Service Program Mobility: A New Paradigm for Mobile Operators' Service Delivery

Henrik Lundqvist[1], Zoran Despotovic[1], Dan Jurca[1], Wolfgang Kellerer
DOCOMO Communications Laboratories Europe
Email: lastname@docomolab-euro.com

***Abstract*—The delivery of emerging, high performance mobile services is traditionally limited to the cellular operator home domain. Due to the centralized nature of service platforms, resource-demanding services suffer from long delays when a user is accessing them when roaming. With the globalization of our society, such data intensive services are enabled at the high cost required by provisioning high-speed inter-domain connections. Delivering those services from the local, visited domain (instead of the home domain), does not only overcome the aforementioned limitations, but also opens new possibilities for service delivery. An operator could expand its service delivery reach to a global coverage, enjoying all the benefits of such a move. We propose Service Program Mobility as a novel paradigm for global service provisioning. Our system architecture realizes dynamic service component migration across operator domains and platforms to offer the same experience to roaming users as they would have in the home network.**

***Index Terms*—Mobile Services, Service Mobility, Service Delivery Platform, Service Roaming, Service Migration, Mobile CDN, In-Network Services, Network Virtualization.**

I. Introduction

Imagine cellular network services that are offered globally and seamlessly without the traditional limitation of having to access the home network no matter where the service requests originate. Then the same user experience can be provided to roaming users as to the home users. This allows a cellular operator to reduce deployment and operation cost while extending the service reach. Or, imagine a scenario in which services are distributed within the home network as to follow the users, and then provided from points in the network that are closest to the requesting users as opposed to central servers. This should lead not only to better user experience but also to more efficient usage of network resources.

We refer to a set of technologies to realize this paradigm as Service Program Mobility (SPM). The vision of SPM is to develop an in-network capability supporting the dynamic placement of service components across domains to improve service performance for the user (Quality of Experience) and reduce infrastructure cost. Roaming is extended from a conventional access roaming paradigm to the roaming of the whole or parts of a service (data, logic and state). As such, SPM will introduce an important paradigm shift in the way how the global mobile operator business (i.e., roaming) is structured. For mobile operators, SPM is an opportunity to provide a richer service offering instead of being just a bit pipe. By integrating the service platform with the network, the service performance will be improved, bringing an advantage for both, users and service providers. In addition, the mobile operator improves the usage of its network resources.

SPM is a part of a wider initiative undertaken by NTT DOCOMO, named Next Mobile Network (NMN), with the goal to define the mobile network of the future. [1] provides an overview of NMN. Although the next mobile network initiative plans its deployment within ten years from now, SPM as an individual project should be deployed within at most four to five years. We also show in this paper how network virtualization technologies, wich are part of the NMN concept, allow a smooth deployment of SPM.

Thus we expect SPM to be part of the service delivery platform of a next generation IP-based mobile network evolving today's centrally organized service delivery platform in 3G mobile networks. Emerging network control and management concepts allowing more flexibility and cost efficiency in the network configuration and operation (such as network and server virtualization) constitute enabling technologies for service program mobility. However, SPM can also be partly realized with enhancements of today's service delivery platforms. This can be achieved by gradually distributing the service platform to multiple datacenters and allowing flexible placement of services using methods similar to Content Delivery Networks (CDNs) as a first step. Services may also be placed in a computing cloud geographically close to a visited network that does not support SPM.

Manuscript received February 15, 2011; revised May 15, 2011; accepted June 15, 2011.

[1] Henrik Lundqvist, Zoran Despotovic and Dan Jurca are now with Huawei Technologies, henrik.lundqvist@huawei.com, zoran.despotovic@huawei.com and dan.jurca@huawei.com.

doi:10.4304/jcm.7.2.106-114

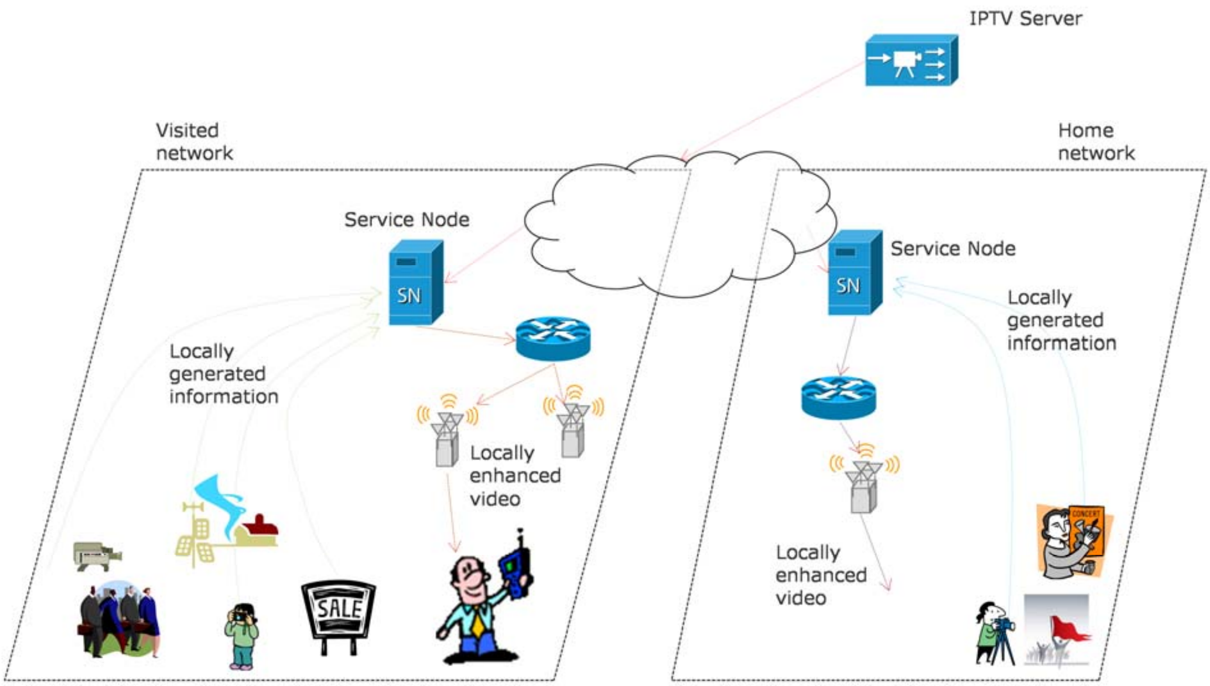

Figure 1. Video with local information enhancement scenario.

We argue that SPM is a key technology that will help the next mobile network satisfy the ever increasingly demanding requirements from its customers, while keeping the operation costs of the overall network low.

II. Benefits and Use Cases

The main advantages of the technology are briefly discussed below:

Improving user satisfaction: By moving services along with the requesting users, we expect to obtain an improvement in perceived quality metrics at the user side. This improvement is mostly noticeable in roaming scenarios, where currently users traveling far away from the home network, experience a high delay in accessing home services. Bringing the service closer to the user in the visited network will increase the performance of the system. For example the service access delay incurred by traversing the visited and home networks, connected by a trans-continental link, can be avoided. A home service relocated in the visited network will allow for local breakout and service access with minimal delay. In some severe cases, such as roaming from Japan in Europe, this can result in a reduction of delay of around 400ms. For many interactive applications this delay is critical.

Saving network resources and cost: Whether in a home network, or roaming scenario, SPM works on the assumption that services can be moved closer to the clients' geographical position. This implies that geographically distributed servers are available to host the services. Moving services from one server to another will naturally incur a migration cost. However, it is expected that such cost is compensated for by the reduction of network resources used, (e.g. bandwidth in the core network, or on the link connecting two mobile networks). Overall, an optimal dynamic placement of services closer to the network access points should greatly reduce the network resource consumption, leading to a better load distribution in the network and a more efficient utilization of the operator's network resources. In turn, this will lead to a lower infrastructure investment and operational cost for the network operator.

Supporting new services: SPM opens up the road to a new paradigm governing an operator's service creation and offering. New services and enhancements of current services will benefit from the deployed SPM technology, and increase the attractiveness of the operator for the mobile clients. One example is service composition across operator domains. For users who are roaming in a visited domain, home network services can be enhanced dynamically with local information. This allows a timely access to local information of current interest. Another example is the dynamic reconfiguration of services based on local access and of the transport network with respect to specific service components that are available at a given network location. Here the goal is not to add additional information, but to enhance the user experience by exploiting locally available service components and combine them with migrated service components, and increase his satisfaction.

Below, we introduce three use cases for Service Program Mobility.

Use case: Enabling new services

Imagine the following scenario (see Figure 1): a video is distributed from a central server to a number of mobile users. Further, the video is enriched with some information that is generated locally, close to the users receiving the video. Examples of such local information might be weather info, local advertisements, traffic jam alerts of even locally generated video such as web cam feeds from a surrounding area.

A straightforward solution to realize such a service is to send the local information to the central server, process it there, i.e. create a set of video streams customized for specific users or areas, and then send those back to their corresponding users.

A more efficient way to provide the described service is as follows: Consider the available locations in the network providing processing units capable of integrating the locally generated information with the centrally originated video and place the respective service components at those locations which are beneficial for execution with respect to some performance metrics. Compared to existing solutions for local information addition to videos, SPM would allow general processing components to be placed inside the network, which leads to a higher flexibility and easier deployment of new services.

Both flexible placement of service components at instantiation time and dynamic re-arrangement of service components onto fitting resources at runtime should allow achieving the best possible service quality under varying service conditions. This requires functionality in the network that derives appropriate processing resources and coordinates instantiation or migration at runtime based on specific information e.g., resources available at specific locations, service requirements and current network status.

Use case: Reducing delay

Imagine a Japanese mobile subscriber traveling to Paris and getting involved in a conversation in French. She is consulting her real-time speech translation service provided in the home network. However, due to the large distance between Paris and Tokyo, she is experiencing too long delay to make the service appropriate and useful. It would be much better if the service itself could follow the user and be provided from a place close to the user location, i.e., from a suitable point in the French operator network in our example.

Use case: Addressing flash crowds

Consider a roaming scenario in which a large number of users move to a hot-spot while roaming in a network abroad (e.g., gathering in a stadium to see a soccer world cup match). Access to home services may become a bottleneck in this case. As a countermeasure, the home network operator may decide to migrate service components to the visiting network, close to the place where the users reside.

III. System Requirements

In the following, we discuss the system requirements for the realization of our SPM concept. Two very basic requirements have to be fulfilled for the introduction of SPM referring to the benefits in Section II:

- Any service delivered to the user based on the SPM platform must fulfill at least one of the following two conditions: 1) user perceives better or equal quality of experience, 2) operator reduces its costs of delivering the service.
- It must be possible to introduce SPM through an incremental process, i.e. no abrupt change of the network operation is acceptable and it should work with all common mobile phones.

Besides these basic requirements, SPM should fulfill the following requirements:

- SPM shall, triggered by an event or by predictions of the service usage, select appropriate components to migrate and perform their migration to a new network position.
- SPM should be able to redirect client requests to new locations following service component migrations.
- SPM should be able to dynamically adapt to the state of the network and optimize resource usage by distributing service components based on the service usage.
- SPM should be transparent to the user, i.e., user should not be able to notice any negative change in the service operation.
- SPM should be secure to use, i.e., should not cause any security hazards to mobile subscribers or network operators.
- SPM should work in a multi-operator environment such that service components can move from a home network to a visited network.
- SPM should be easily tunable as to support various cost / performance trade-offs that the operator of the platform might find appropriate to implement.
- SPM should easily adopt any future changes in the network layer, i.e., should not be bound to any specific version of any specific network architecture.

IV. SPM Architecture

In this section we will first describe some general considerations and prerequisites for the architecture and then move to more specific inter-operator and intra-operator solutions.

General aspects and main architecture function blocks

SPM enabled network nodes (containing extra storage, processing power and the necessary functionality) spread in the operator’s network are a condition towards realizing the goal of transparent service migration. Such nodes can be enhanced core network nodes such as G-/S-GSN in 3G or P-/S-Gateways in LTE core networks, or

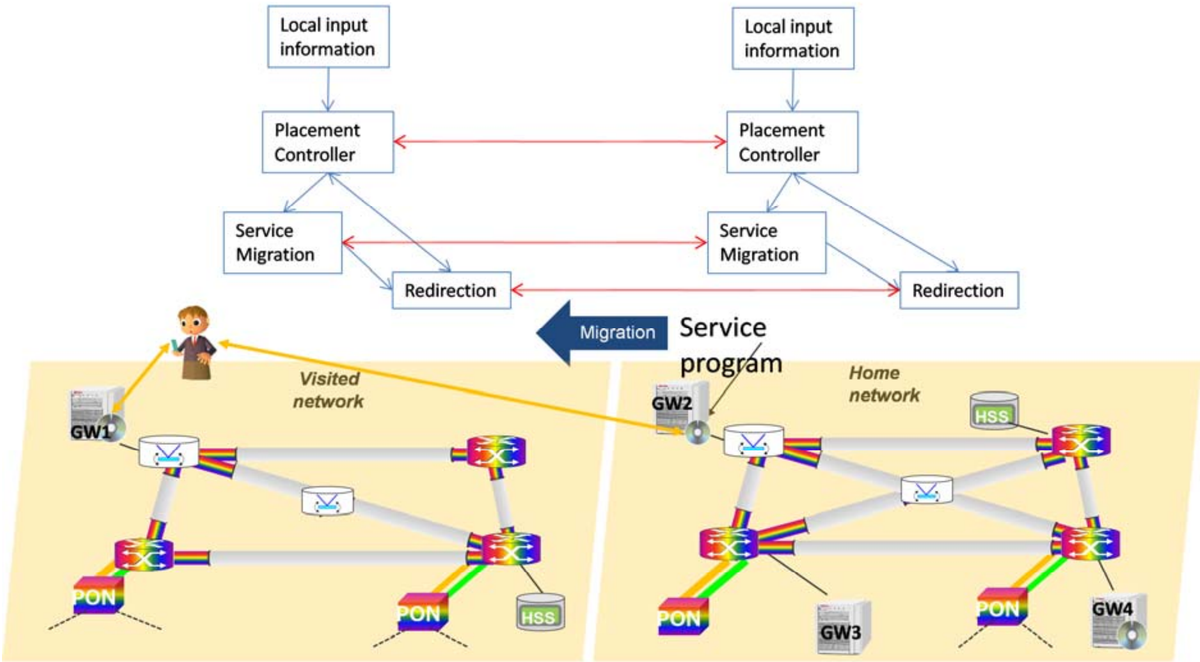

Figure 2. Inter-operator architecture and interfaces.

servers attached to core network nodes or enhanced service platform servers.

In addition to the SPM enabled nodes, the SPM architecture consists of functions to optimize the placement of the service programs. The main function blocks are illustrated in Figure 2. The optimization requires detailed information about the service requirements, the current service usage in the network, network load information and other types of information. As indicated in Figure 2 this information can be provided locally in each network. It is not necessary that every operator use the same placement policies or input information. Different trade-offs between performance and complexity may be preferred. Further function blocks realize mechanisms for migration of services and redirection of clients and service components.

Functions for static and dynamic migration

A first step to fulfill the requirements listed in Section III, is to place services when a new service request arrives. This includes the placement of the SPM enabled nodes in the network topology, the right assignment of a service to a SPM node based on the current network state and directing the client to the server. A service may be composed of several components, and some of the components may require a unique instance for each client, while others may be shared between multiple clients, for example a database. Clearly there are different trade-offs for these cases; for the shared components the usage and location of all clients should be taken into account. whereas for single user service instances it is sufficient to consider the location of the single client.

To realize the full potential of the concept, SPM supports dynamic migration of running services. To handle this, the SPM architecture consists of protocols for the transparent migration of live services. When users are consuming a service in real time, its migration becomes problematic due to the lack of availability of the given service during the migration time. Furthermore, service starvation during this transitory state has to be compensated. Mechanisms include pre-buffering functions at the service location, and possibly at the user equipment, bursting functions for fast service migration, functions for the synchronization of content and user access between the original and new service location, and finally functions for transparent, on the fly, user connection redirection to the new service location. Virtualization technology can be of used to mitigate or solve at least some of these issues. Placed in a virtual container, the service can be moved between network locations, while keeping its state saved during the transition(see Section IV for more details).

Inter-operator Aspects

To design a complete system architecture for SPM we need to consider the inter-operator collaboration when a user is roaming in a visited network. The interfaces between different operators need to be standardized, but the internal solution of each operator can differ. Some examples of protocols for the interfaces are mentioned in the following sections. To facilitate the analysis we may divide components into information providing, decision making and enforcement functions. The information providing functions gather data from an operator's infrastructure about running service processes, server load, network congestion etc, which an operator typically does not want to share with its competitors.

Therefore, it is preferable to avoid inter-provider interfaces between the information providing functions.

The decision about where to place a service component can be divided into two steps, first the home network operator decides whether the service should be run in the home network or by a different operator, then the operator that runs the service decides where in its infrastructure to run the service. For the first step it is necessary for the home network to compare the costs and benefits of running the service itself or utilizing services of a different operator. Although the optimization goals and algorithms of different operators may differ, it is necessary with an interface that provides information that can be used for this comparison. The interface should be designed to allow a large freedom for different optimization solutions.

The enforcement functions for migration of a service and for redirection of the clients and service components also require collaboration across the operator borders. Hence, there is a need for inter-provider interfaces also for these functions, but exactly what they should look like depends on the specific application, whether IMS is used and other considerations. The interfaces should be flexible enough to support the introduction of new services. However, there may be some common components, for example security aspects could be supported by the same protocols. In Figure 2 the inter-operator interfaces are shown in between the two domains.

Operator Internal Architecture

Figure 3 shows the system architecture of a possible implementation of the SPM functionality within one operator network. While Figure 2 shows the main SPM function blocks, this is a more detailed system architecture figure. The elements of this system architecture and their respective roles are described in detail in the next section.

Several of the functions can be based on elements already existing in the current networks, for example the accounting and billing function can be realized as an extension of current Accounting and Billing systems, the monitoring can use information already collected by existing network monitoring, and the client redirection could make use of Domain Name System (DNS).

Detailed description of components

This section presents in detail the functionality of each system architecture component (Figure 3), and the interaction and information exchange between the component functions.

The Monitoring Function (MF) and Prediction Function (PF) play an important role in SPM. The MF monitors relevant parameters related to the usage of any specific service. Examples include the number of users using the service at specific locations, traffic generated by the service usage and so on. The PF analyzes the data collected by the MF and discovers relevant trends present in the data. It then decides which data and in what form should be fed into the Placement Controller Function (PCF). The main goal is to enable pre-emptive placement of service components. For example, if we see a clear trend of a large number of users getting concentrated in a specific area then it makes sense to ship appropriate service components (close) to that area. The prediction can combine such instantaneous changes with knowledge of time-of-day patterns to predict the future usage. The MF sends information on relevant parameters to the other components either periodically or when a certain predefined event occurs.

The choice of which service component to move is made by a separate function, the Service Component Selection Function (SCSF). The decision is based on predictions on future service usage from the PF and the information about the current state of the network and the servers from the MF. The SCSF can, for example, choose to move or duplicate services where the quality requirements are under the risk of not being fulfilled based on the input information from MF and PF, or services that are not utilizing resources efficiently. The SCSF keeps information about the requirements each service puts on the execution environment, for example, in terms of hardware capabilities, and provides this information to the Topology Discovery Function (TDF). It also takes into account security aspects, for example restrictions on duplication and placement due to content or software ownership.

The Topology Discovery Function (TDF) discovers a set of candidate migration locations for a service component. The TDF contains a database of possible locations to be used for service placement, either in the operator’s network, or in the network of peering operators. Each location is characterized by its IP address and by a system configuration description and resource capabilities. The information in the database is to be matched against the service requirements received from the SCSF. The locations which fully satisfy the service requirements form a subset of possible migration locations, which is forwarded to the Placement Controller Function (PCF). The TDF ensures the hardware/software compatibility between the current location of the service and possible future locations. The information in the TDF should be updated whenever there is an update of the topology, for example introduction of a new server. More dynamic load information is maintained by the monitoring function.

The Accounting and Billing Function (ABF) is the function that has access to the migration pricing/cost information and SLA agreements between the operator and other partner operators (e.g., a database). This cost can either be a unitary cost (per event/amount of data/time) or a fixed pre-established cost (e.g., flat-rate resource rental between operators). Based on this information, the ABF is able to compute the price of migration of a given service to a given destination, either a location in the operator's network, or a location in a different operator's network. To this end, the ABF receives information from the PCF about the service requirements (size, necessary system configurations, transfer bandwidth, migration delay requirements) and the envisioned destination locations for the service. The computed cost for the possible locations is forwarded to the PCF. Before and after the service migration, the ABF collects usage information, and produces billing information taking into account user subscriptions and service placement.

The Placement Controller Function (PCF) calculates an optimal assignment of service components to processing locations. It is given as input an array of costs of running any component of any service at any specific location. It also takes as input the goal that needs to be satisfied such as minimization of traffic in the network, minimization of delay observed by the end users or any combination of these two, possibly discounted by the potential costs of the actual migration of the service components. It gives as output the new assignment of service components to processing nodes, or, if incremental operation is in place, a set of changes that need to be performed in order to switch to the desired state from the current one.

The Relocation Enforcing Function (REF) is the function that manages the service migration from its current location to the new location. It needs information of the service ID and service requirements, the list of users that currently access the service, the service state with respect to these users from the SCSF. It also needs information about the source and destination IP addresses of the current and future location of the service. Based on service requirements, the REF reserves the necessary resources at the destination entity. The REF then signals the source and destination entities in order to establish an end-to-end data path based on the service characteristics (e.g., migration bandwidth required, tolerated end-to-end delay for the migration, maximum loss probability, migration priority). The data path can be represented by a data tunnel (e.g., using VPN, or GTP protocols), with the possibility of integrity protection and ciphering the data transmitted (e.g., key agreement capability between the source and destination). When the service migration is complete, the REF function signals the destination entity the service parameters required in order to install the service/run the service at the expected configuration. Finally, once the service migration is complete, and the service is run at its new location in the desired state and configuration, the REF informs the Client Redirection Function (CRF) of the IP address of the new location of the service, the service identifier and the list of affected users which need redirection to the new service location.

Once a service has been relocated, the client connections are redirected from the old location to the new location by the Client Redirection Function (CRF). For this purpose the CRF uses information from the REF about the new location of the service for a user session, and implements some mechanism for redirecting the client. Several different methods can be considered depending on the specific service, for example using SIP REFER messages, HTTP redirection or DNS based redirection through modification of URLs. In the case of seamless migration of real-time services there is an additional need to synchronize the server and the client when the client is redirected. Typically, this may require application specific solutions, such as pre-loading additional data and state information to the client or the new server in anticipation of the redirection.

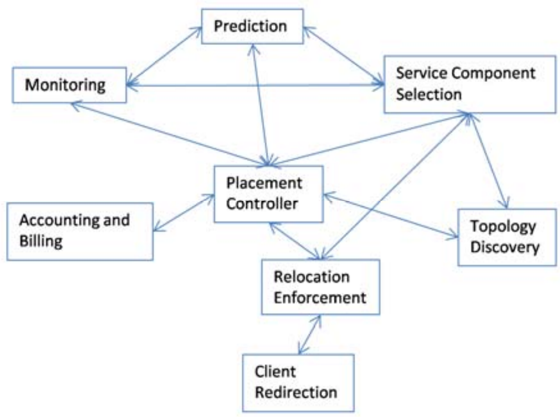

Figure 3: SPM system architecture

Figure 4 illustrates a typical procedure of migration of a component between two servers and the operation of the presented architecture, i.e., the message exchange between the involved components. The message flow is described as follows:

Let us take the second use case as described in Section II as a typical example. A Japanese mobile subscriber is requesting a speech translation service from his UE while roaming abroad. Initially this request reaches the original application server AS1 in the subscriber's home network (right side of Figure 4). The roaming situation is detected triggering the SPM functions through a *service migration request*. The Placement Controller decides the new location of the speech translation service program as AS 2 in the visited network and requests relocation from the Relocation Function, which requests a service transfer from AS 2 to AS 1. After the transfer of the service program is completed and the speech translation is ready to be used in the visited domain, the Relocation Function receives an *acknowledgement*. Then the Redirection Function is requested to redirect the service data flow

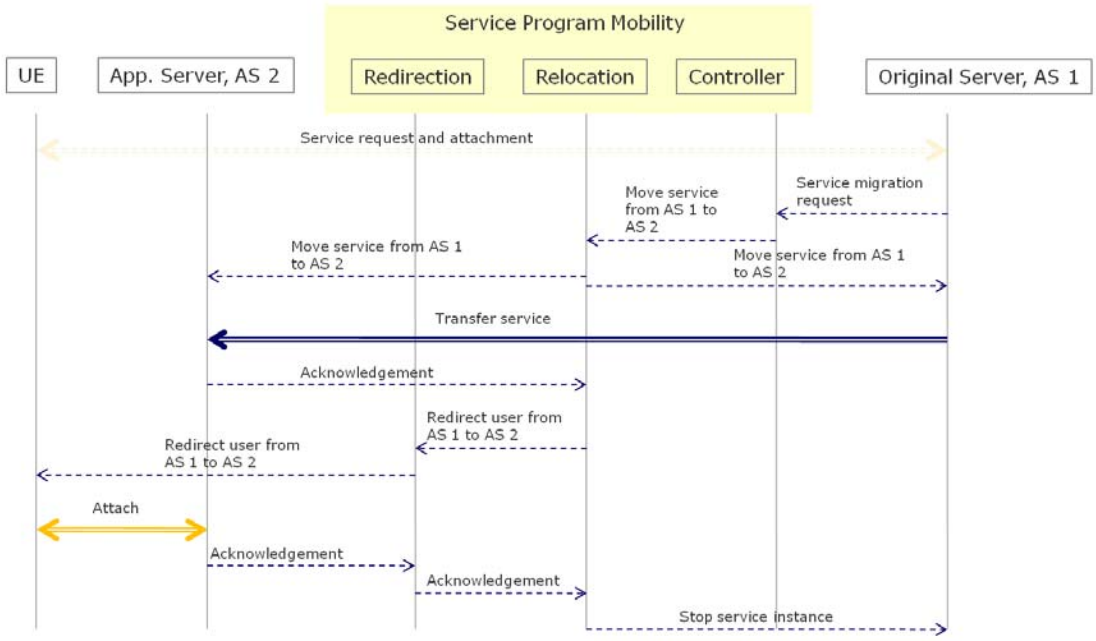

Figure 4. A Minimal SPM Procedure

from the UE to AS 2. This can be achieved through interfacing with the UE as shown in Figure 4, or through a DNS update in the visited domain. After the UE attaches to the closely located translation service program the Relocation Function has completed the service program migration successfully and stops the service instance on the AS 1.

V. SPM IN THE NEXT MOBILE NETWORK

NTT DOCOMO's vision of a Next Mobile Network [1] includes the introduction of a reconfigurable mobile network concept. In particular, this concept strives for an advanced control plane governing virtualized network resources. Virtualization technologies in service platform, core network and radio access network allow for a more flexible and dynamic allocation of network resources to different network services or coexisting virtual operators.

In view of such development SPM can be realized in a much more flexible way as operators can use such network virtualization mechanisms together with an advanced control plane in order to expand their (virtual) network and service coverage across their home network borders. With respect to SPM deployment this means that the inter-domain interfaces (shown in Figure 2) will not be needed any more as the operator can have the full control of its virtual network. The requirement of redirecting clients is shifting to a realization in the virtual network architecture designing how clients can access different virtual networks and services in one physical network domain.

Network virtualization plays an important role in the Next Mobile Network. The main drivers of network virtualization for network operators are sharing costs and investments for network equipment as the same physical resources (e.g., base stations, routers) can be used by multiple operators. This is followed by better utilization of network resources. A natural step following network virtualization is server virtualization. The two concepts give operators possibility to dynamically change the shape of their networks and tune their operation according to the policies they find appropriate. The concept of service

program mobility is a perfect fit for virtualized environments, networks and servers. Once virtualization is in place, it becomes easy to migrate services according to defined policies. Optimal placement of services, matching the demand and maximizing whatever goals the operators find critical (user perceived quality) become the main problem to solve.

VI. RELATED WORK

Moving services to a visited network to provide better services for a mobile user has previously been proposed in [2][3]. A set of requirements and functions to support service roaming are described. Moreover, business and legal aspects are described in [3]. The proposed architecture in these papers is built around a service roaming manager which handles the service roaming between home and visited network. The architecture proposed in this article generalizes this concept to handle both service roaming and service placement within the home network.

Also Content Delivery Networks (CDN) [4][5] are highly relevant to SPM. They distribute content among a set of locations at which a CDN provider operates, trying to reduce delay experienced by the user and bandwidth needed for distribution of the content.

CDNs follow two general approaches: overlay and network approach [6]. In the overlay approach, servers and caches create an overlay network, while core network components such as routers or switches play no role. In contrast, in the network approach the network components implement the cache and request forwarding features and thus realize the CDN functionality at the network layer. We believe that the latter approach is more appropriate for a network operator, for which accessing and programming network components is not an issue. Our ongoing project NetServ [7] makes a first step towards realizing the network approach to CDN by investigating how such functionalities can be effectively realized in the current network equipment. The techniques used by CDNs could be seen as a first step towards an SPM solution. However, SPM goes further in supporting service execution of complex services in the network and will be based on a tight integration into a mobile network.

Network support for media delivery in a mobile network has also been addressed in [8]. The approach taken is to introduce service-specific overlays that can provide network-based processing and adaption capabilities. Moving computing into the network instead of running programs at a client is commonly referred to as cloud computing. For mobile users the potential for energy saving has by cloud computing has been analyzed in [9]. This aspect is important because of the limited battery capacity of mobile terminals. The conclusion from [9] is that the energy saving potential depends on the application. In particular applications with high processing and limited communication demands can gain significantly from offloading the processes to the cloud.

Network virtualization technology can be used as a layer of indirection between the physical infrastructure and the network services deployed by the network operator. The technology is composed of node virtualization and network link virtualization technologies combined together. There are many methods available for link virtualization: lambda level L1 VPN, L2/L2.5 tunneling and L3 tunneling, e.g., [10]. All these levels can be holistically managed by GMPLS. OpenFlow and Flow Visors can also virtualize links/routers at flow level [11]. At the same time, well established node virtualization technologies can be deployed for gateway and server farm virtualization (e.g., Amazon's EC2 cloud).

As described in V., within the context of SPM, virtualization is considered as an enabler technology, which facilitates the migration of services and applications. Virtualization technology is supposed to enable large-scale global services, which users can consider as local to their home network at all times, independent of their location. A similar statement holds for the operator, it deploys global services based on the same principles as the local ones. Here we refer to the German National Project GLAB COMCON [12], for example. The main use-case of this research project addresses the problem of service component mobility with network virtualization technology. The project adheres to the principles and benefits of virtualization as presented above, in order to achieve the transparent migration of network services and applications, according to the user mobility patterns.

The Network of Information (NetInf) concept aims at establishing new network mechanisms which enable the information-centric networking paradigm, and take into account the many-fold expected growth in user generated content and network traffic. The same solutions may be possible to apply to the service program mobility issue through a generalized network caching technology. Ideas in this direction have recently been proposed under the name service-centric networking [13]. For example optimal caching placement and population solves a very similar problem to optimal processing placement and the required management functionality for enabling seamless service migration is also similar to that of information centric networks.

VII. Conclusion and Future Work

This article presents a novel paradigm for delivery of emerging, high performance mobile services based on the concept of migration of services close to the places where they are consumed.

The future work will focus on further defining the execution platform and, in particular, investigate how to provide its inherent SPM functionality. This requires overcoming a number of key challenges, the most urgent of which are:

Defining efficient service component mapping algorithms executed at service instantiation time (side condition: need for instantiation of millions of service components in parallel);

Defining efficient service component re-placement algorithms executed at service runtime;

Monitoring of network and service status to provide educated input for service component mapping and re-placement algorithms;

SPM requires a service component oriented service model to enable the movement of distinct service functionality to resources located at beneficial parts of the network. As a consequence there is need for efficient media forwarding between service components as media data will potentially be processed by a number of subsequent service components;

API definition to request service component instantiation and service composition.

References

[1] M. Yabusaki, Y. Okumura, M. Tamura: Next Mobile Network Architecture, International Conference on Communications (ICC 2010), Cape Town, South Africa, May 2010.

[2] J. Zoric et al., Service and information roaming – architecture and design aspects, WWRF17, November 2006.

[3] J. Simón, S. Thakolsri, S. Alapnes, J. Zoric, Technical and business aspects of service roaming, WWRF20, April 2008.
[4] Mukaddim Pathan and Rajkumar Buyya. A Taxonomy of CDNs. In Rajkumar Buyya, Mukaddim Pathan, and Athena Vakali, editors, Content Delivery Networks. Springer, 2008.
[5] George Pallis and Athena Vakali, Insight and perspectives for content delivery networks. Communications of the ACM , Volume 49, Issue 1, pp. 101-106, 2006.
[6] Lazar, I. and Terrill, W. Exploring content delivery networking. *IT Professional*, 3(4), pp. 47–49, 2001.
[7] NetServ Project Homepage, www.cs.columbia.edu/irt/project/netserv/
[8] F. Hartung et al. Advances in Network-Supported Media Delivery in Next-Generation Mobile Systems, IEEE Communications Magazine, August 2006.
[9] K. Kumar, Y.-H. Lu, Cloud Computing for Mobile Users: Can Offloading Computation Save Energy?, IEEE Computer, Mar. 2010.
[10] D. Benhaddou and W. Alanqar. Layer 1 Virtual Private Networks in Multi-domain Next-Generation Networks. *IEEE Communications Magazine*, pages 52–58, April 2007.
[11] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, J. Turner, OpenFlow: Enabling Innovation in Campus Networks. *ACM SIGCOMM Computer Communication Review*, volume 38, issue 2, pages 69-74, April 2008.
[12] G-LAB Project Homepage, http://www.german-lab.de/ → COMCON Project.
[13] T. Braun, V. Hilt, M. Hofmann, I. Rimac, M. Steiner, M. Varvello, Service-Centric Networking, 6th GI/ITG KuVS Workshop on "Future Internet“, Nov. 2010

Henrik Lundqvist received his Ph.D. in telecommunications engineering from KTH, Royal Institute of Technology, Sweden and the Master of Science in electrical engineering and applied physics from Linköping University, Sweden. He is a Researcher at DOCOMO Communications Laboratories Europe in Munich, Germany. He has previously worked as Research Scientist at NEC Laboratories Europe, held an ERCIM fellowship at NTNU, Norway and worked as Design Engineer at Nokia. His current research interests are service delivery in mobile networks, Internet architecture and resource control.

Zoran Despotovic is a Manager at NTT DOCOMO Euro-Labs. His research focus is mobile network architectures and distributed provisioning of services in mobile networks. His research interests also include distributed systems in the context to TCP/IP networks and peer-to-peer systems in particular. He holds a Dipl.-Ing. Univ. degree in Electrical Engineering from the Belgrade University (Serbia) and a PhD degree from the Swiss Federal Institute of Technology (École polytechnique fédérale de Lausanne, EPFL – Lausanne, Switzerland).

Dan Jurca received his M.S. degree from the "Politehnica" University of Timsoara, Romania in 2002. In 2007 he received his Ph.D. degree in electrical engineering from the Swiss Federal Insitute of Technology (EPFL), Lausanne, Switzerland, with his main topic of research focused on real-time adaptive multimedia streaming. Between 2007 and 2008 he worked as a post doctoral researcher in the Royal Institute of Technology (KTH), Stockholm, Sweden, on distributed network management aspects. He is currently working as a researcher for DOCOMO Communications Laboratories Europe GmbH, Munich, Germany, inside the Ubiquitous Networking Research Group. His research interests focus on reconfigurable mobile networks, multimedia streaming and QoE, and application and user mobility.

Wolfgang Kellerer is Director and head of the network research unit at NTT DOCOMO’s European research labs in Munich, Germany. Before he joined DOCOMO in 2002, he has worked at the Institute of Communication Networks at Munich University of Technology (TUM). In 2001 he has been a visiting researcher at the Information Systems Laboratory of Stanford University, California. His current research interests are next generation mobile networks including optical communication, network virtualization, multimedia transport, advanced mobility solutions, overlay networks and service platforms.