

Towards Real-time Management of Virtualized Telecommunication Networks

John Keeney
 Network Management Lab,
 Ericsson,
 Athlone, Co. Westmeath,
 Ireland
 john.keeney@ericsson.com

Sven van der Meer
 Network Management Lab,
 Ericsson,
 Athlone, Co. Westmeath,
 Ireland
 sven.van.der.meer@ericsson.com

Liam Fallon
 Network Management Lab,
 Ericsson,
 Athlone, Co. Westmeath,
 Ireland
 liam.fallon@ericsson.com

Abstract—The idea of virtualizing network functions is driven by recent advances in network-focused hardware. In 2012, several large telecommunication operators issued a call to action for Network Function Virtualization (NFV)¹. The underlying idea is that every network service currently delivered on proprietary, application specific hardware should be deliverable using virtual machines. This means that routers, firewalls, load balancers and other network devices run virtualized on commodity hardware. Consequently, ETSI is extending this idea to mobile networks. Here, parts of the core and the radio access network will be virtualized. The immediate advantage is that any Virtual Network Function (VNF) can now be deployed, re-deployed and un-deployed in the same way as any traditional virtual machine. Thus, NFV will result in more dynamic and agile networks than seen heretofore. However, this will raise a number of serious issues in managing these future networks. In this paper, we examine issues and challenges in orchestrating these virtualized functions and their interconnections to provide a more agile mobile telecommunication network.

I. INTRODUCTION

Traditionally the business model for network equipment vendors has been to sell the necessary software, equipment, and appliance nodes to realize and manage large-scale telecommunication networks. However, given recent advances in hardware virtualization and SDN, it has now become feasible to replace bespoke dedicated network appliances, often operating on proprietary hardware nodes, with Virtual Network Functions (VNF) running on industry standard, off-the-shelf high volume servers, storage devices, and switches.

In this paper we focus on Network Function Virtualization (NFV) [2] for mobile telecommunications networks. For example, the 3GPP EPS Network Architecture [1] shown in Fig.1 is comprised of a collection of networked dedicated hardware appliances. These dedicated nodes include: eNBs, base stations in the Radio Access Network; MMEs, the control nodes that manage eNBs to provide a radio network; SGWs, to route user traffic to and from eNBs; PGWs, act as the gateways between the LTE network and external Packet Distribution Networks; and HSSs, databases to store user and subscription information. Most nodes maintain standardised connections to

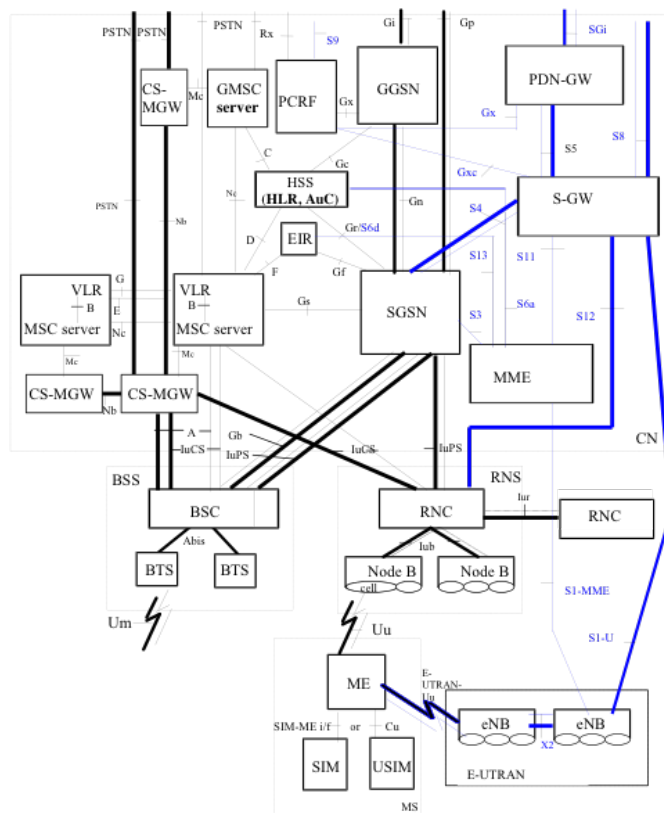


Fig. 1. The 3GPP Evolved Packet System (EPS) Network Architecture [1]

each other, for example, eNBs are connected to their neighbours using X2 connections. eNBs are connected to MMEs and SGWs using S1 connections, MMEs are interconnected using S10 connections, HSSs are connected to MMEs using S6a connections, PGWs are connected to SGWs using S5 and S8 connections. Moreover, mobile networks are increasingly heterogeneous and may be running four or more radio technologies (such as 2G, WCDMA, LTE, Wireless LAN), and may have very complex radio architectures with macro, micro, pico, and even femto cells [3]. This leads to networks with a plethora of different node types and connections types, with tens-of-thousands of physical nodes and connections for even

¹http://portal.etsi.org/nfv/nfv_white_paper.pdf

a modest network.

One of the main advantages of providing VNFs to fulfil role of physical networking appliances is that these VNFs can be deployed, redeployed and undeployed much faster than physical nodes. Although the architecture of the Core Network (CN) will not need to change initially and only small changes may be needed in the Radio Access Network (RAN), moving from bespoke appliances to virtual appliances will require significant revisions to Operation Support Systems (OSS) and Business Support Systems (BSS), not just to continue network Operation and Management (OAM) tasks as before, but to handle and exploit the dynamicity and flexibility made possible by NFV. To manage a virtualized network, operators will need to manage the VNFs themselves as well as the underlying cloud infrastructure upon which the VNF runs. It will also be necessary to maintain a consistent network deployment in virtual environments and between virtual and physical environments.

There are likely to be many more virtual nodes in a virtualized network than in a network with the same characteristics built using dedicated hardware appliances. This is because virtual nodes are unlikely to have performance characteristics equivalent to dedicated hardware appliances and because, for manageability, many small virtual nodes with single functions are likely to be deployed rather than large virtual nodes with multiple functions. These considerations mean that NFV will result in more large-scale, dynamic and agile networks than seen heretofore. However, this will raise a number of serious issues in managing these future networks. Just some of the the problems that need to be addressed include: extracting the functionality of the nodes into modular VNFs, having VNF interoperating with physical network nodes, mapping VNFs to physical servers, maintaining and transferring state between functions as they are deployed or undeployed, and maintaining fault tolerance in the event of problems with the hardware hosting VNFs. In this paper we examine issues in planning, orchestrating, and managing these functions and their interconnections to provide a more agile mobile telecommunications network.

II. BACKGROUND AND RELATED WORK

Existing dynamic cloud applications tend to focus on increasing or decreasing the number of VMs or visualized applications they are running as load changes [4][5]. However, many cloud-based applications do need not to consider handling of long-lived shared state across application component boundaries where applications may have interrelationships and connections that must be explicitly managed when applications or component instances are added or removed. In contrast to most stateless distributed/cloud architectures such as RESTful web services [6], network functions and their management generally use stateful session semantics to interact with each other and the management infrastructure. Network functions establish sessions with other network functions, but such sessions are often very long lived and may be carrying connection-oriented user-plane traffic. One example of such

a connection-oriented stateful connection approach can be seen with management applications that connect to managed network elements to collect statistical information [7], carry out configuration operations [8] or receive event streams [9]. A particular instance of a management application using such long-lived stateful sessions must manage the session state for the duration of that connection session. It is, therefore, inherently difficult to exploit flexible distributed architectures such as cloud platforms for dynamic network function virtualization because network session state between VNFs must be managed when the number of VNF instances change. In order to change the connections between VNF instances old sessions must be allowed to close naturally, or else be torn down re-established according to the new VNF topology.

Network operators are already defining and realizing their strategies around network virtualization. AT&T published their Domain 2.0 strategy in 2013 [10]. In this strategy document, NFV Infrastructure (NFVI) is considered as being similar to traditional cloud technology whilst leading to differing functional and engineering trade-offs. Three main areas are introduced: footprint and distribution; separation of data, control, and management; and throughput intensity. As well as the challenges discussed in this paper, security is emphasized as a major challenge for NFV. Operators and vendors also beginning to deploy and test cloud infrastructure. At the OpenStack summit in May 2014, Verizon presented their NFV strategy and their experiences in deploying OpenStack [11]. The Verizon use case focused on a content distribution network with several OpenStack environments. Some areas requiring further work include: improved high-availability of infrastructure and support functions, hypervisors still lack required monitoring and control functionality, and network virtualization is not yet entirely satisfactory for production systems.

Telecom network equipment providers are however making progress in the NFV domain, with some interesting examples below. Ericsson have recently announced their ECS cloud platform, which will deploy, manage, and host NFV applications, with virtualized versions of network nodes and applications already available [12]. NSN have demonstrated a number of core network VNFs running in a cloud environment [13], while NEC, Huawei, ZTE and Alcatel-Lucent have also announced support for virtual core network nodes [14][15][16][17]. However techniques for orchestrating VNFs remains immature in all cases.

Recent initiatives in standards organisations have demonstrated the need for coordinated management and orchestration of VNFs.

The ETSI NFV Industry Specification Group (ISG)² have proposed an NFV Orchestration component (see Fig.2) as part of their NFV MANO workgroup. This component is responsible for VNF lifecycle management including: registration and on-boarding of VNFs and service templates into a catalogue; managing instantiation of VNFs; managing the

²www.etsi.org/nfv

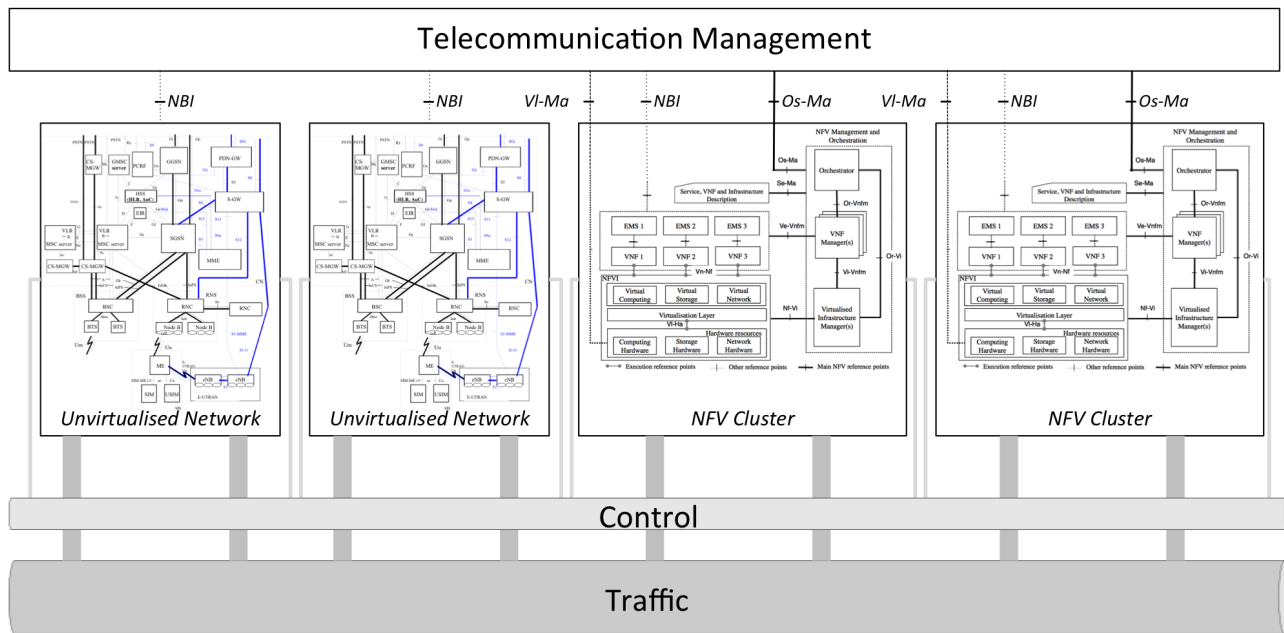


Fig. 3. Management of Virtualized and Unvirtualized Telecommunication Networks

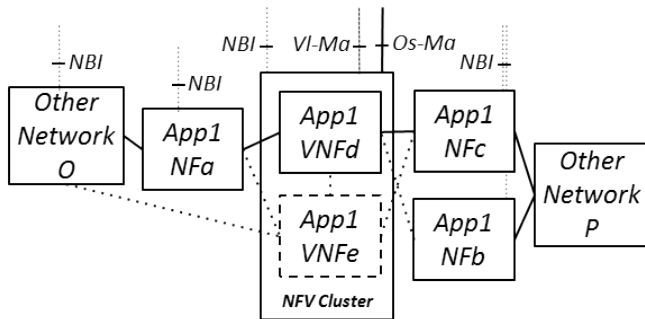


Fig. 4. Abstract Network Application on Real and Virtualized NFs

service. Once operational the fundamental operation of the node and its connections will remain stable and change rarely.

When network functions are virtualized, a telecommunication management system must adapt to handle the dynamic creation, modification, and deletion of network functions and connections in the networks it is managing. To illustrate the challenge virtualization introduces, consider the trivial case of introducing, modifying, or removing a single VNF in the idealized network application shown in Fig.4. The management system monitors each already running physical or virtual network function $a..d$ of the network application using the various NBI interfaces, as well as handling NFV virtualization and orchestration. As the load on the application increases, it triggers an increase in capacity of the network application. It plans introduction of VNF_e , any necessary synchronisation of state between VNF_d and VNF_e , together with all necessary connections, depicted as dashed lines in Fig.4. The management system then uses the $Os-Ma$ interface towards NFV orchestration to spin up VNF_e . It then uses the

NBI towards network functions $a..d$ as well as Other Network O to configure the connections to VNF_e . Modification and deletion scenarios follow a similar sequence of execution.

IV. CHALLENGES

Virtualization eliminates the dependency on hardware to provide network functionality. NFV requires a homogeneous execution environment and the management interfaces required to enable elastic deployment and interconnection of VNFs. Pooling should be used for the infrastructure and the VNFs themselves, resulting in a large, shared, and agile NFV Infrastructure (NFVI). This NFVI should provide resiliency, load balancing, and distribution facilities, along with common management functionality to be monitored, configured and audited. This means that we can categorize the management challenges for the Management and Orchestration (MANO) of an NFV framework as follows:

- 1) Understanding and managing the complexity of the NFV infrastructure via the $Nf-Vi$ reference point: monitor and pool the infrastructure required to deploy, coordinate, modify and undeploy VNFs, and issue notifications of performance or configuration related issues within the support infrastructure that will effect VNF instances.
- 2) Understanding and managing the underlying hardware infrastructure hosting the NFVI via the $Vi-Ma$ and $Vi-Ha$ reference points: add, configure, monitor and remove hardware resources according to the requirements of the VNFs.
- 3) Managing VNF instances, VNF pools, and VNF clusters as software instances via the $Ve-Vnfm$ reference point: manage instantiation and initial configuration of VNF instances, including required connections to other

VNFs or non-virtualized parts of the network, configure forwarding graphs, manage all other aspects of the VNF life-cycle.

- 4) Managing the connections between the VNFs in form of NFV Forwarding Graphs (FG) as a logical abstraction realised by the software instances via the *Os-Ma* and *Se-Ma* reference point: manage the creation of an FG (including the registration of related VNF instances and their bindings to it), configure connectivity across clusters and pools, and monitor for the required service level agreements.

Monitoring and managing the hardware hosting the NFV Infrastructure will actually be substantially easier than managing the plethora of bespoke hardware currently required, but monitoring and managing the much more complex and dynamic NFV Infrastructure and the VNFs themselves has new technical and organisational challenges, some of which are outlined below:

- NFV Infrastructure - monitor and repair a much more complex infrastructure: In a virtualized environment, unlike in a dedicated hardware environment, monitoring may not have full knowledge of all the components being used to deliver a given network function. Take the example of monitoring a heartbeat for a network node. On dedicated hardware, internal monitoring agents will make sure that all parts of a network node are operational and send alarms in case of faults. In a virtualized environment, the heartbeat of each of several different components that are involved in executing a VNF must be monitored:
 - 1) physical nodes and connections (e.g. servers and switches in a data centre)
 - 2) operating system of the physical node (e.g. Linux services running on servers)
 - 3) cloud compute node (e.g. an OpenStack Nova compute node) and its connectivity to other cloud services (e.g. OpenStack connections to the API, management, and storage network)
 - 4) container hosts or hypervisors on the cloud compute node
 - 5) container instances, or virtual machines created by the hypervisors on the cloud compute node
 - 6) the actual VNF running within one of those virtual machines or container instances

To understand the conditions of the NFV infrastructure, the heartbeat information from all those six components need to be collected, aggregated, correlated and processed. Repair actions need to be selected (automatically) for any given combination of heartbeat signals, which requires a deep understanding on what those combinations describe.

- VNF instance management - monitor and repair a much more dynamic environment: with the ability to dynamically create, start, suspend, change, stop, and delete software realizing network functions, managing these net-

work functions now has significant additional challenges. These challenges include:

- Consistency of configurations - each VNF instance, pool, and cluster requires an initial configuration but is also subject to ad-hoc re-configuration. This makes it very difficult to maintain a consistent understanding of the current configuration of VNFs and requires novel mechanisms in managing and maintaining consistent configurations.
- Troubleshooting - Given the complexity of the monitoring of the NFV infrastructure, troubleshooting VNF instances requires understanding to deduce the cause of faults in NFV infrastructures and VNF instance/pool/cluster configuration. Faults might only relate to VNF instances but might also propagate from an infrastructure root cause. The dynamism of the virtual environment creates new challenges for root cause analysis.
- Reliability and availability - A mobile network must provide *five 9's availability*⁶. This means that, for any given VNF instance, management must ensure that enough resources are available at any given time to balance load across VNFs or to replace a VNF and take over all associated traffic. This means that the state information of traffic connections must be transferred seamlessly from one VNF to another.
- Managing NFV Interconnections - maintain consistency between abstractly defined Forwarding Graphs (FG) and the actual instantiated and configured VNFs. The FG represents the topology of a set of VNFs that fulfil a network service. It also represents the required state for the interconnected VNFs. Maintaining a correct current set of VNF instances requires continuous synchronization between the abstract FG (required state) and the current running set of VNF instances.

Two main organisational challenges also present themselves. Firstly, each part of an NFV framework can be owned separately, and secondly, each part of this framework can (will) involve heterogeneous hardware and software artefacts from multiple vendors. To allow for high-availability, NFV management must be able to consistently and securely monitor, control and troubleshoot across organisational domains (domains of ownership, administrative domains) and across a plethora of different hard-/software components in real time. While network and software management have approaches for solving these problems, virtualization extends the complexity of problem identification and resolution drastically.

As well as the challenges above, management of an overall telecommunication network must now consider a virtual (or partially virtual) network. This means that, for instance for the simple use case discussed in section III, we need novel ways to plan and optimize mobile networks (note: real network scenarios are much more complex). The following subsections discuss some of those challenges.

⁶Available 99.999% of the time

1) *Network Application Monitoring*: It will be necessary to monitor the NBIs of each network function as well as the *Vi-Ma* and *Os-Me* interfaces to decide when to trigger re-dimensioning of the network application. To do this, the management system must be aware of the network application domain and its deployment topology and have rules to decide which criteria will trigger re-dimensioning operations. In early systems, it is likely that re-dimensioning will be triggered manually.

2) *Planning Network Application Re-Dimensioning*: It will be necessary to model the required condition of the network application before re-dimensioning, and then build a plan of operation and instructions to be issued on the appropriate interfaces in order to execute re-dimensioning and achieve the required condition. Again, the management system must be aware of the network application domain and topology and handle the state changes required in each affected node in order to carry out re-dimensioning operations. Planning of network re-dimensioning may be executed by off-line planning tools in early systems.

3) *Re-Dimensioning Execution*: It will be necessary to control and coordinate the sequence of operations to execute re-dimensioning operations over the correct interfaces in the correct sequence. For example, when adding a VNF, the VNF must be spun up in the early steps of the operation whereas when removing a VNF, the VNF must be spun down towards the end of the operation. Each operation must have transaction support so that aborted re-dimensioning operations do not leave network operations in an inconsistent state.

V. CONCLUSION

In this paper we have focused on some the challenges for Network Function Virtualization (NFV) for mobile telecommunications networks. While the main advantages of providing VNFs is that they are much cheaper and their deployment is much more flexible than physical networking appliances, NFV will require significant changes to the way networks are deployed, operated and managed. Such changes are required, not just to provide network and service solutions as before, but also to and exploit the dynamicity and flexibility made possible by NFV. With virtual networks, operators will need to manage not only VNFs themselves but also the underlying infrastructure upon which the VNF depends. It will also be required to maintain consistent network deployments where virtual and physical networks coexist.

It is no surprise that the entire telecommunications industry is actively addressing these challenges. It is clear however that NFV presents new and exciting opportunities arising from more dynamic and agile network and service infrastructures than seen heretofore.

REFERENCES

- [1] 3GPP, "Network Architecture," TS 23.002, 3GPP, June 2014.
- [2] ETSI, "Network Functions Virtualisation (NFV): Architectural Framework," ETSI, Tech. Rep. NFV 002, October 2013.

- [3] S. Landstrom, A. Furuskar, K. Johansson, L. Falconetti, and F. Kronstedt, "Heterogeneous networks – increasing cellular capacity," *Ericsson Review*, no. 2, pp. 34–38, 2011. [Online]. Available: http://www.ericsson.com/res/thecompany/docs/publications/ericsson_review/2011/Ericsson-Review-2011-2.pdf
- [4] J. A. Wickboldt, L. Z. Granville, F. Schneider, D. Dudkowski, and M. Brunner, "Rethinking cloud platforms: Network-aware flexible resource allocation in iaas clouds," in *Integrated Network Management (IM), 2013 IFIP/IEEE International Symposium on*, May 2013.
- [5] B. B. Nandi, A. Banerjee, S. C. Ghosh, and N. Banerjee, "Dynamic sla based elastic cloud service management: A saas perspective," in *Integrated Network Management (IM), 2013 IFIP/IEEE International Symposium on*, May 2013.
- [6] L. Richardson and S. Ruby, *RESTful Web Services*, 1st ed. O'Reilly Media, Inc., 2007.
- [7] IETF, "An architecture for describing simple network management protocol (snmp) management frameworks," RFC 3411, IETF, Tech. Rep. RFC 3411, Dec. 2002. [Online]. Available: <http://tools.ietf.org/html/rfc3411>
- [8] —, "Netconf configuration protocol," RFC 4741 (Proposed Standard), IETF, Tech. Rep. RFC 4741, Dec. 2006. [Online]. Available: <http://tools.ietf.org/html/rfc4741>
- [9] P. Gustas, P. Magnusson, J. Oom, and N. Storm, "Real-time performance monitoring and optimization of cellular systems," *Ericsson Review*, no. 1, pp. 4–13, January 2002. [Online]. Available: http://www1.ericsson.com/res/thecompany/docs/publications/ericsson_review/2002/2002011.pdf
- [10] AT&T. (2013) AT&T vision alignment challenge technology survey - AT&T domain 2.0 vision white paper. Online Document. [Online]. Available: https://www.att.com/Common/about_us/pdf/AT%20Domain%202.0%20Vision%20White%20Paper.pdf
- [11] F. Oliveira. (2014) Verizon experience building openstack clouds with red hat. Online Video. Verizon. [Online]. Available: <http://superuser.openstack.org/articles/summit-session-verizon-s-nfv-strategy-lab-specs>
- [12] Ericsson. (2014, February) The Real-Time Cloud. Ericsson. [Online]. Available: <http://www.ericsson.com/res/docs/whitepapers/wp-sdn-and-cloud.pdf>
- [13] G. Csatai and T. Laszlo, "NSN Mobile Core Network Elements in Cloud," in *Communications Workshops (ICC), 2013 IEEE International Conference on*, June 2013, pp. 251–255.
- [14] NEC. (2014) NEC Virtualized Evolved Packet Core – vEPC. NEC. [Online]. Available: http://www.nec.com/en/global/solutions/nsp/nfv/doc/vEPC_WP.pdf
- [15] Huawei. (2014, February) Network Function Virtualization: Rebuilding Network Functions and Open Architectures. Huawei. [Online]. Available: http://www.huawei.com/mwc2014/en/articles/hw-u_319944.htm
- [16] ZTE. (2014, February) ZTE Launches Cloud UniCore Solution. ZTE. [Online]. Available: http://www.zte.com.cn/en/press_center/news/201402/t20140224_418322.html
- [17] ALU. (2014) The Journey to Packet Core Virtualization. ALU. [Online]. Available: <http://resources.alcatel-lucent.com/asset/174234>
- [18] C. Mobile. (2011, October) C-RAN: The Road Towards Green RAN. China Mobile. [Online]. Available: http://labs.chinamobile.com/cran/wp-content/uploads/CRAN_white_paper_v2_5_EN.pdf
- [19] NGMN, "Suggestions on Potential Solutions to C-RAN," Deliverable, NGMN, Tech. Rep., January 2013. [Online]. Available: http://www.ngmn.org/uploads/media/NGMN_CRAN_Suggestions_on_Potential_Solutions_to_CRAN.pdf
- [20] MCN. (2014, July) Mobile Cloud Networking EU FP7 Project. MCN. [Online]. Available: <http://www.mobile-cloud-networking.eu>
- [21] iJOIN. (2014, July) Interworking and JOINt Design of an Open Access and Backhaul Network Architecture for Small Cells based on Cloud Networks EU FP7 Project. iJoin. [Online]. Available: <http://www.ict-ijoin.eu>
- [22] D. Sabella, P. Rost, Y. Sheng, E. Pateromichelakis, U. Salim, P. Guitton-Ouhamou, M. Di Girolamo, and G. Giuliani, "RAN as a Service: Challenges of Designing a Flexible RAN Architecture in a Cloud-Based Heterogeneous Mobile Network," in *Future Network and Mobile Summit (FutureNetworkSummit), 2013, July 2013*, pp. 1–8.