PERMIT: Network Slicing for Personalized 5G Mobile Telecommunications

Tarik Taleb, Badr Mada, Marius-Iulian Corici, Akihiro Nakao, and Hannu Flinck

The authors discuss the need for the deep customization of mobile networks at different granularity levels: per network, per application, per group of users, per individual users, and even per data of users. They also assess the potential of network slicing to provide the appropriate customization and highlight the technology challenges.

ABSTRACT

5G mobile systems are expected to meet different strict requirements beyond the traditional operator use cases. Effectively, to accommodate needs of new industry segments such as healthcare and manufacturing, 5G systems need to accommodate elasticity, flexibility, dynamicity, scalability, manageability, agility, and customization along with different levels of service delivery parameters according to the service requirements. This is currently possible only by running the networks on top of the same infrastructure, the technology called network function virtualization, through this sharing of the development and infrastructure costs between the different networks. In this article, we discuss the need for the deep customization of mobile networks at different granularity levels: per network, per application, per group of users, per individual users, and even per data of users. The article also assesses the potential of network slicing to provide the appropriate customization and highlights the technology challenges. Finally, a high-level architectural solution is proposed, addressing a massive multi-slice environment.

INTRODUCTION

Mobile networks are nowadays architected to serve all mobile users, ensuring some degree of service-level differentiation but with no specific tailoring of the functioning to the specific user needs. However, statistics demonstrate that users do not behave all in the same way: 53 percent are light mobile phone users, 24 percent exhibit medium usage behavior, and the remaining 23 percent are heavy mobile phone users [1]. Even among heavy mobile users, usage patterns of data-intensive mobile applications, that is, those related to social, news, and video, vary considerably [2, 3]. From these statistics and others, it becomes apparent that having the same mobile network architecture serving all mobile users, let alone all mobile applications, despite the diversity they exhibit in their attitudinal response to mobile services, have to be rethought.

Furthermore, a mobile user usually subscribes to a single mobile operator that provides the delivery of all the mobile services. In addition, a single mobile network usually ensures communication for all service types, regardless of the suitability of its available functionality to deliver these services with acceptable quality of experience (QoE) and network efficiency. Due to the uniformity of the network, all users of the network are charged based on the same bandwidth consumption model, which fails to capture the specifics of the usage of applications with large overhead and makes too expensive the network for individual large-scale sensor deployments.

Last but not least, current mobile core networks are offering a uniform ubiquitous service for all the connected devices. Even if a mobile user moves far away from the mobile core network infrastructure, he/she remains serviced by the same core network, even in the case of a highly decentralized mobile system [4]. This feature may impact numerous emerging advanced mobile services with strict latency and jitter requirements (e.g., augmented reality and self-driving vehicles). High latency and jitter degrade such mobile services, rendering their respective devices unusable, ultimately turning users away and impacting revenues. Also, a large number of users do not move from a specific network area, their mobility support being only an additional overhead on the network.

To cope with the above, this article advocates the need for customizing mobile telco services, through providing a functional differentiation for the different user requirements. Herein, the intention is to leverage the emerging technologies in the areas of network function virtualization (NFV) [5], software defined networking (SDN), and cloud and edge computing for providing a single infrastructure on top of which multiple versions of the same software, generically named slices, customized to have specific behaviors are running, through this removing the overhead of the uniform network service. Furthermore, in order to account for different layers of granularity, this article underlines the need and the technical possibilities to support a very large number of slices as well as their appropriate deployment according to the momentary location of the subscribers.

This article is organized in the following fashion. We present a quick overview on network slicing and highlights its utility for the Personalized Mobile Telecom (PERMIT) vision. The overall PERMIT framework is portrayed. We introduce the PERMIT slice orchestration system and discusses its challenges. The article then concludes.

NETWORK SLICING

The term "network slicing" has captured much attention within research communities and the industry, as well as standards development organizations (SDOs), such as the Next Generation

Digital Object Identifier: 10.1109/MCOM.2017.1600947

Tarik Taleb is with Aalto University and Sejong University; Badr Mada is with Aalto University; Marius-Iulian Corici3 is with Fraunhofer Fokus; Akihiro Nakao is with the University of Tokyo; Hannu Flinck is with Nokia Bell Labs.

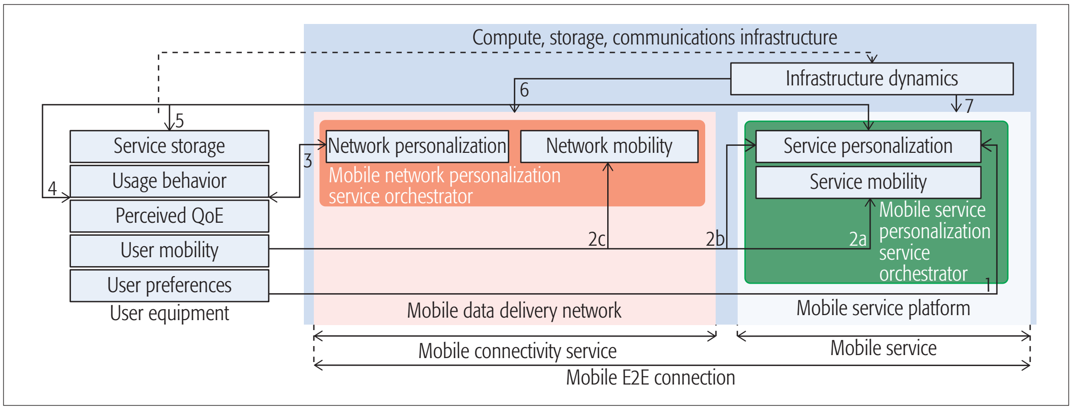

Figure 1. PERMIT virtual mobile network architecture.

These lightweight VMNs are expected to have the flexibility to serve up to a granularity of a single service for one individual user, multiple services for one individual user, or a group of users, although due to scalability reasons, most probably users with same services and behavior will be grouped into the same slice.

Mobile Network Alliance (NGMN), Third Generation Partnership Project (3GPP), and International Telecommunication Union — Telecommunication Standardization Sector (ITU-T). Although the definition of network slicing is still under heavy discussion, it generally means an isolated collection of resources and functions implemented through software programs on top of the resources, flexibly allocated on demand in order to enable quality of service (QoS) guarantee for the network requirements as well as in-network processing along end-to-end communications.

In PERMIT, network slicing (slicing hereafter) is considered to be one of the most important concepts to realize personalization of mobile networks for users. Although slicing in the mobile networking context has often been addressed for enabling different classes of communications (e.g., enhanced mobile broadband, massive machine-type communications, and ultra reliable and low latency communications), in PERMIT, slices can be instantiated per user, and in the most extreme case, per device and/or per application.

It is a well-known fact that the concept of slice in networking was first introduced in the overlay network research efforts, such as PlanetLab, in 2002. At that time, a slice was defined as an isolated static set of resources allocated for a group of users who "program" network functions and services over their overlay network, overlaid across "the planet." Since various network virtualization testbed efforts (e.g., GENI, VNode, FLARE, and Fed4Fire) have inherited the concept of slices as a set of programmable, dynamically allocated resources to tailor new network services and protocols, it is quite natural to use the slice concept in PERMIT to further personalize the access and in-network edge processing for mobile network users.

The PERMIT Framework

Figure 1 schematically depicts the main components of the PERMIT framework. It mainly consists of two orchestrators, the mobile network personalization service orchestrator (MNP-SO) and the mobile service personalization service orchestrator (MSP-SO). PERMIT also envisions some changes to the user equipment as detailed below. The MNP-SO and MSP-SO entities can run separately or jointly on dedicated hardware or as software on virtual machines (VMs) with adequate characteristics. These two entities incorporate all necessary intelligence for mobile service personalization and mobile network personalization, respectively. They are decision making entities that decide how mobile services and the lightweight virtual mobile network (VMN) [6, 7] to transport them shall be personalized to the current and anticipated needs of a mobile user or a group of mobile users. In their decision making procedures, both MNP-SO and MSP-SO take into account the underlying infrastructure dynamics, processing and storage capabilities of user equipment, and users' contextual information like mobility and resource usage patterns.

In the envisioned PERMIT framework, lightweight VMNs are expected to run as a slice on one or multiple instantiated virtual resources [6, 8]. These lightweight VMNs are expected to have the flexibility to serve up to a granularity of a single service for one individual user, multiple services for one individual user, or a group of users, although due to scalability reasons, most probably users with same services and behavior will be grouped into the same slice. This flexibility can be attained by composing the lightweight VMN of uncorrelated building blocks that can be freely and dynamically combined or separated as per the requirements of the target mobile services and the needs of the serviced users [6].

Indeed, in PERMIT, a mobile network component (a network function) is defined in terms of its application logic and data, as depicted in Fig. 2. The application logic is decomposed into compute "blocks," including one basic block that provides minimal core network functionalities and multiple added value blocks, extending the basic components to provide additional network services such as communication reliability, QoS, and mobility support or accounting. The blocks run either at the network edge or in the cloud, depending on latency, bandwidth, resilience, and security requirements.

The composition of the blocks should follow the latest frameworks in service composition such as the ones based on micro-service application programming interfaces (APIs) [9] or on event bus communication. Both of these ensure the loose coupling between the main functionality and the other modules, permitting the addition of new functions on demand and even during runtime, as well as a comprehensive separation of the liabilities. The PERMIT architecture will further optimize these mechanisms to address the end-to-end delay while processing a request through multiple compute blocks.

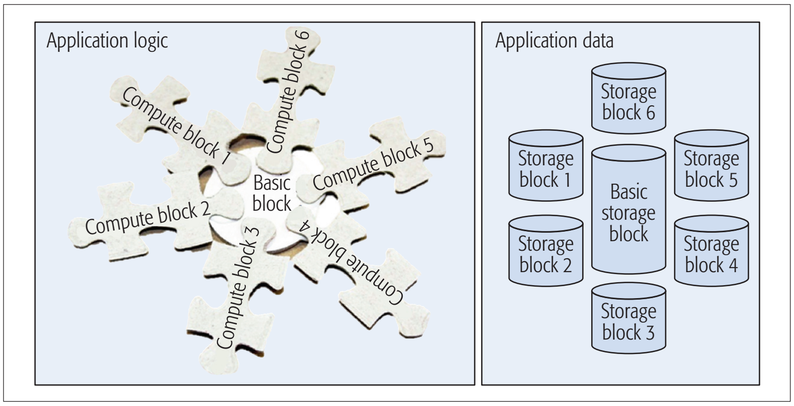

Figure 2. Network functions consisting of application logic and application data.

The PERMIT architecture enables the configuration of the in-service parameters related to the network functions as well as to the composition of the network functions within end-to-end services. This approach enables the creation of lightweight, customizable, and truly elastic mobile networks with network services/blocks that can be adapted to the users' needs. It will also enable seamless decomposition of application logic and data, allowing them to be moved to more convenient network locations. It is worth noting that careful attention shall be paid to how to gracefully compose the different functionalities within the network functions, especially considering the performance of multi-vendor software components, in order to provide a retraceable service level as well as to separate liabilities in case of failures.

At the protocol level, the lightweight feature of VMNs can also be achieved by simplifying a number of procedures typically required for mobile networks like authentication, authorization, and accounting (AAA) and charging functions. Dependencies on data anchoring (i.e., the packet data network gateway in the Evolved Packet System, EPS) and mobility anchoring (i.e., the serving gateway in EPS) concepts shall be relaxed if not completely replaced, as in the follow me cloud concept [10, 11]. With customizable VMNs, currently unsupported communication modes that mimic connectionless communication over shared media become possible over mobile networks.

In PERMIT, the personalization of mobile networks for a mobile user or a group of mobile users is achieved by anticipating the needs of the mobile services of this user or this group of users. Indeed, once the needs of a mobile service or a set of mobile services received by an individual user or a group of users are anticipated, the right VMN with the right characteristics (e.g., composing building blocks, total number of VMs involved, the locations of their respective DCs, their respective CPUs/memory/storage) can be identified so that VMNs can scale up and down as per the assessed needs. Another aspect of VMN personalization consists of its mobility to a different data center when required. For this purpose, it is possible to leverage different algorithms and mechanisms [12] that decide on and enforce the VMN mobility as per the mobility patterns of the served mobile users and/or the dynamics of the underlying infrastructure, in such a way that the "mobile network," serving a user or a group of users, follows their mobility. This decision may be based on several possibly conflicting attributes/criteria such as the mobile service type (e.g., delay-sensitive), the perceived quality of experience (QoE), the migration cost, the activity level of the users, the usage behavioral patterns, the mobility patterns, and the dynamics of the underlying communication infrastructure (e.g., for load balancing). Inputs, used for VMN personalization and VMN mobility, and relevant to users' mobile service usage behavior, perceived QoE, a user's mobility, and dynamics of underlying communications infrastructure are schematically depicted in Fig. 1 through arrows 2c, 3, and 6.

The personalization of the mobile service depends first of all on the user preferences for the service delivery (arrow 1 in Fig. 1) and on its mobility (arrow 2). Based on insights on the user behavior and on its perceived QoE, the network is customized according to the user needs (arrow 3). As multiple users may have the same network requirements, the customization can be seen as a user classification problem toward the appropriate cluster of users that have the optimal handling of the communication requirements. Similar to network customization, a service customization may also be executed (arrow 4). As the users may come from different customized networks, mainly due to their multi-application terminals, the customization of the applications should consider the customization of the network as a given parameter. A further step in the customization is the distribution of the service data (arrow 5) to the user equipment (UE) when needed or when the network conditions are appropriate, depending on the specific user behavior.

Finally, the customization highly depends on the availability of the infrastructure resources for the specific customization (arrows 6 and 7). The decisions of both the MNP-SO and MSO-SO depend on the possibility of the infrastructure to support their needs at the specific location. As the PERMIT architecture assumes a very large number of slices (i.e., up to one for each network user), it is possible that the network infrastructure will not have enough momentary resources to handle the subscriber communication. In this case, the subscribers should be classified in a default communication class for which the processing is handled by a central data center within a default slice similar to the current network infrastructure.

THE PERMIT VMN SLICE ORCHESTRATION SYSTEM

PERMIT aims to achieve elasticity, flexibility, dynamicity, scalability, manageability, and efficiency beyond the current network level by building on-demand VMN slices, customized to the service requirements, thus much reducing the network overhead.

The PERMIT architecture is envisioned to function for different slices ranging from slices of individual users to Internet of Things (IoT) application slices as well as for verticals such as industrial control systems, autonomous driving, virtual reality, or video streaming. The resulting system will consist of numerous network slices, running in parallel and composed for end-to-end service delivery (Fig. 3).

Each slice consists of a set of virtual network functions (VNFs) within both the control and data

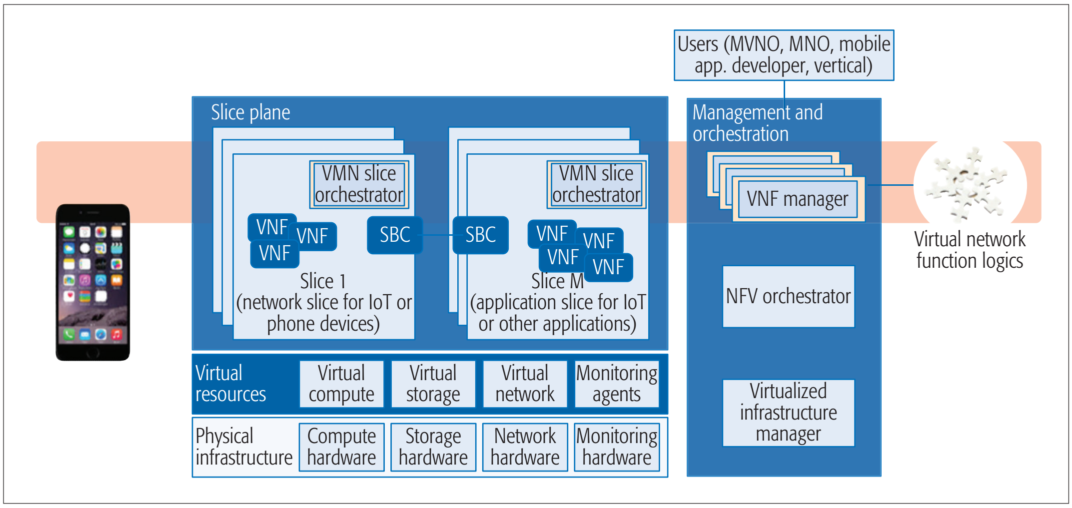

Figure 3. PERMIT VMN Slice Orchestration System: diverse VMN slices running in parallel and serving diverse verticals or individ-ual users.

The role of the VNF manager is to transmit network function placement [12] and scaling requirements to the NFV orchestrator, as well as to transmit to the VNFs the dependency parameters in order to enable the communication between the different VNF Components within a slice.

planes, customizable to the particular service types or vertical market needs or personalized to the individual end user, as presented earlier. From the perspective of the orchestration, the VNF components of each slice are seen as software functions that may be composed and customized, thus transparent to the functions they handle.

The VNFs accommodate the intrinsic features of the slices and their changing requirements, such as scaling up to match sudden growth in their traffic or smooth mobility to another network location. Also taking into account the requirements of the service delivery in terms of latency, reliability, and security together with the large number of slices, resource control becomes highly complicated. With the deployment of a large number of software services across multiple cloud and edge data centers, the complexity of the system increases beyond the capabilities of a single orchestration node. Moreover, a centralized network orchestrator may not be able to make the appropriate decisions and enforce them in due time, as it needs to handle a large amount of runtime operations, especially related to the sharing of the common data path environment. These delay and scalability limitations can be overcome through a distributed orchestration system where parts of the orchestration functionality are delegated to the edge nodes [13]. Data path sharing between the different slices and the decisions that require immediate response, such as network function failures, are particularly suitable for such delegation.

Figure 3 shows a high-level architecture of the PERMIT VMN slice orchestration system. The physical infrastructure consists of hardware for computing, storage, networking, and monitoring. These equipments can be administrated by the same entity or could belong to different domains. The slice orchestration plane of the architecture include images of VNFs, which represent the software version of existing network equipment. These VNFs could consist of building blocks (Fig. 2) designed in a clean-slate fashion or as components of existing network equipment. The VNF slice orchestration system is the main component of the architecture. It creates slices of VNFs for an individual user or a group of end users of a vertical. These slices can be created following pre-defined blueprints or in a fine-grained fashion, taking into account inputs relevant to end users' mobile service usage behavior, perceived QoE, and mobility, as discussed earlier. The slice orchestration system can be owned by a cloud provider, a mobile operator, or a new stakeholder. The users of the system can be vertical providers (e.g., automotive and IoT service provider), a mobile application developer, or an individual end user wishing for personalized mobile telecommunication service. Users can communicate to the slice orchestration system via well defined northbound interfaces (e.g., following Open Mobile Alliance, OMA, guidelines).

In order to separate the orchestration concerns, the following PERMIT orchestration levels are considered. First, a basic NFV resource orchestrator is considered, able to broker the available virtual resources, as provided by a virtualized infrastructure manager (VIM), to the different slices. The NFV orchestrator receives resource allocation requests from the VNF managers, one for each slice, which are aware of the specific slice logic. The role of the VNF manager is to transmit network function placement [12] and scaling requirements to the NFV orchestrator, as well as to transmit to the VNFs the dependency parameters in order to enable the communication between the different VNF Components within a slice.

To reduce the complexity and to appropriately manage the slice-specific operations, a VMN slice orchestrator is added to the architecture having the role of managing the functionality within the specific slice, including fault, configuration, accounting, performance, and security (FCAPS),[1] adapted to the dynamic resource environment. This includes the acquisition of monitored data on the specific service agents within the slice, a composition logic for the VNF components, enabling the appropriate processing flow allocation according to the momentary available resources as well as the interaction with other slices within the system. Another important role of the VMN slice orchestrator could also be offering flexible service function chaining (SFC) as a service, indicating the forwarding graph/path that a set of VNFs should be following within the respective slice. For this, the management system as well as the communication plane are extended with a slice border

[1] ISO/IEC 10040, 1998 — Information Technology — Open Systems Interconnection Systems Management Overview.

From the users' perspectives, PERMIT will facilitate a fully personalized and elastic end-to-end mobile connection service, which will provide mobile users with easy and efficient access to advanced mobile services. Indeed, with PERMIT, a mobile user may have his mobile network fully personalized to his current and anticipated needs or to the requirements of his mobile services.

control (SBC), which is similar to the session border controllers in the current architecture, and enables the filtering and the classification of both the inbound and outbound data traffic, the application-level firewall, as well as the appropriate forwarding to the components within the slice or to the SBC of the peer slice.

Through this separation of concerns within multiple functions, the possible policy conflicts are mitigated, although due to the shared resources, it may happen that the exact requirements of a specific slice would be partially fulfilled by the infrastructure. Considering the large number of slices, the alternative to provide the information on available resources to each VNF manager is highly complex.

From the users' perspectives, PERMIT will facilitate a fully personalized and elastic end-to-end mobile connection service, which will provide mobile users with easy and efficient access to advanced mobile services. Indeed, with PERMIT, a mobile user may have his/her mobile network fully personalized to his/her current and anticipated needs or to the requirements of his/her mobile services. True service elasticity will be attained: services of heavy mobile phone users shall never be throttled. Fair charging models can also be achieved: light and medium mobile users will have their respective VMNs running on smaller VMs, which will enable them to be fairly charged only for what they have indeed consumed. A group of users of a vertical or receiving a particular service may have a mobile network fully customized to their needs and the requirements of their mobile service. This shall enable the much desired service-tailored mobile networking concept. Furthermore, different VMNs with the right processing features can be created for different services as per the specifications of each service. Instead of being "locked in" the same mobile network for the delivery of all service types, users will then have the flexibility of subscribing to the most suitable VMN slice to receive a particular service type. Accordingly, subscription to multiple VMNs for multiple service types becomes possible. This requires UEs to have the capability to simultaneously connect to and steer mobile traffic across multiple VMN slices, optimally created for a set of services (Fig. 4). A default connectivity slice will always be assumed for each user, available for applications that use non-customized service delivery. A UE shall be able to discover existing VMN slices and request the creation of a new one. The creation of a new VMN slice for a particular service can be UE-initiated (following a set of rules and policies) or network-controlled when discovering that the current slice used for the application is inefficient for the delivery of the service to the user or to a group of users. A UE shall also have the ability to leave a VMN slice, join an existing one, or upgrade an ongoing one, either upon a UE request or upon a network-based classification of the UE.

The slice concept can be extended further to support the slicing of resources within a single UE. In the current smartphone market, it is quite common that UEs come with application vending facilities such as Google Play on Android phones, iTunes App Store for iPhones, iPads, and iPods, and proprietary application marketplaces on other smartphones. Applications on top of UEs must be examined and approved by their respective vendors and must run in the sandboxes they have prepared. We posit that the application marketplaces and the sandboxes as execution environments should be implemented, isolated within separate slices, so that one can have multiple, different, personalized execution environments and available application suites within a single UE. With this envisioned UE slicing concept, one may have multiple personalized containers within a single UE, so one may have different security and privacy contexts such as private usage and corporate usage as well as automatic personalized updates directly from the software developers. Considering data contamination and privacy breaches often observed in the mixed use of a single UE for private and public matters, it makes a lot of sense to introduce isolation between different usages. For example, one may want to isolate the applications with personal data (e.g., medical record and bank account information), those with confidential data (e.g., corporate confidential information), and those with public data (e.g., general web browsing applications) exposed to rather wild and rogue environments where security breach is often observed. If multiple UEs are carried to avoid such mishaps, it is reasonable to consolidate them into a single one using slicing techniques. Alternatively, one may have multiple UEs with different operating systems, so one may benefit from different application suites and execution environments. Embedded operating system virtualization and network virtualization technologies have already advanced to support such a concept of UE slicing [3, 16]. However, to the best of the authors' knowledge, PERMIT is the first to consider end-to-end slicing, including UE slicing, mobile network slicing, edge computing, and cloud computing. The authors are also already aware of the challenges of defining the granularity of slicing, such as the necessity of individual UE-level slicing and the feasibility of slicing in mobile operators. Although compelling application use cases drive such decisions, the authors have already conducted a feasibility study of application-specific slicing concepts [14, 15]. We plan to address these challenges in our future research work.

In PERMIT, a user can have his/her personalized VMN and his/her personalized mobile services constantly following him/her. In this way, PERMIT will support a wide gamut of high-quality services, customized to users' preferences, behavior, and mobility features. Emerging devices, such as Microsoft HoloLens, will largely benefit from the PERMIT approach, particularly when their users are moving in dense smart cities onboard smart connected vehicles.

PERMIT shall also define novel and promising business opportunities for cloud providers, especially in the area of providing value-added services beyond the basic infrastructure sharing. It also represents an innovative and ambitious solution to open up mobile networks and revolutionize the mobile networking principles, going from large-scale ubiquitous uniform connectivity service to highly efficient service support tailored to the specific needs of individuals. In PERMIT, a new set of business stakeholders may also emerge, orchestrating the mobile service and mobile network personalization for individuals and groups of users as well as orchestrating the interaction between the mobile service slices.

Conclusion and Future Research Directions

In this article, we propose the PERMIT approach, which is expected to act as a catalyst for structural changes to the current communication system configuration, whereby both the mobile delivery network and the mobile services it supports are personalized for each individual user, or alternatively customized for groups of users/verticals.

In PERMIT, service personalization goes beyond the traditional approaches whereby service personalization is based on the classical users' preferences, explicitly indicated by the users or deducted through collaborative filtering from other means (e.g., social networks), resulting in a specific parametrization of the uniform network. In PERMIT, users' mobility patterns, their mobile service usage behavioral patterns, and the dynamics of the underlying communications infrastructure are all taken into consideration for acquiring both network and service personalization, treating service personalization and networking customization as two flexibility-enabling and complementary components. To always ensure short response times for emerging advanced mobile services, the mobility of mobile services and the personalized mobile networks are both enabled toward the proximity of the respective mobile users across the overall mobile service area, as per the mobility features of mobile users, and in a seamless and cost-efficient manner. This will make both the mobile delivery network and the mobile services — after being personalized — constantly follow their respective mobile users.

With such an approach, service personalization and network personalization can take place dynamically and interactively. This also enables the much desired service-tailored mobile networking concept, achieving a so far unprecedented level of flexibility in service-specific optimizations, fine-grained network resource slicing, and transparent capacity scaling. In this regard, scalable programming of data plane and data paths and fine-grained mobility management of various services, considering both centralized and distributed approaches, are needed. This defines a promising research area that will stimulate the relevant community of researchers.

With such an approach, service personalization and network personalization can take place dynamically and interactively. This also enables the much-desired service-tailored mobile networking concept, achieving a so far unprecedented level of flexibility in service-specific optimizations, fine-grained network resource slicing and transparent capacity scaling.

Acknowledgments

This work was partially supported by the TAKE 5 project funded by the Finnish Funding Agency for Technology and Innovation (TEKES) and in part by the Finnish Ministry of Employment and the Economy. It is also partially supported by the European Union's Horizon 2020 research and innovation programme under the 5G!Pagoda project with grant agreement no. 723172.

References

[1] Y. Jin *et al.*, "Characterizing Data Usage Patterns in a Large Cellular Network," *Proc. SIGCOMM CellNet Wksp.*, Helsinki, Finland, Aug. 2012.

[2] P. Du and A. Nakao, "Application Specific Mobile Edge Computing through Network Softwarization," *IEEE CloudNet '16*, Pisa, Italy, Oct. 2016.

[3] A. Nakao, P. Du, and T. Iwai, "Application Specific Slicing for MVNO through Software-Defined Data Plane Enhancing SDN," *IEICE Trans. Commun.*, vol. E98-B, no. 11, 2015, pp. 2111–20.

[4] T. Taleb, K. Samdanis, and A. Ksentini, "Supporting Highly Mobile Users in Cost-Effective Decentralized Mobile Operator Networks," *IEEE Trans. Vehic. Tech.*, vol. 63, no. 7, Sept. 2014, pp. 3381–96.

[5] ETSI GS, "Network Function Virtualization (NFV) Management and Orchestration," NFV-MAN 001 v. 0.8.1, Nov. 2014.

[6] T. Taleb *et al.*, "EASE: EPC as a Service to Ease Mobile Core Network," *IEEE Network*, vol. 29, no. 2, Mar. 2015, pp. 78–88.

[7] T. Taleb, "Towards Carrier Cloud: Potential, Challenges, & Solutions," *IEEE Wireless Commun.*, vol. 21, no. 3, June 2014. pp. 80–91.

[8] T. Taleb, A. Ksentini, and A. Kobbane, "Lightweight Mobile Core Networks for Machine Type Communications," *IEEE Access*, vol 2, Oct. 2014, pp. 1128–37.

[9] I. Nadareishvili *et al.*, *Microservice Architecture: Aligning Principles, Practices, and Culture*, O'Reilly, 2016.

[10] T. Taleb and A. Ksentini, "Follow Me Cloud: Interworking Federated Clouds and Distributed Mobile Networks," *IEEE Network*, vol. 27, no. 5, Sept./Oct. 2013. pp. 12–19.

[11] A. Ksentini, T. Taleb, and F. Messaoudi, "A LISP-Based Implementation of Follow Me Cloud," *IEEE Access*, vol. 2, Oct. 2014. pp. 1340–47.

[12] T. Taleb, M. Bagaa, and A. Ksentini, "User Mobility-Aware Virtual Network Function Placement for Virtual 5G Network Infrastructure," *Proc. IEEE ICC 2015*, London, U.K., June 2015.

[13] A. Aissioui *et al.*, "Elastic and Distributed SDN Controllers for Follow-Me Cloud," *IEEE Access*, DOI 10.1109/ACCESS.2015.2489930, vol. 3, Nov. 2015.

[14] C. Dall *et al.*, "The Design, Implementation, and Evaluation of Cells: A Virtual Smartphone Architecture," *ACM Trans. Comp. Sys.*, vol. 30, no. 3, Aug. 2012.

[15] J. Andrus *et al.*, "Cells: A Virtual Mobile Smartphone Architecture," *Proc. 23rd ACM Symp. Operating Systems Principles*, Oct. 2011.

[16] A. Nakao and P. Du, "Application and Device Specific Slicing for MVNO," *1st Int'l. Science and Tech. Conf. SDN & NFV*, Moscow, Russia, Oct. 2014.

Additional Reading

[1] A. Ksentini *et al.*, "On Using Bargaining Game for Optimal Placement of SDN Controllers," *IEEE ICC '16*, Kuala Lumpur, Malaysia, May 2016.

Biographies

Tarik Taleb (tarik.taleb@aalto.fi) is currently a professor at Aalto University, Finland, leading the MOSA!C Lab. He worked as a senior researcher at NEC Europe Ltd until April 2015. Prior to that, he worked as an assistant professor at Tohoku University, Japan. He received his B.E. degree in information engineering with distinction, and his M.Sc. and Ph.D. degrees in information sciences from Tohoku University in 2001, 2003, and 2005, respectively.

Badr Mada (badr.mada@aalto.fi) is currently conducting his doctoral studies at Aalto University. His research focuses on mobile edge computing and open source networking. He received his Bachelor degree in mathematical and computer science and his Master degree in software engineering from University Mohammed V, Rabat, Morocco, in 2014 and 2017, respectively.

Marius Corici (marius-iulian.corici@fokus.fraunhofer.de) has been a senior researcher in Fraunhofer FOKUS's NGNI department for 12 years. He is currently deputy director of the Software Networks division in charge of research and innovation in the areas of 5G, NFV, SDN, and massive IoT, and the development of the correspondent software toolkits, sustaining the industry and academia R&D to obtain and demonstrate meaningful results with high impact toward standardization.

Akihiro Nakao (nakao@iii.u-tokyo.ac.jp) received his B.S. (1991) in physics and M.E. (1994) in information engineering from the University of Tokyo. He was at IBM Yamato Laboratory, Tokyo Research Laboratory, and IBM Texas Austin from 1994 to 2005. He received his M.S. (2001) and Ph.D. (2005) in computer science from Princeton University. He is a full professor and also a department chair at the Interfaculty Initiative in Information Studies, Graduate School of Interdisciplinary Information Studies, University of Tokyo.

Hannu Flinck (hannu.flinck@nokia-bell-labs.com) is a research manager at Nokia Bell Labs, Espoo, Finland. He received his M.Sc. degree (1986) and Lic. Tech. degree (1993) in computer science and communication systems from Helsinki University of Technology, Finland. His current research agenda includes cloud technology, SDN, and content delivery in fifth generation mobile networks.