Middleware platform support for the realisation of advanced business models in beyond 3G environments

Nikos Houssos (1), Vangelis Gazis (2), Maria Koutsopoulou (3), Nancy Alonistioti (4)
(1,2,3,4) Communication Networks Laboratory, Department of Informatics & Telecommunications, University of Athens, 157 84, Athens, Greece
E-mail: (1) nhoussos@di.uoa.gr, (2) gazis@di.uoa.gr, (3) mkoutsop@di.uoa.gr, (4) nancy@di.uoa.gr

Abstract

The field of telecommunication value-added services provision is undergoing significant changes during the last years. The mobile services value chain is already impacted and more dramatic developments are expected in view of the 4G mobile systems, whose advent is expected in the medium to long term. This contribution presents a non-network operator centric business model for mobile service provision, as well as a middleware platform that can provide the technological means for its application in practice.

1 Introduction

Service provision in telecommunication networks (e.g., fixed telephony networks, 2G mobile systems) has traditionally been based on network operator-centric paradigms (Lazar, 1997). Important negative effects of this approach include tedious and time-consuming (in the order of months or years) processes for the introduction of a new service as well as a limited range of available services (a significant part of these services concerns supplements to the basic voice service like freephone, call diversion, etc.), which are available only to the subscribers of a single mobile network operator.
Efforts for enabling the provision to end-customers of value-added services that are more user-centric and content-oriented have been under way since late 1990s. The most successful among those efforts have led to the significant popularity of mobile services in the Japanese market (i-mode, EZ-Web, J-Sky) and recently (mid-2002 to early-2003) to similar activities elsewhere (e.g., i-mode Europe, Vodafone live!). Despite their limited innovation with respect to technological issues, these efforts are an important step forward, since third parties have the chance to enter the market and widen the range of available content/applications, while service time-to-market has decreased.
All the above models are dominated by the network operator, a fact that is reasonable given the vertically integrated network infrastructures of today. However, this status could be challenged by efforts for realizing the vision of environments commonly termed as “beyond 3G” or “4G” (Pereira, 1999), where user-centric, ubiquitous services, supported by computing capabilities embedded in all kinds of tangible objects and provided over heterogeneous, reconfigurable network infrastructures, will significantly impact almost every aspect of peoples’ daily activities. In such a case, the requirements for seamless and optimized service delivery over heterogeneous access technologies (e.g., cellular, WLAN, DVB), make current approaches not adequate and create the need for novel business models where the network provider does not hold the leading role (Houssos et al, 2002) (UMTS Forum 2002) (Mylonopoulos et al, 2003). Although recognizing that any path towards beyond 3G systems should be evolutionary, this paper describes a model that can support the realization of the 4G vision, based on a business actor in a mediating role that provides the “glue” between network operators and application developers/providers and facilitates their efficient co-operation. Moreover, this contribution aims to present insights on how this paradigm can be supported by a distributed middleware platform.
The rest of this document is structured as follows: Section 2 describes our proposed scheme, including the business model and middleware platform aspects. Section 3 identifies and presents key technical issues and solutions that enable our overall approach, while Section 4 evaluates the proposed paradigm and discusses its feasibility from a business point of view.

2 A scheme for beyond 3G service provision

2.1 Business and service provision model

Figure 1 presents an illustration of the service provision model proposed and the supporting infrastructure. The key innovation is the introduction of a mediating platform that enables service delivery through the seamless collaboration of different market players. From a business perspective, the main actors in such an environment will be the following:

- *Mobile User*: The mobile user is the actual consumer of services.
- *Mobile operator*: This is the entity that operates the network infrastructure for mobile user access to services and will typically also provide third-party access to its network through a standardized API (e.g., analogous to OSA/Parlay (Parlay Group, 2003) (3GPP, 2003) or JAIN SPA (Keijzer et al., 2000)). Networks of several types (e.g., 3G cellular, WLAN, DVB) can be “plugged-in” the service delivery architecture. Naturally, a single operator may administer different types of networks.
- *Value-Added Service Provider (VASP)*: This is the provider (and typically also the developer) of the end-user application. The application may comprise components that reside in user equipment (e.g., mobile terminal) as well as in enterprise servers.
- *Platform operator/provider*: This is the entity that owns and administers the middleware platform for service provision. The platform provider establishes business-level agreements with all the other main types of actors: mobile

users, network operators and VASPs. An important innovation concerns the fact that the platform operator maintains the customer relationship with the end-user, undertakes the generation of a single user bill for connectivity as well as service consumption and handles the sharing of the resulting revenue among network and service providers.

The service provision model supported by the platform is roughly the following: Before a service becomes available to end-users, it is automatically deployed based on VASP-provided detailed service metadata. This procedure may include sophisticated actions like reconfiguration of the underlying networking infrastructure. From then on, the user is able to dynamically discover, download and execute the service, without the need for a prior subscription with the VASP. Eventually, the charges for service consumption are included in a single user bill and the corresponding revenue is shared among the platform provider, VASP and network operator.

2.2 Platform functionality

The utility of the platform to the various mobile service provision stakeholders is illustrated in the use cases of Figures 2, 3 and 4 and elaborated in the following.

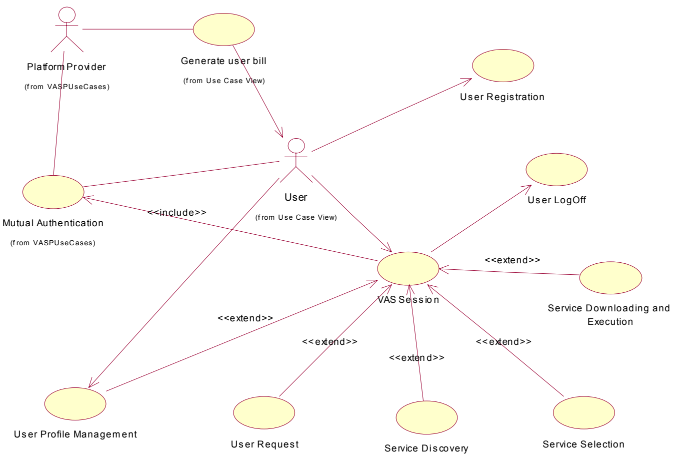

Figure 1. Mobile user use cases.

Figure 1 depicts UML diagrams of use cases pertaining to the mobile user. Before each other interaction with the platform the user should register with it (*User Registration*). From then onwards he/she is offered the possibility of establishing with the platform *sessions* (*RCSPP Session*), in the frame of which he/she can have access to value-added services. Each session begins essentially with the mutual authentication (*Mutual Authentication*) between user and platform and expires with the explicit or implicit (e.g., based on idle time) user log-off (*User Logoff*). A session consists typically of service discovery (*Service Discovery*), selection (*Service Selection*), adaptation and packaging (*Service Adaptation and Packaging*), downloading and execution (*Service Downloading and Execution*).

The above-mentioned operations can be repeated any number of times within the frame of a single session and are highly customized to situational parameters. For example, the service listings presented to the user during service discovery include only the applications that can be delivered given the current context; services are excluded if they require features not presently available, like a specific operating system, java execution environment or network type. The set of applicable contextual data contains, at least, terminal capabilities, network characteristics and user preferences and status.

An innovative user-related use case that deserves attention is the adaptation and packaging (Houssos et al., 2003). Essentially, it refers to the bundling all the software components and other supporting resources (e.g., images, etc.) required for executing a service in a single archive and provide it for download to the mobile client. The generation of the single archive is based on identifying the service components that should be included in this particular application configuration and possibly also selecting the more appropriate implementation of a component (if more than one are available). The above crucial decisions are reached by examining the context of the particular service selection request and adapting the service configuration accordingly.

Within a session, the user is also able to directly manage his/her personal profile (*User Profile Management*), which is maintained by the platform. Last but not least, the platform provider issues for each user an itemized bill (*Generate User Bill*), reflecting the charges resulting from consumption of network resources as well as value-added services.

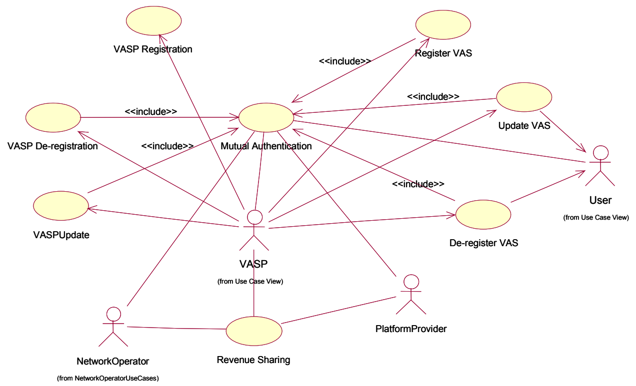

Figure 2. VASP use cases.

VASP use cases are shown in Figure 2. The platform keeps track of VASP data, such as identification and security information and tariffing and revenue sharing policies. VASP data is communicated to the platform during the *VASP Registration* operation and can be updated (*VASP De-registration*) and removed (*VASP Update*) at any time. At least a part of the aforementioned operations shall be performed off-line. This could include, for instance, negotiations about the terms and conditions (e.g., general revenue sharing schemes and rules) of the VASP-platform provider business relationship.

An important platform feature is the automatic deployment of applications over a variety of networks of different types (*VAS Registration*). The corresponding operation is accomplished as follows: the VASP formulates service metadata in a formal way (e.g., according to an XML DTD or Schema) and provides it to the platform, which is intelligent enough to interpret it, decide on the actions that need to take place and trigger their execution.

A service profile may contain a wide variety of data, such as the following:

1. General information about the service (e.g., description, category, keywords, VASP information).
2. Targeted networks for deployment (if these are specific, otherwise the desired fine-grained properties should be specified, see 5).

3. Data describing the service software architecture. The latter follows a high-level model (Houssos et al., 2003), which is independent of implementation details and therefore does not place heavy constraints on the service developer. A modular and component-based structure of service software is highly desirable, since it enables the easier identification of potential deployment alternatives and the corresponding customisation of the application on a per request basis. The service software architecture description includes information concerning the physical distribution and location of service components, which is required, among others, for billing purposes (e.g., monitoring of corresponding network traffic flows).
4. Requirements from user terminals or in general customer premises environments (e.g., HANs, PANs).
5. Requirements from network infrastructure, which can be distinguished in pure technical requirements and business/technical requirements. The former concern only the technical characteristics of a network, such as its type (e.g., GSM/GPRS, UMTS, WLAN), available bandwidth and QoS support. The latter contains data related to the administrative entity of the network and its revenue sharing policies. Most importantly, however, it includes information on application requirements for access to network functionality through standardised open interfaces, such as Parlay/OSA and JAIN. Important elements of this information are the specific standardized APIs (e.g., call control, user location, charging, presence) that are used by the application, as well as the access conditions of use that are required (e.g., available number of method calls per second).
6. VAS-specific information regarding tariffing/billing and revenue sharing between VASP and platform provider.
7. Support for user preferences (e.g., language, pricing level) and status (e.g., geographical location).
8. Security data.

A VASP can, at any time, stop offering a particular application through the platform (*VAS De-registration*).
Furthermore, it is reasonable to expect that some attributes of the service may change over time. For example, new versions may become available, the physical allocation of application components may be modified or new tariffing schemes and price levels may be introduced. Such modification can be remotely triggered by the VASP through the *VAS Update* use case.
Note that the implementation of this operations requires sophisticated logic, since certain types of changes can affect the overall service provision process in various ways, which need to be identified. For instance, when a VASP server moves to another network address, appropriate reconfiguration of the network metering and billing system may need to take place so that application usage traffic is correctly monitored. Moreover, an increase in the price of an application may lead a user that is currently executing it to immediately stop using it. All these actions are enabled and co-ordinated by the platform.
A crucial use case relates to the sharing of the income of application consumption by end-users among the platform provider, network operator and VASP. This function is accomplished by the platform (*Revenue Sharing*), which maintains a kind of clearinghouse for all service provision stakeholders.

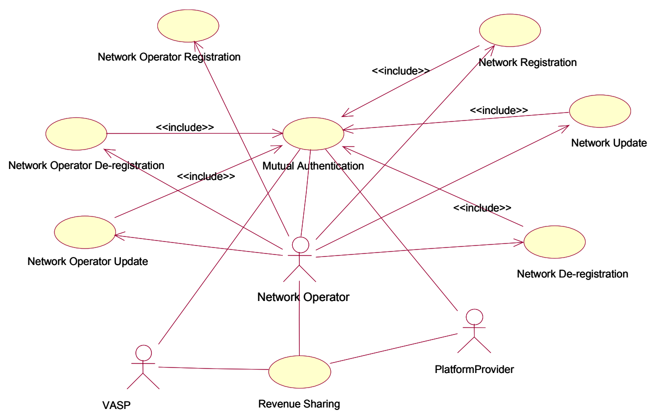

Figure 3. Network operator use cases.

Platform use cases for the network operator are depicted in Figure 3. In the same way as VASPs, network operators need to be registered with the platform. The corresponding data management operations are *Network Operator Registration*, *Network Operator De-registration* and *Network Operator Update*. They involve exchange of data required for security and revenue sharing functions and they are carried out partially off-line.
Network operators can register their network(s) (*Network Registration*) with the platform and thus enable the end-users that are connected to their infrastructure to access the whole range of value-added services that are offered by platform-affiliated VASPs. In analogy with VAS Registration, this operation involves the formal specification, by the network operator, of network metadata. The latter is subsequently interpreted by the platform, which identifies and triggers appropriate actions (e.g., transfer of an application session over the new network, change of price levels for the consumption of a service).
The network profile consists of disparate data, which may include:

1. Network identification and general description information, like type (e.g., GSM/GPRS, UMTS, WLAN), physical location and network operator data.
2. Technical characteristics, like type of bearer, supported bandwidths, topological and coverage information, QoS levels supported.
3. Characteristics of hybrid business/technical nature, such as revenue sharing policies. Moreover, Information of major significance concerns support of open interfaces (e.g., Parlay/OSA, JAIN) that enable access to selected network functionality. The aforementioned data comprises a bulk of parameters related to which specific APIs are supported and under what terms (e.g., third-party entities that may access them or more fine-grained conditions, like maximum allowed frequency of invocations per method). The availability of open interfaces and the relevant terms of use can be an important differentiator for network operators, since it facilitates advanced, flexible service provision through context collection and network reconfiguration capabilities.

The platform offers also operations for updating network information (*Network Update*) as well as for removing networks from the system (*Network De-registration*).
Note that, as mentioned previously, the *Revenue Sharing* is another use case that involves the network operator.

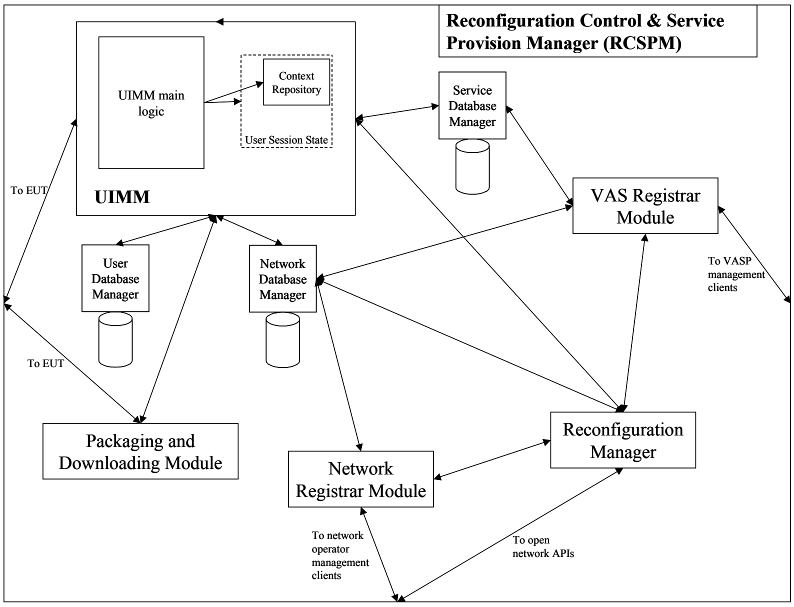

Figure 4. RCSPM internal architecture and interfaces.

2.3 Platform architecture

The proposed platform comprises a number of middleware components that are physically distributed in various administrative domains. They are briefly described below:

The *Reconfiguration Control and Service Provision Manager (RCSPM)* is the central platform component in that it co-ordinates the entire service provision and management process. The RCSPM is described in more detail later in this section.

The *Charging, Accounting and Billing (CAB)* system is responsible for producing a single user bill for service access and apportioning the resulting revenue between the involved business players.

The *End User Terminal Platform (EUT)* resides in the user terminal and includes functionality such as service downloading management, GUI clients for service discovery and selection, capturing of event notifications as well as service execution management. The EUT contains a context repository, which is a subset of the primary copy residing within the RCSPM.

The *VASP* component of the platform is located in the VASP domain (e.g., enterprise servers) and provides graphical clients (e.g., web-based) through which VASPs are able to remotely trigger management operations (e.g., VAS registration).

Similarly, the *NetworkOperator* component of the platform is located in the network operator domain and is used as a user interface from which the network provider-related operations are initiated.

The most important component of the platform is the RCSPM, whose internal architecture is depicted in Figure 4. The RCSPM comprises the following modules:

The **User Interaction Management Module (UIMM)** is responsible for providing the user with a highly personalizable, context-aware mobile portal. It manages user sessions with the RCSPP and co-ordinates user-related operations like service discovery, selection, adaptation and downloading as well as user profile management (user data is persistently stored in an appropriate database). The UIMM is the place where dynamically updatable user session information is maintained. This information is accessible to clients (e.g., applications) through generic open APIs.

The **VAS Registrar Module (VASREGM)** is responsible for interacting with 3rd party service providers. Through the VASREGM the platform operator provides VASPs with a way to automatically deploy their services. The VASP compiles a profile, encoded in XML, of service attributes. Based on these attributes, the VASREGM co-ordinates service deployment, including various actions like reconfiguration of the underlying infrastructure and uploading of service components to the RCSPM. The service provider is able to manage (add/delete/update) its services via a convenient web interface.

The **Network Registrar Module (NetREGM)** enables network operators to register their network with the platform. The network providers feed into the NetREGM a descriptor of their infrastructure, according to a common, extensible network information model. Based on this descriptor attributes, the NetREGM triggers all the required (automatically carried out) actions like storage of appropriate data in network profile database as well as instant deployment of the available VAS that are compatible with the newly registered network(s). In the same way as the VASREGM, NetREGM functionality is remotely accessible to network operators via a convenient web interface.

The **Reconfiguration Manager (RCM)** undertakes network, platform and service reconfigurability. The RCM is responsible for executing the appropriate reconfiguration actions on the underlying network during VAS management procedures (registration/de-registration/update), triggered by the VASP. The RCM also comprises a generic adaptation module (Houssos et al., August 2003) that is used for supporting adaptation through functions like intelligent profile matching.

The **Packaging and Downloading Module (PDM)** is responsible for dynamically creating, in a context-aware manner, a single downloadable application bundle.

The RCSPM also includes database managers that provide interoperable access to the persistent service, user and network profile repositories hosted by the platform. In particular, the maintenance and management of network profile information is vital to situation-awareness support, since it enables the instant availability of data regarding the constantly changing infrastructure that is close to the user as he/she moves within beyond 3G, ubiquitous/ambient computing environments.

3 Key technical aspects

The implementation of the platform functionality described above is far from trivial; the platform architect has to face various design alternatives and crucial technical issues. In this section, we present key design choices and technical enablers that are crucial for the development of the platform.

General features that are critical in implementing various parts of the system are the following:

- Employment of open interfaces for the communication between other actors and with the underlying infrastructure. This characteristic is required so that portability of services and system functions, as well as easy interoperability/interworking is achieved.
- Generality in terms of profiles and algorithms supported. The platform is required to interpret and process a vast variety of data, according to a diversity of algorithms, suitable for different situations. Thus, the corresponding profiling and algorithm loading tasks should be implemented in a run-time extensible fashion, so that the platform can accommodate a wide range of not a priori known environments.

Moreover, with respect to particular platform functionality, important aspects are discussed in the following paragraphs.

3.1 Context management

A context management system designer faces a variety of technical challenges. The most important of them are:

- Identification of the type of entity (e.g., person, application, terminal?) that context information is bound to.
- The location and distribution of context data, including synchronization techniques in case replication is employed.
- The retrieval of context information from its original sources.
- The representation of context information.

The proposed platform includes mechanisms for resolving these issues (Houssos et al., September 2003). Key aspects of our solution are the following:

- Context information is related not with the user, an application or their combination, but with a *user session* with the platform.
- A single complete and authoritative copy of the environmental information concerning each user session is maintained; this copy is kept throughout the session in the internal of the RCSPM and in particular within the UIMM. The context information maintained by the UIMM is accessible to platform components and external services through generic, open APIs.
- The UIMM is responsible for retrieving context from all possible sources. The latter mainly include persistent repositories (like the databases maintained within the RCSPM), as well as (registered with the platform) network infrastructures accessible through open interfaces (e.g., OSA/Parlay, JAIN SPA). The UIMM retrieves data from the latter not only in a "pull" approach, but also is able to receive event notifications form the corresponding open API gateways. The contact information of the latter (e.g., network addresses, access terms of use) is included in the network profile information provided by the network operator at the time of registration of the corresponding network with the platform.
- Context parameters are represented with a scheme that ensures generality (Houssos et al., September 2003). Furthermore, they are dynamically "annotated" with certain meta-attributes, which can be very useful for ensuring that the environmental data will be used in an appropriate way. Examples include timestamps and security/privacy settings.

3.2 Reconfigurability/Adaptability management

Reconfigurability/adaptability management is based on a variety of information, including service, network, terminal profiles as well as meta-data about the various actors (network operators, user, VASPs) and their policies. The design is such that high-level events (input from a user or a modification of an actor's policy) can trigger the appropriate actions to a service or down to the level of the network element.

Certain tasks regarding service adaptability and network reconfiguration are performed *transparently* to the service. Thus, the service logic is totally agnostic of them. This enables the support of applications (e.g., legacy services) whose developer were not aware of the existence of the platform or of open interfaces to underlying networks.

A crucial feature, as also mentioned previously in this section, is support by the system of arbitrary profile and algorithm data that is not known at development/compile time. This feature can be implemented through the application of principles of generic programming and use of advanced facilities of modern object-oriented languages (e.g., reflection, dynamic loading, run-time type identification in Java) (Houssos et al., August 2003).

3.3 Charging/Accounting/Billing

A critical part of the platform is the Charging/Accounting/Billing (CAB) function, which is implemented as a separate service (Koutsopoulou & Alonistioti, 2003), not coupled with the network infrastructure of a particular connectivity provider. Thus, it can be administered by a trusted third-party, like the platform operator.

CAB functions are based on intelligent collection and correlation of traffic monitoring data originating from diverse types of networks, thus allowing the use of heterogeneous infrastructures for service delivery. Network monitoring can be tailored to CAB needs, so that the required information is collected without affecting the performance of the network operation. Moreover, billing and pricing policies are dynamically reconfigurable, in a secure manner, by authorised stakeholders. All the above-mentioned tasks are accomplished through open interfaces that may be implemented as extensions of the Parlay/OSA interfaces.

The overall billing function is layered; a separation exists between transport, service and content charges. Different techniques can be applied for each layer, with the implementation of content-based charging being probably the most challenging technical issue (Koutsopoulou & Alonistioti, 2003) (Volubill, 2003).

4 Evaluation and feasibility of business model

An important business consequence of our approach is that the service provision process is co-ordinated by an entity other than the mobile operator. This paradigm enables third parties to become specialized in the service platform operator role, partner will various network operators and thus offer VASPs fast service deployment over more than one networks through only a single business-level agreement. In an environment where a wide range of end-user applications, developed by a multitude of VASPs, are provided over multiple types of networks our approach provides significant gains in terms of efficiency within the mobile service provision value chain. In particular:

- One-step deployment of a single service over multiple networks (and even different network types, e.g., UMTS and WLAN) substantially reduces deployment costs for VASPs. Moreover, the delegation of interaction with the mobile operator for access to open network APIs also significantly contributes to VASP overhead minimization.
- Dynamic delivery to the mobile user of a richer repertoire of applications becomes easier, since the available services can be discovered, downloaded and accessed in a unified and customisable way through a portal-like interface.
- After a single registration operation, network operators can have a bulk of services delivered over their infrastructure, generating for them substantial revenue from transport charges.
- Automation of service installation and deployment drastically reduces service-to-market costs for network operators and VASPs. The ability for easier and faster deployment leads to availability of a larger number of services and thus to more traffic and income for operators. Moreover, the reduced time between the development and the commercial exploitation of a service leads to faster return of investment for VASPs.
- The automatic service adaptation to context parameters (e.g., terminal and network capabilities, user location and preferences) off-loads from the VASPs the cost of implementing the required adaptation decision logic on their own.

An important issue, however, is whether a dynamic marketplace as described above in this section, can be created, given that until today (and most probably for several more years to come) mobile service provision has been dominated by network operators. We believe that this scenario has realistic chances to come true, in the case that:

- Different types of economically viable access technologies proliferate in the near future.
- Service and content developers worldwide manage to create a large number of applications that are useful, affordable for end-users and portable to a variety of network and terminal infrastructures.

There are already signs that both these prerequisites will be fulfilled in the coming years, like the drastic increase of 802.11 WLAN hotspots worldwide and the wide adoption by VASPs of portability enabling industry standards like 3GPP MExE, OMA UAProf and W3C CC/PP. This is also recognised by international fora and research communities (UMTS Forum 2002) (Mylonopoulos et al, 2003). However, only time will tell whether these factors will become strong enough within a reasonable timeframe (e.g., in the range of 5 years).

Another question relates to what kind of business players can take on the service platform operator task. Such an undertaking requires significant effort, for gaining customer recognition (from both end-users and VASPs) as well as establishing non-exclusive business relationships with multiple operators. Entities that have the potential to overcome these difficulties and successfully carry out the platform operator function are enterprises that have an already established customer relationship and have built a strong brand (e.g., ISPs and wired Internet portal providers). Moreover, companies that due to their main function possess a direct channel to the user (e.g., media companies like broadcasters and newspaper publishers) may have the best potential in this role, since they can more easily promote their offering (at least to end-users).

5 Acknowledgements

The concept of a mediating service provision platform, although in a significantly more limited scope than presented in this paper, has been developed and validated through extensive prototyping and trials during project IST-10206-MOBIVAS (mobivas.cnl.di.uoa.gr). Part of the work included in this paper has been performed in the framework of the project “ANWIRE” (www.anwire.org), which is funded by the European Community under the contract IST-2001-38835.

References

Lazar A. (1997),, *Programming Telecommunication Networks*, IEEE Network Magazine, October 1997.

Pereira J. (1999),, *Beyond third generation*, WPMC 1999, Amsterdam, The Netherlands.

Houssos N., Gazis V., Alonistioti A. (2002),, A flexible management architecture for the support of advanced business models in 3G mobile service provision", 1st International Conference on Mobile Business, Athens, Greece, 8-9 July 2002.

UMTS Forum, Report 21, *Charging, Billing and Payment views on 3G Business Models.*

Mylonopoulos N. et al., *Emerging Market Dynamics in the Mobile Services Industry*, IST-MOBICOM project white paper, available from www.mobiforum.org.

Houssos N., Gazis V., Alonistioti A. (2002), *Application-Transparent Adaptation in Wireless Systems Beyond 3G*, 2nd International Conference on Mobile Business, Vienna, Austria, 23-24 June 2003.

The Parlay Group home page, http://www.parlay.org.

3GPP TS 23.127: *Service aspects; The Virtual Home Environment*
.

Keijzer J., Tait D., Goedman R. (2000), *JAIN: A new approach to services in communication networks*, IEEE Communications Magazine, January 2000.

N. Houssos, et al. (2003), “Advanced adaptability and profile management framework for the support of flexible mobile service provision”, *IEEE Wireless Communication Magazine*, August 2003.

Koutsopoulou M. Alonistioti N. (2003), *Framework for Charging and Billing for Reconfigurable Services*, in Dillinger M. et al. (Eds.), *Software Radio: Architectures, Systems and Functions*, John Wiley & Sons.

Volubill Dialog Control and Charging Platform (DSCP), URL http://www.volubill.com, June 2003.

Houssos N., Alonistioti N., Merakos L. (2003), *Towards efficient support of context-awareness in mobile systems*, PIMRC 2003, Bejing, China, September 2003.