware services involve the solution of multi-dimensional adaptation problems which requires appropriate service management support by the VHE. This contrasts an alternative approach in which it is the application that, having access to system resources through the VHE, holds the service adaptation logic.

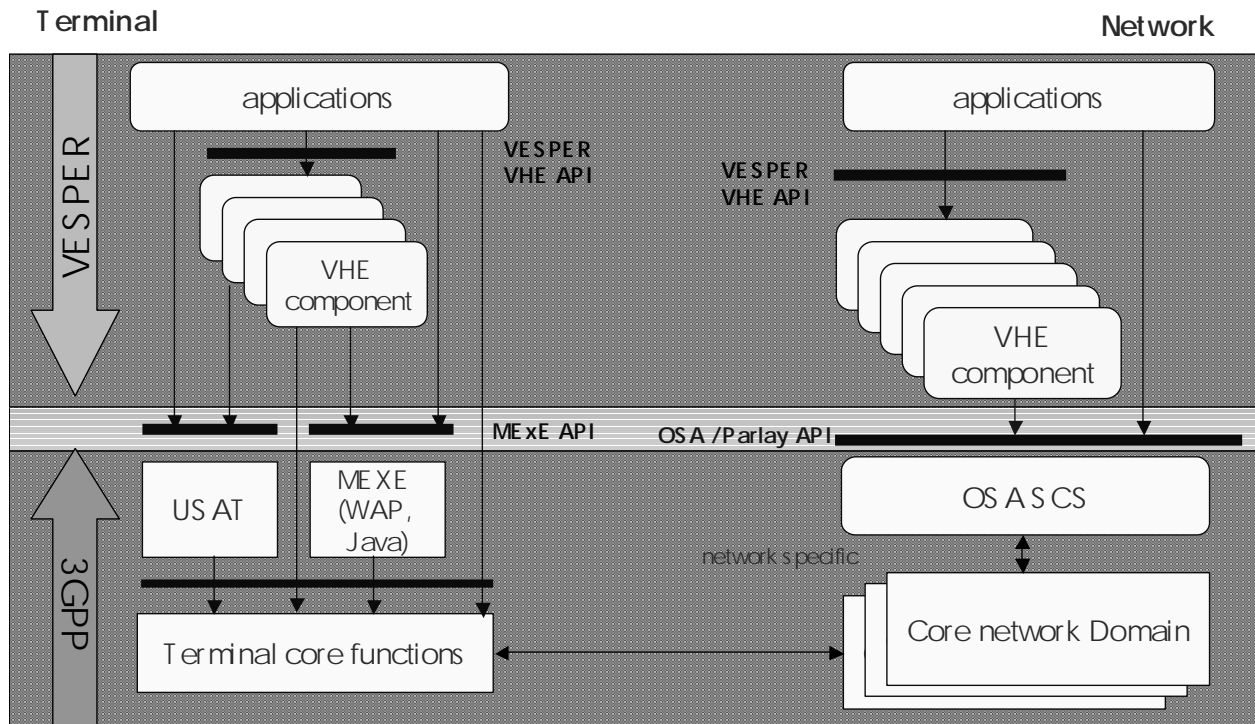


Figure 1. The VESPER VHE functional architecture.

Our approach to the VHE is complementary to that followed by 3GPP. While it adopts interfaces, APIs, and tools that are being developed under the 3GPP umbrella, the functional architecture of the IST VESPER project focuses on the development of and experimentation with core VHE components (**Figure 1**). Components are meant to offer to the service developer and administrator high level APIs in order to alleviate them from tasks that can or are better performed by the VHE. VESPER components include the ones for Access, Profile, Discovery, Account, Subscription, Authentication, Security, Service Interaction, Upload, Session, Location, Connection, and Adaptation [IST00b, IST01]. In this article, we focus on the latter three components. Below, we first illustrate VHE services including context-driven and adaptable services. Then, we focus on the design of a VHE middleware system offering run-time service adaptation management functionality to those services.

3 CONTEXT-AWARENESS AND SERVICE ADAPTATION IN THE VHE

To better illustrate some of the features of VHE services we make use of the following use case scenario. James, subscribes with his VHE provider to two complementary services, a *news service* and a *road traffic service*. These services are maintained by a content provider who keeps the necessary information up to date. They are also customizable by the user through a simple VHE, web-based interface. Adaptation with regard to context – i.e. user preferences, location, time, terminal, and network conditions – is supported by the VHE service management middleware.

James configures the news service in such a way as to receive different types of news, depending on time of the day and day of the week. If he activates the service on any weekday during the morning James wants the latest world business developments; in afternoon hours he would be interested in the local news; in the evening he would like a summary of the world news of the day; finally, on a week-end day James only wants to know about sport. It is the VHE that will handle such user preferences and make sure that James receives the relevant news depending on the time of service activation.

In addition, the VHE needs to deliver the news content in a format which suits James terminal characteristics and the available network QoS level. James uses a desktop with a 56kbps modem when he is at home; a desktop connected to a 100Mbps local area network when he is at work; and a Java-enabled mobile phone while he is commuting by car. The service management middleware determines James' context and deduces the appropriate service parameters to apply before initiating the service. In such a way, he will receive the news in the form of a low-quality video and audio stream when sitting at home; a full-quality video and audio stream when sitting at work; and a stereophonic audio-only stream while on his car.

The middleware will also provide *service continuity*, allowing James to temporarily suspend / resume service delivery as he changes terminal. Run-time adaptation is also required to cater for network QoS fluctuations while James is on the move. As the QoS decreases, the VHE middleware receives an event notification from the network and triggers service reconfiguration to lower the bit rate, resulting for instance in a mono audio version of the service.

Finally, James configures the traffic service in such a way as to be notified of any local traffic hot spots as soon as he

connects to the service. The service adaptation support required for this service is analogous to the previous one. In addition, the VHE needs to deliver only a selection of the hot traffic spots of the country which matches James location.

The above case highlights the need to support context-awareness including various forms of adaptation. *Context-aware services* possess the capability of sensing and receiving feedback from the environment, or context, in which they operate, thus offering to the user potential benefits such as greater responsiveness, flexibility, adaptability, and customizability [Mandato02]. There are many ways in which a service can take advantage of its context. *Context-driven service configuration*, allows the configuration and delivery of a service (service components, parameters and content) in a different way, depending on the context. For instance, different codecs may be needed at client/server side to deliver the news content to James, depending on his preferences (e.g. audio/video quality), terminal capabilities (e.g. display resolution, color depth, network card), and available end-to-end QoS level.

Another important feature is *context-driven run time adaptation*. In this case the service will exploit context information to trigger appropriate adaptation actions with the ultimate goal of meeting the user expectations regardless of location and terminal changes as well as network fluctuations. One or more adaptation actions may be possible for a given service and context. The service content may be adapted in order to accommodate network-changing conditions – e.g. the video stream compression rate of the codec may have to be dynamically modified in response to a varied end-to-end QoS level. Alternatively, there may be an option for the service to dynamically request additional network resources in order to satisfy James’s request for increased video resolution. Also, multimedia content adaptation is necessary when James switches between Desktop and mobile phone, in order to satisfy the VHE requirement of session continuity.

Context-awareness can also be exploited for *time- and location-driven service activation*. James’s news service is an example of time-dependent service, since the information that is actually delivered changes depending on time. Instead, the traffic service is an example of a location-dependent service since the traffic hot spots notified to James vary depending on his geographical location.

Having discussed VHE services and highlighted their requirement for adequate management, the next section describes the design and prototype of our VHE middleware system with particular focus on mechanisms to handle the context (including mechanisms for sensing and re-configuring it) and to trigger service adaptation actions.

4 A VHE MIDDLEWARE FOR ADAPTIVE SERVICE MANAGEMENT

4.1 Functional Architecture

The VHE service management functionality consists of a set of software components logically located between the network and the service management layer, as in **Figure 2**. Context-awareness and adaptation are managed by a group of context-handling components and by the adaptation component, respectively. The former components play a twofold role. On the one hand they gather the context in order to adapt the service to it. Context-gathering components include the location, time, terminal capability, user preferences, service configuration parameters, and network performance monitoring components. On the other hand, context-handling includes functionality needed to influence the context in order to accommodate service requirements and user preferences. The QoS management functionality is used to negotiate network resources and allocate them to individual services. The context can also be influenced through the ‘service settings’ component which controls parameters such as video stream compression ratio, image resolution, audio sampling period, and so on.

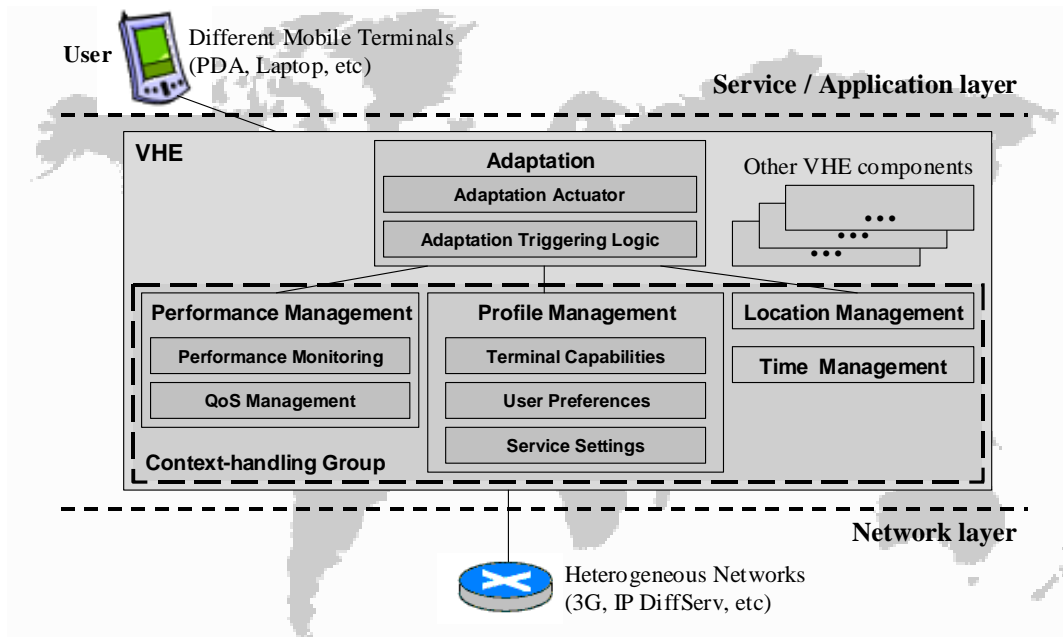


Figure 2. Architectural view of the VHE service management functionality.

While the context-handling components provide the instruments to sense and influence the user and service environment, the decision system resides in the adaptation component. It is here that the multi-dimensional adaptation problem is solved. The Adaptation Triggering Logic (ATL) incorporates the service adaptation algorithm which, by making use of context information, decides on the appropriate adaptation action, triggering the relevant operations through the Adaptation Actuator (AA). The AA mediates between the ATL, which is system independent, and the other components. In such a way, the ATL can be more readily re-programmed in order to accommodate new service adaptation strategies as needed.

Adaptation is performed in two ways. First, at service activation time, the adaptation component analyses the context in relation to the user's preferences and triggers the appropriate service configuration actions (*context-driven service configuration*). Then, at service execution time, the adaptation component relies on real-time notifications of any changes, in order to consider additional adaptation actions (*context-driven run time adaptation*). At run time, adaptation relies mainly on notifications coming from the network, through the network performance monitoring component, and from the terminal. As the user switches between different terminals, the Terminal Capabilities component derives the physical capabilities of the new terminal through the Content Capability/Profile Preferences (CC/PP) protocol (currently being specified by W3C) [Ohto99, Nilsson00]. To experiment with the above forms of adaptation, we have prototyped a system and realized an adaptable video streaming service which are described below.

4.2 The Prototype System

Based on what has been described so far in this article our prototype system involves the components and interfaces of communication shown in Figure 3.

Starting with the network aspects of the system three interfaces of communication are specified for the interactions with the performance monitoring and QoS management components:

- `IpAppPerformanceMonitoringEvent`: Allows notifications and reports to be sent to the VHE adaptation component keeping it informed of network conditions affecting the service.
- `IpPerformanceMonitoringControl`: Allows the configuration of the performance monitoring task.
- `IpQoSManagementRequest`: Allows the dynamic configuration of network QoS guarantees.

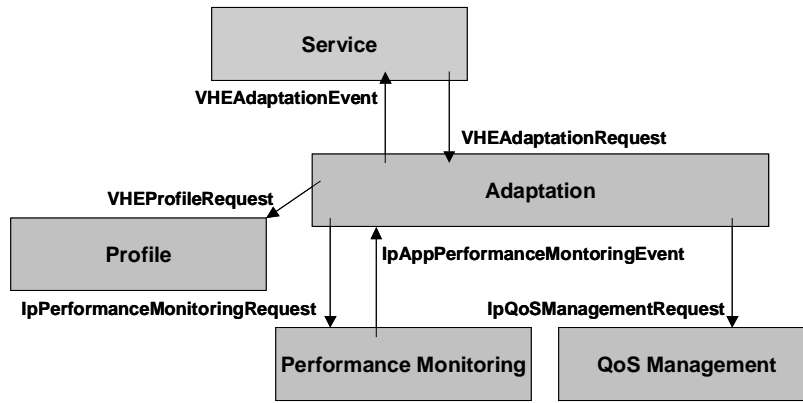


Figure 3: Prototype system design

This set of interfaces is compliant with the design guidelines of the Parlay group, as specified in Figure 4. This choice is motivated by the strong industrial support currently enjoyed by the Parlay group. However, we have had to design and develop from scratch the performance notifications and dynamic QoS configuration facilities of our system, which are crucial to the VHE adaptation process and are not currently specified in Parlay.

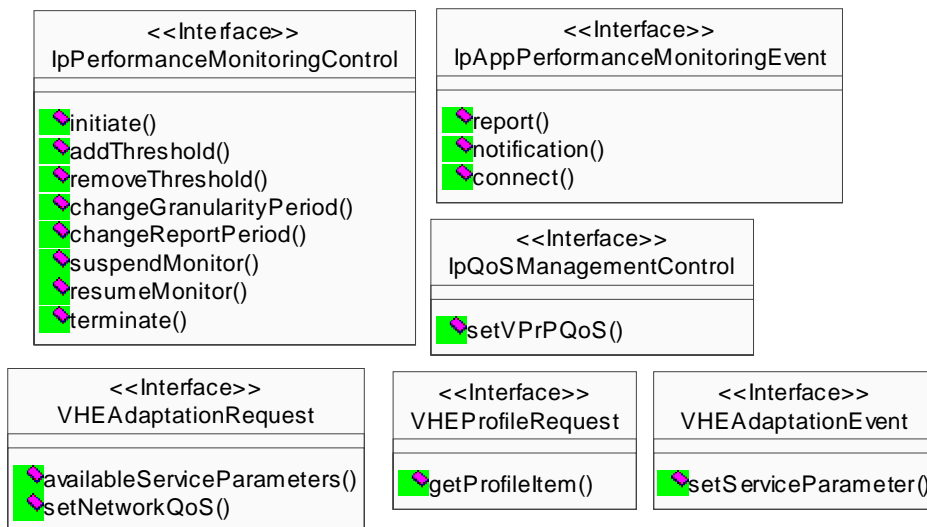


Figure 4: Details of system interfaces

In addition to the network QoS management interfaces, the VHE adaptation component requires the following interfaces:

- **VHEAdaptationEvent**: Provides a generic way for the adaptation component to instruct on the adaptation of a service parameter.
- **VHEAdaptationRequest**: Provides a generic way for the service to request the adaptation of the network QoS guarantees offered. In addition, it provides a means for the service to inform the VHE adaptation of the service parameters available for adaptation.
- **VHEProfileRequest**: Allows the adaptation component to obtain profile parameters including terminal characteristics, user preferences and service settings.

Our prototype system was tested on an IP/DiffServ network environment although our interfaces do not assume any particular network technology, terminal or service. We are also making no assumptions on a specific engineering approach chosen for the implementation of functionality. In our case and because of the dynamic and ideally programmable VHE network environment our approach involved the migration of management logic in the appropriate network nodes through the use of mobile agent technology.

4.3 Case Study: Adaptable Video Streaming

In order to test the behavior of our system we have developed an adaptable multimedia streaming application suitable for stressing the network QoS aspects of the adaptation process, as depicted in Figure 5. This application can be considered as being part of the use case introduced in Section 3.

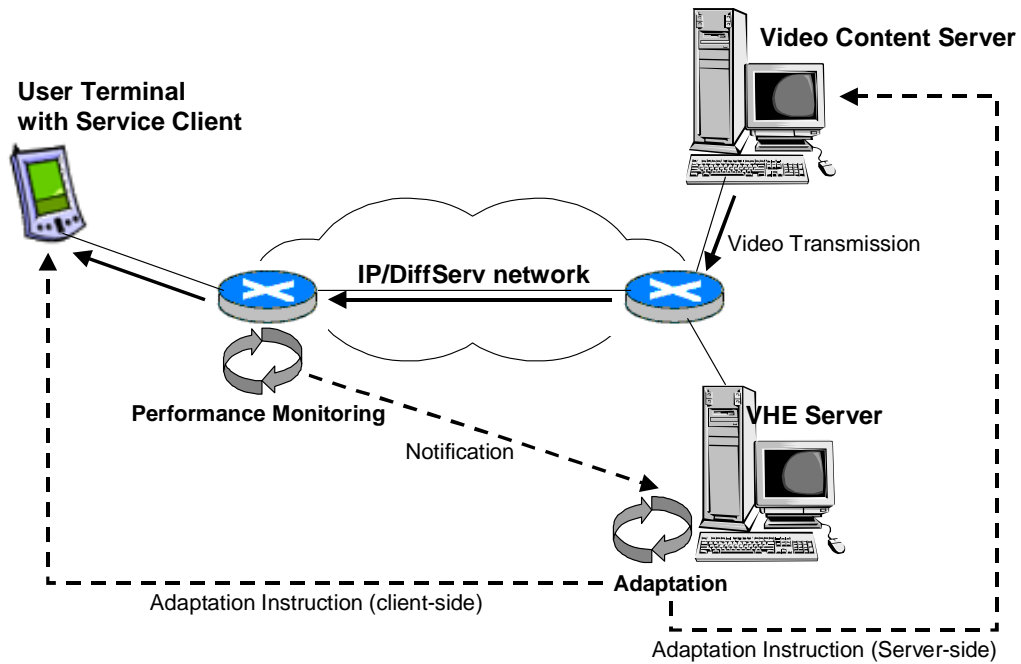


Figure 5: An adaptable video streaming service

The service involves the transmission of a video stream from a video content server to the user's terminal. Upon initiation, the service is responsible for informing the VHE adaptation component of the specific service parameters supported for adaptation (e.g. bit rate of video transmission, frame size, etc.). The adaptation component will subsequently use the input received from the context-handling components to set up these service parameters before commencing the video transmission.

The user of the VHE is given the capability to suspend a service session from the current terminal and resume it from a different one; in which case the VHE will respond as depicted in **Figure 6**. The VHE, obtains the new terminal characteristics, analyzes the information along with the user preferences and service settings, and finally instructs the service to adapt its terminal-dependent parameters. For example, this is what would happen to James as he prepares to leave home and switches his current session from his home desktop to his mobile terminal. As soon as James logs from his mobile, the VHE will adapt terminal-dependent, service parameters such as: frame size, display color depth, network buffer size, CPU usage priority, and so on.

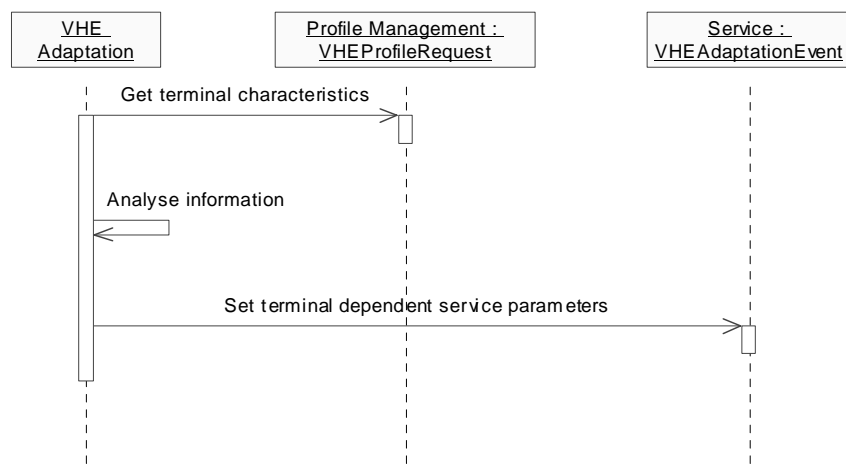


Figure 6: VHE terminal adaptation scenario

Two forms of network adaptation are supported by our system. In the first case, service adaptation is triggered by the network (**Figure 7**). The network performance monitoring system detects changes in the network which are bound to affect service delivery and sends an appropriate notification to the adaptation component. This, in turn, processes such a notification against service settings and user preferences and sends specific adaptation instructions to the service. Referring to James again, adaptation to network may be required as he travels, if a reduced QoS level is compensated by the VHE with a signal directed to the service, which requests a reduction in bit rate.

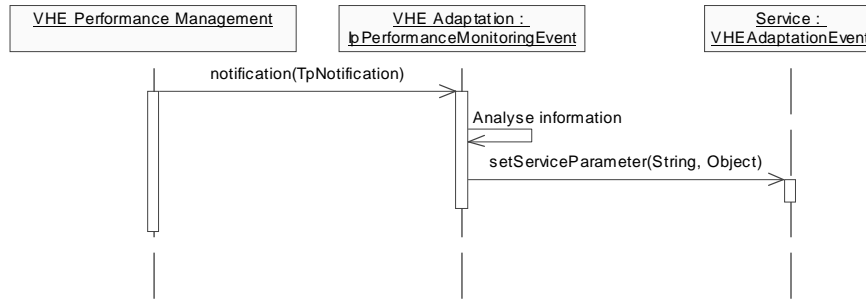


Figure 7: VHE network adaptation triggered by changing network conditions

A second form of network adaptation is triggered by the user, who is allowed to dynamically request a better service QoS level, reflecting on a change in network QoS (**Figure 8**). In this case, the adaptation component checks the validity of the user’s request and subsequently maps the request into a suitable network QoS. In our system, user demands are mapped into appropriate IP/DiffServ QoS levels that are handled by the QoS management component. James, may take advantage of this facility if he is not satisfied with the quality of the streaming application. He may simply request an increased frame rate, a better image resolution, or a high-fidelity sound. His request would be received by the VHE that would, in turn, trigger network adaptation through the QoS component.

Clearly, an additional form of adaptation may involve a combination of the above.

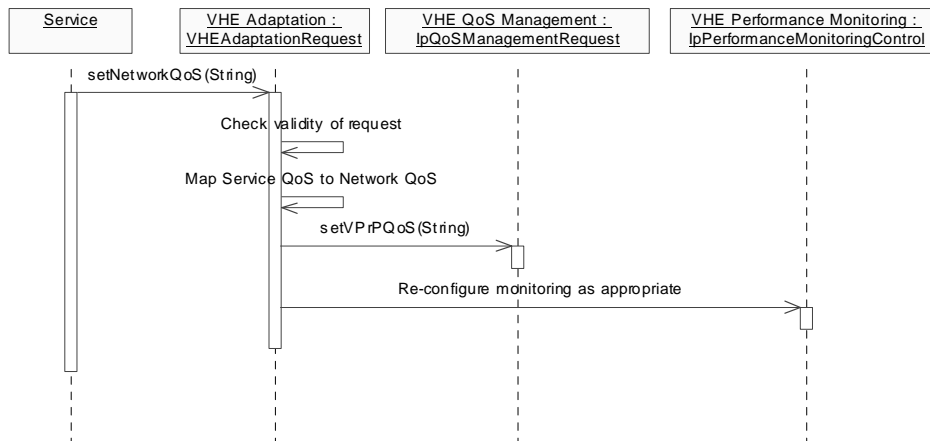


Figure 8: VHE network adaptation triggered by the user

Let’s now assume a simple example that allows us to illustrate the operation of a simplistic algorithm for adaptation triggering logic. The scenario involves a video stream encoded at 64kbps, a terminal that can receive data through a 128kbps link and a VHE adaptation component performing service-to-network adaptation. In this case, the following may take place, as depicted in **Figure 9**:

1. Initially the video can be transmitted at the appropriate 64kbps rate. The Adaptation Triggering Logic (ATL) sets a network performance threshold at 10% lower or 51.2kbps.
2. At time t1 the threshold level is exceeded, the ATL is notified and takes two actions. First, it reduces the bit rate of transmission by 10% (at 51.2kbps). Subsequently, it places a new threshold at 10% higher than the crossed threshold.
3. At time t2 the 64kbps threshold level is exceeded. The ATL reduces the bit rate of transmission by 10% (at 38.4kbps). Subsequently, it removes the 51.2kbps threshold and adds a threshold at 10% higher than the crossed threshold.
4. At time t3 the threshold is cleared. The ATL increases the bit rate of transmission by 10% (at 51.2kbps). Subsequently, it removes the 76.8kbps threshold and adds a threshold at 51.2kbps.

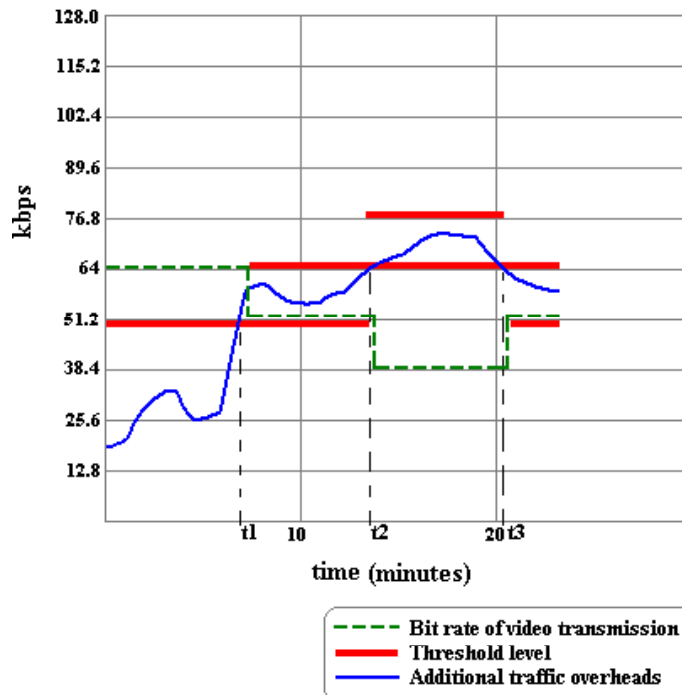


Figure 9: Time diagram of the effects of VHE adaptation

Clearly, more sophisticated ATL algorithms may be embedded into the system, leading to systems of different complexity, reactivity, and user satisfaction. If the ATL is seen as a programmable control system of a VHE, a variety of solutions can be designed for a rich set of services.

5 FROM PROTOTYPE TO REALIZATION

There are two major factors that will influence the successful management of context-aware, adaptable services in the VHE – feedback and reactive ability. Feedback aids a management platform in deciding the optimal method of adapting services, and is performed in our management system by the context-handling functionality. Reactive ability involves the actual adaptation process of the service. A successful service management platform should support and offer services the means of multi-dimensional adaptation. This section discusses the realization strategies of our platform in the VHE and a comparison of their impact on the two abovementioned factors for service management.

We adopt two realization strategies for our system that are based on the general interpretation of the VHE business model in the research community. The traditional VHE business model is based on 3GPP’s standards on the VHE and emphasizes the role of the network operator in providing VHE functionality and management. Case a) in **Figure 10** shows our system embedded in the network operator domain and working alongside the Parlay/OSA APIs and Service Capability Servers (SCS) as was explained in section 2. Our service management middleware can be seen as providing extensions and an additional service capability functionality to the Parlay/OSA framework, and can therefore be implemented as an SCS. Current network operators are selective over the service providers that can gain access to their user’s private profile and context through such APIs. Service providers that have such exclusive access are represented in the 3GPP business model as Home Environment Value-Added Service Providers (HE-VASP). As the nature of our system has a larger share of the responsibility in the adaptation process than the VHE service, the service does not need to directly access a user’s context or profile to perform the adaptation. The bonus of supporting user context confidentiality in our system allows a Value Added Service Provider (VASP) that does not have full access to the user’s profile through the Parlay/OSA APIs to provide advanced adaptable services within the confines of the 3GPP business model. Placing our system and the core of the adaptation logic in the network operator’s domain would increase the efficiency of gathering the user context as this could be done without the need of going through an intermediary as would be required if the middleware was placed externally to the network operator’s domain. As such, the contextual information may also be more accurate as the information would not be remapped to cater to an open external API provided by the network operator. For example, the accuracy of a user’s location in a cellular network for location-based services can be based on the network topology (e.g. cell number). Such accuracy is wasted if the user location is translated to the local PSTN area code as exemplified by the current practice of supporting location-based WAP services by network operators.

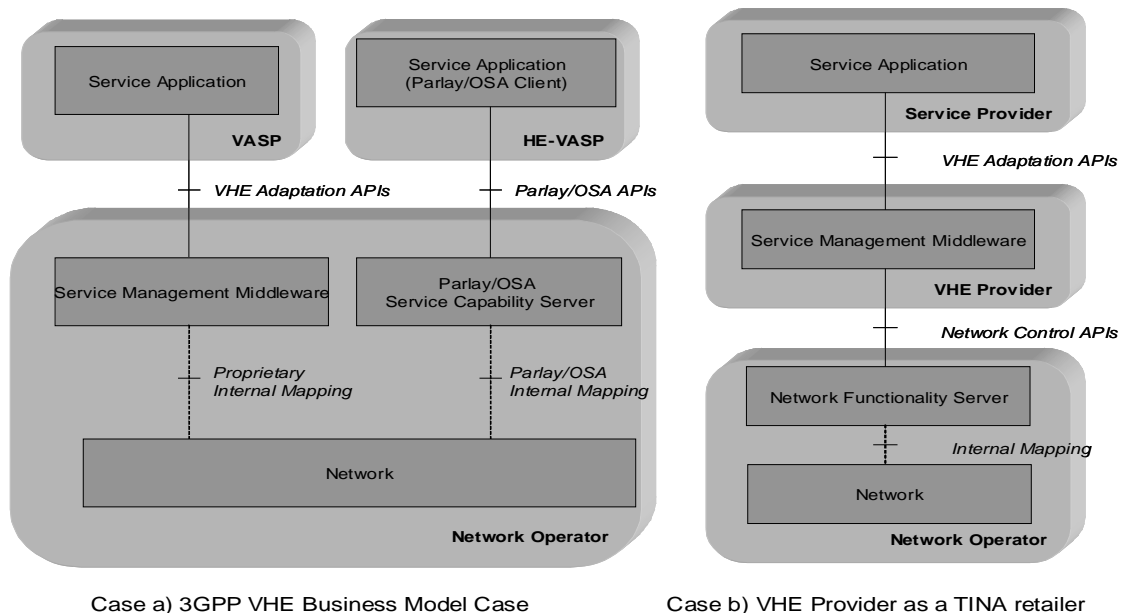


Figure 10. Realization strategies for the VHE service management middleware

The other interpretation of the VHE business model shown in case b) of **Figure 10** is derived from that of the TINA consortium and relies on an open service provisioning environment where the VHE provider represents the role of the TINA retailer and is responsible for managing the user home environment [Farlay]. In this model, our system would be housed within the VHE provider domain and would require network operators to allow access to the user contextual data. The benefit of this model is an improvement in scalability and extensibility as our system can operate without being tightly bound to the intrinsic capabilities of a particular network type or network operator. Therefore, the success of our system is dependent on the functionality of the API used between the VHE provider and the network operator. In this scenario, it would hence be possible to extract the user's terminal capabilities through the API without knowing whether it originated from a CC/PP proxy, WAP gateway or MExE server located within the network operator's domain. The advantage of using such an approach is that such contextual information would inherently be translated to provide openness and independency from the technology used in gaining it, thus resulting in a generalization of the information gained. The authors, however, do not envision this scenario to be adopted readily in the near future as the business politics of customer ownership is seen to be a hindering factor by network operators.

6 CONCLUDING REMARKS

Our development of a prototype VHE adaptation system gave us practical insights on a number of interesting issues. One of our aims was to design a system based on state-of-the art technology that would facilitate the development of context-aware, adaptable applications and services. We have achieved that by relieving the application from the complex tasks of context- and adaptation-handling. In doing so, we have followed the VHE architectural approach which neatly separates the domains of content, service, and network provisioning. The result was a middleware layer which manages service adaptation, offering to the service a bi-directional control mechanism. On the one hand, the service can dynamically request additional resources in order to increase user satisfaction. On the other hand, the service may be instructed to adapt to real-time decrease of resources in order to achieve a graceful degradation at user level.

Our work complements the on-going standardization efforts lead by 3GPP by suggesting a way for materializing the VHE concept rather than just indicating APIs that will facilitate its realization. Prototyping the service adaptation middleware has taught us a number of lessons. First, this approach does facilitate service development by concentrating the complexity – i.e. adaptation and context-handling – in the middleware, outside of the application itself. We have validated this claim by implementing an adaptable multimedia streaming application which relied on the interfaces depicted in **Figure 3** and **Figure 4**.

Another important advantage of the middleware-based solution is that it allows us to easily modify the service adaptation behavior or extend its functionality to reflect on changing requirements. This can be achieved by re-programming the ATL component whilst keeping the service logic unchanged. In fact, services rely merely on triggers or signals exchanged with the adaptation component. It is the ATL that produces those triggers and manages the service context. This separation of management functionality from the application has also the advantage of enabling sophisticated multi-dimensional adaptable services.

The main performance issues that arise from our experimentation relate to the latency of the adaptation control loop. For its correct operation, the system has to process real-time contextual information and trigger the appropriate adaptation action in a timely fashion. In these lines, we have measured the period of time elapsed between the generation

of a network notification and the moment in which the application adapts. Sample measurements taken on the testbed of **Figure 5** produced an average response time of 34 milliseconds which is acceptable for most adaptable services. Clearly, our system operates in ideal conditions since it is bound to a LAN rather than being evaluated on a global Internetwork. However, our results give an indication of the lower-bound achievable on response time. In a real operating environment, a VHE will also have to face the challenge of scalability. Scalability issues deserve further attention, although our architecture is distributed in nature and assumes several VHE providers. Distribution and appropriate resource allocation to VHE provider is key to scalability.

The service management middleware described herein is part of a wider project, IST VESPER, aiming to specify, implement and validate a complete architecture for the VHE concept. VESPER is currently implementing a system composed by several components required by a VHE. Future work includes the integration of our service adaptation middleware system in VESPER. This will allow the experimentation with large-scale, integrated fixed and mobile networks and the experimentation with more advanced, adaptable telecommunication services.

REFERENCES

- [3GPP00] 3GPP Technical Specification 22.057 v5.0.0, "Mobile Station Application Execution Environment (MExE); Stage 1 (Release 5)", October 2000.
- [3GPP01a] 3GPP Technical Specification 22.127 v5.0.0, "Service Requirement for the Open Service Access (Release 5)", June 2001.
- [3GPP01b] The 3GPP Technical Specification TS 22.121 v5.1.0, "The Virtual Home Environment (Release 5)", June 2001.
- [3GPP01c] 3GPP Technical Specification 22.078 v5.1.0, "Customised Applications for Mobile network Enhanced Logic (CAMEL); Service description, Stage 1(Release 5)", January 2001.
- [3GPP01d] 3GPP Technical Specification 22.038 v5.2.0, "USIM/SIM Application Toolkit (USAT/SAT); Service description; Stage 1 (Release 5)", June 2001.
- [EURESCOM00] EURESCOM Project P920-GI: "VHE concept description, scenarios, and protocols", Deliverable 1, June 2000.
- [Farlay] P.Farley,R.Minetti, M. Mampaey, "Retailer Reference Point", TINA-C specification version 1, available at <http://www.tinac.com/specifications/documents/ret111.pdf>
- [GSM98] GSM MoU: "3G Service Requirements and Concepts", PRG TG.21, V3.1.0, April 1998.
- [Haas01] Z.J. Haas, *Designing Methodologies for Adaptive and Multimedia Networks*, IEEE Communications Magazine, pp.106-107, Vol. 39, N.11, November 2001.
- [IST00a] Information Society Technologies VESPER project, Homepage at <http://vesper.intranet.gr>.
- [IST00b] IST VESPER project (IST-1999-10825): "VHE Requirements", Public Deliverable D21, November 2000.
- [IST01] IST VESPER project (IST-1999-10825): "VHE Architecture Kernel Specification", Public deliverable D32, April 2001.
- [ITU98] ITU-T: "Report of the meeting held in Geneva from 5-22 May 1998 Part II". , Draft Recommendation Q.1711, May 1998.
- [Mandato02] D. Mandato et al, *CAMP: A Context-Aware Mobile Portal*, IEEE Communications Magazine , pp. 90-97, January 2002.
- [Nilsson00] M. Nilsson, J. Hjelm, H. Ohto, "Composite Capabilities/Preferences Profiles: Requirements and Architecture", WWW Consortium Working Draft WD-CCPP-ra-20000721, available at <http://www.w3.org/TR/2000/WD-CCPP-ra-20000721/>
- [Ohto99] 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>.
- [Parlay01] The Parlay Group, Homepage at <http://www.parlay.org>.
- [UMTS97] UMTS Forum Report 1: "A Regulatory Framework for UMTS", June 1997.