Mobile Netw Appl (2018) 23:521–528
DOI 10.1007/s11036-017-0905-y

SDN and NFV as Enabler for the Distributed Network Cloud

Marco Hoffmann [1] · Michael Jarschel [1] · Rastin Pries [1] · Peter Schneider [1] · Admela Jukan [2] · Wolfgang Bziuk [2] · Steffen Gebert [3] · Thomas Zinner [3] · Phuoc Tran-Gia [3]

Published online: 22 September 2017

Abstract SDN and NFV gained significant momentum within the last years. Although widely used in research labs and cloud environments, SDN has not yet been deployed in mobile telecommunication networks. In this paper, we focus on use cases driving mobile network evolution towards cost-efficient IT-based solutions using standardized hardware and software-based concepts, such as SDN and NFV. Two SDN use cases for mobile networks are described. One deals with the disaggregation of mobile network gateways and the other with SDN-enabled security concepts and applications for mobile networks. Based on the SDN use cases, the paper highlights open issues and challenges for integrating IT concepts in future telecommunication networks.

Keywords Network cloud · SDN · NFV · Telco-centric IT infrastructure

1 Introduction

In the last years, we have witnessed tremendous research efforts in the broad area of Software Defined Networking (SDN) [1] Despite its popularity in the research community, SDN deployments in operational service provider networks have not been realized for several reasons, including interoperability with legacy systems as well as critical aspects of security and system scalability [2]. However, the increasing cost pressure has led to a "softwarization" and "cloudification" of network functions in general [3]. Network operators already deploy entities in clouds, such as the Mobility Management Entity (MME), the IP Multimedia Subsystem (IMS), and the Home Subscriber Server (HSS) [4]. Upcoming services such as autonomous driving, eHealth, massive IoT, and industrial communication especially in 5G might further increase the cost pressure and thus, the demand for outsourcing network functions to the cloud [5]. This evolution has positioned SDN not only in the role of an enabler for Network Functions Virtualization (NFV) [6], but also more broadly, for network clouds.

Two key features led to the popularity of SDN and NFV in the first place: flexibility to run on commodity hardware, and easy and extensible implementations of network functions through software-based solutions that indeed can save operational and management costs due to automation. Just as SDN

✉ Marco Hoffmann
marco.hoffmann@nokia-bell-labs.com

Michael Jarschel
michael.jarschel@nokia-bell-labs.com

Rastin Pries
rastin.pries@nokia-bell-labs.com

Peter Schneider
peter.schneider@nokia-bell-labs.com

Admela Jukan
a.jukan@tu-bs.de

Wolfgang Bziuk
w.bziuk@tu-bs.de

Steffen Gebert
steffen.gebert@informatik.uni-wuerzburg.de

Thomas Zinner
zinner@informatik.uni-wuerzburg.de

Phuoc Tran-Gia
trangia@informatik.uni-wuerzburg.de

[1] Nokia Bell Labs, Munich, Germany

[2] TU Braunschweig, Braunschweig, Germany

[3] University of Würzburg, Würzburg, Germany

has brought the separation of packet forwarding logic from the routing hardware, NFV goes even further and utilizes software techniques to remove network functions from dedicated network hardware by migrating them to the cloud. The implementation of network functions like firewalls, performance monitors, and load balancers purely in software, as opposed to integration into specialized hardware middle boxes, enables network providers to utilize commodity hardware and cloud computing. This way, the network functions can be implemented in a more cost-effective, flexible, and vendor-independent fashion. Leveraging both concepts, SDN to separate routing and forwarding in the core network, and NFV to implement network functions as cloud solutions, network operators can achieve new levels of flexibility, system agility as well as elasticity.

General challenges and opportunities of NFV are shown in [7, 8]. They both describe the state-of-the art in network virtualization and show different use cases for applying NFV. Use cases addressed are the virtualization of the Radio Access Networks (RAN) and the virtualization of the evolved packet core. The papers however only slightly touch the biggest challenge of NFV and SDN, namely the I/O performance problems when using commercial off-the-shelf (COTS) hardware. In addition, practical network design considerations are missing.

In this paper, we discuss these practical network design considerations for telecommunication networks embracing cloud computing concepts and solutions in the service chain. We present two examples and illustrate the evolution and challenges of softwarization and cloudification in 5G networks. The first example focuses on the implementation and performance of disaggregated mobile network gateways and the second addresses SDN-based security enforcement in mobile networks. Afterwards, we give a brief overview of the performance challenges of using COTS hardware and provide a classification of implementation approaches for network functions with respect to virtualization support, performance, scalability and complexity.

The rest of this paper is structured as follows. In the next section, we describe the need for softwarization from a mobile telecommunication network perspective. We discuss the various increasing service demands and the need for the network operators to shift network parts to commodity hardware for cost reduction and to reduce the management and operation overhead. This is followed by a detailed discussion of the two use cases. At the end of this paper, we present our vision on network evolution including the hardware challenges. Finally, we draw conclusions.

2 Softwarization

Mobile computing technologies as well as the development of low-cost sensors and big-data data processing in Machine-to-Machine (M2 M) communication have facilitated a myriad of the so-called Internet-of-Things (IoT) services in a broad spectrum of industries. From the health care sector, Smart-X sectors (Transportation, City, Home) to environment monitoring, autonomous driving and linked manufacturing, communication and networking technologies are evolving together with their IT counterparts, in both synergetic and competitive fashion. From carrier Data Centers (DCs), i.e., the Network Cloud, to the public cloud infrastructure for processing and storage capacities, and the so-called *fog computing* through pervasive mobile connectivity of end-devices, the requirements for capacity and a sheer number of end points have increased to the unprecedented levels [9]. At the same time, latency remains the challenge, not only due to the ever more stringent requirements for real-time services, but also for the communication latency between the clouds and the users. This is putting a great pressure on service providers to reduce any excess costs and overheads, as more and more demanding services need to be provided over the same or a gradually improved infrastructure.

The reduction of costs is the main motivation for service providers to consider new technologies. Software-based technologies with a high degree of automation have proven to streamline the operation of services in the IT domain. This resulted in reductions of service launch times, required administration personnel, and number of maintained configurations in the order of magnitudes. SDN and NFV are no exception. Just as SDN has brought the "softwarization" through separation of packet forwarding logic from the routing hardware, NFV concept has gone even further by migrating functions directly to the telco cloud. Both, SDN and NFV are thus effectively the enablers of the so-called network cloudification. Figure 1 illustrates the evolution of telecommunication networks towards network cloudification through SDN and NFV concepts and finally to a conglomerate of centralized, large clouds as well as distributed clouds. As the figure shows, there are two possible ways from the integrated hardware as starting point towards cloudification. The first focusing on decoupling the layers through SDN and the second moving existing solutions into the cloud and making them scale horizontally by applying NFV. The paths are not isolated from each other and switching the focus is possible at any stage. In the following, we will present two use cases for the first direction, a decoupled LTE mobile network gateway and their enhancement with SDN-based security features.

3 SDN use cases for mobile networks

3.1 Use case of LTE mobile network gateways

In contrast to central functions of mobile networks, most of which are nowadays already running virtualized on

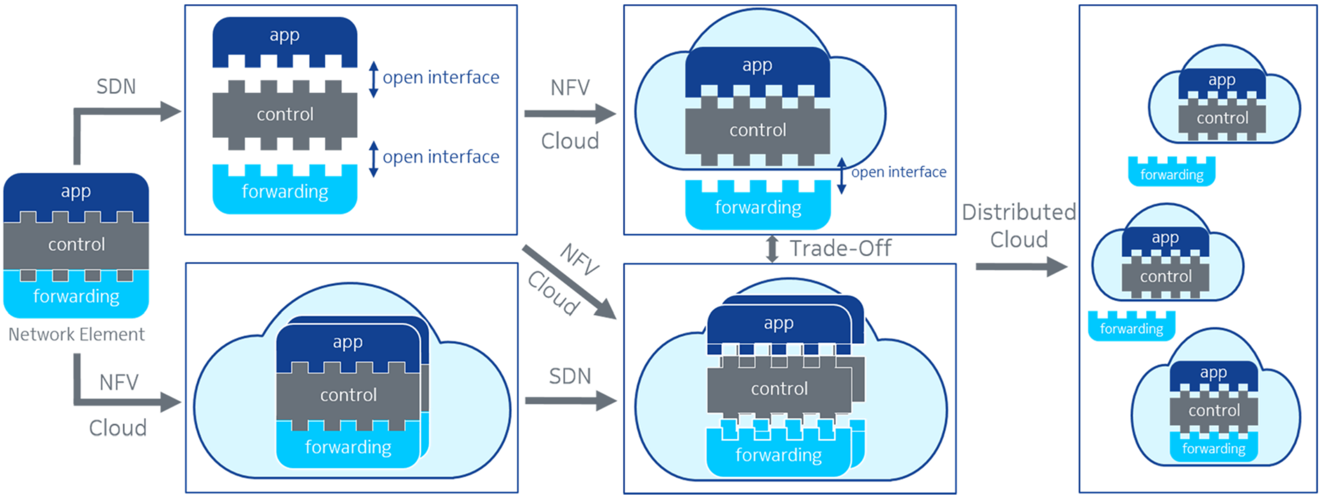

Fig. 1 Evolution from decoupling over cloudification to distributed clouds

commodity hardware in the cloud, the user plane traffic still requires handling by specialized hardware to ensure high network performance. The introduction of SDN and NFV presents the opportunity to relocate the user plane workloads towards standard equipment gradually.

LTE mobile network gateways [10], i.e., the Serving Gateway (SGW) and the Packet Gateway (PGW), are examples of SDN and NFV applicability in mobile networks. The fundamental purpose of the gateways is to receive and forward traffic from the user equipment to the Internet and vice versa. They further serve as “mobility anchor” for the user plane, allowing users to roam between base stations while maintaining uninterrupted connectivity. Additional functionality includes providing billing data to the operator’s charging systems. Today, every operator has deployed typically only a relatively small number of optimized high-performance hardware gateways in strategic locations of their network. While this is sufficient to handle today’s user traffic, the expected exponential increase in mobile network bandwidth usage, different mobility behaviors, and future latency sensitive applications will require more flexibility to deliver efficient services.

Therefore, instead of forwarding the traffic to a few fixed anchor points in the network, it will be beneficial to further distribute the traffic over multiple points-of-presence and provide additional capacity when needed. SDN enables the separation of the mobile gateways’ control and user plane. This way, a cloud-hosted Virtual Network Function (VNF) can serve as scalable control plane. At the same time, an SDN switch can function as user plane. The location and number of SDN switches serving in the field is adaptable as it depends only on the control connection towards the cloud. Figure 2 shows such an SDN/NFV function chain for the mobile gateways as described in [11]. The figure illustrates the SGW and PGW as logical entities within the dashed grey boxes. Each gateway consists of a programmable switch for forwarding, i.e., user plane, and a control connection to an SDN controller in the cloud. In our realization of this use case, the control connection uses OpenFlow 1.3 with TLS encryption as southbound protocol. The protocol is further augmented with a vendor extension to permit the handling of packets using the GPRS Tunneling Protocol (GTP). GTP serves to distinguish between User Equipment (UE) bearers on the links between eNodeB and SGW as well as from SGW to PGW. The extension leverages the OpenFlow eXtensible Match (OXM) mechanism introduced in OpenFlow 1.2. The data path implementation, i.e., the switch, is realized by the eXtensible DataPath Deamon (xdpd) [12]. It provides a pipeline of individual OpenFlow flow tables for flow installation. The SDN controller consists of a hierarchy of reusable proxy OpenFlow controllers acting as data paths to the north and controllers to the south. Each controller layer uses a dedicated set of flow tables from the data path to implement the functionality of the corresponding network layer protocol, i.e. Ethernet, IP, and GTP. Up to this point, the two gateway implementations are identical, which means they can be instantiated as VNFs in the cloud from the same basic template. The application running on top of the controller instances provides the actual functional logic. In case of the SGW, the app is denoted as SGW-C in the figure, whereas it is denoted as PGW-C for the PGW. The applications communicate with the highest controller in the hierarchy (GTP) using a RESTful interface over a TLS encrypted connection.

This architecture offers the possibility to place a controller instance close to the user plane device, which enables fast status updates and responses at the lower layer, whereas the central application functionality can be hosted in a central data center. As can be seen in Fig. 2, the PGW SDN controller also runs the MG-C application (Mobile Guard Control Application), which is related to the second use case shown in the figure and discussed in the following.

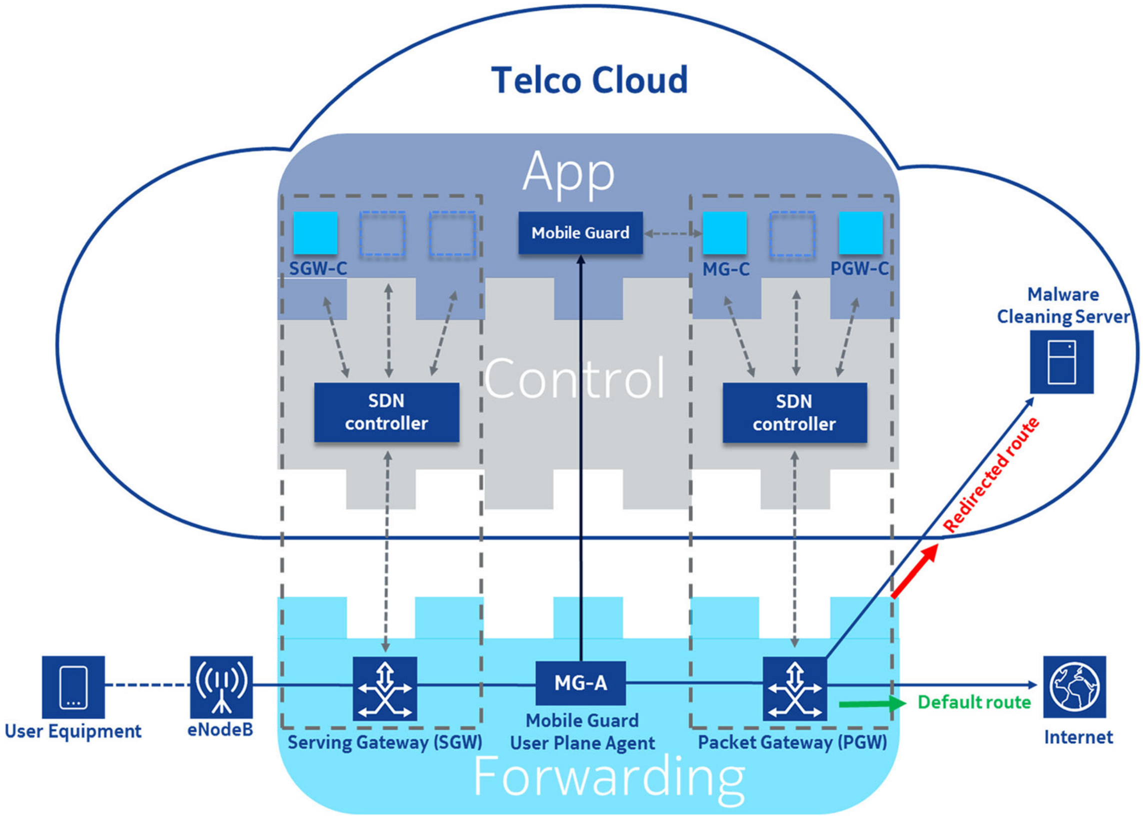

Fig. 2 Use Case LTE mobile gateway SDN/NFV function chain (SGW: Service Gateway, MG-C: Mobile Guard Control, PGW: Packet Gateway)

3.2 Use case with SDN-enhanced security features

We extend the previous use case of an SDN/NFV function chain by incorporating a security VNF into the function chain between the two mobile gateways as shown in [13].

The scenario we look at is a user of a mobile operator's network that uses a tablet device infected with malware. The malware hides inside a mobile app. The security VNF, i.e., Nokia's Mobile Guard security appliance [13], detects the infection using its user plane agent located between SGW and PGW. Typically, at this point the Mobile Guard would issue a warning to the user via SMS or email. However, this would allow the malware to have Internet access until the user reacts to the notification and removes it from the device, which bears the risk of it spreading further or in the worst case disrupting network operation. Therefore, we leverage SDN's northbound API to notify the MG-C app running on the PGW SDN controller of the attack and redirect all traffic coming from the device towards a malware cleaning appliance in the operator cloud. The appliance then offers an app for cleaning the tablet of the user. After the user has downloaded the app and used it to remove the malware from the tablet, the Mobile Guard will instruct the MG-C on the PGW SDN controller to forward the tablet's traffic normally again. The PGW-C application runs in parallel to the MG-C application on the controller. The applications are isolated from each other. Thus, the added Mobile Guard functionality does not interfere with the normal operation of the mobile network. This example highlights one of the major advantages of the software-based approach. The extension or alteration of the gateway functionality is possible at run time while the operation remains virtually unaffected.

4 Open issues and challenges

Automatically accommodating a significant number of diverse and demanding services is a difficult task. Solving it requires a more granular and customizable network architecture and at the same time maintaining the benefits and simplicity of commodity hardware and IT processes.

Therefore, we envision the network cloud, which is enabling a converged network architecture, a process based on three key ideas:

First, breaking SDN out of the overlay and moving it into the network as a primitive. Second, applying NFV principles beyond the data center, and third, bringing the service to the customer instead of the data to the cloud.

Realizing this use case requires an adaptive cloud system that does not follow the principle of "one-size-fits-all", but rather tailors to specific network tasks in terms of functionality and locality. Figure 3 gives a simplified overview of such a cloud system. It consists of three tiers each containing a different type of data center.

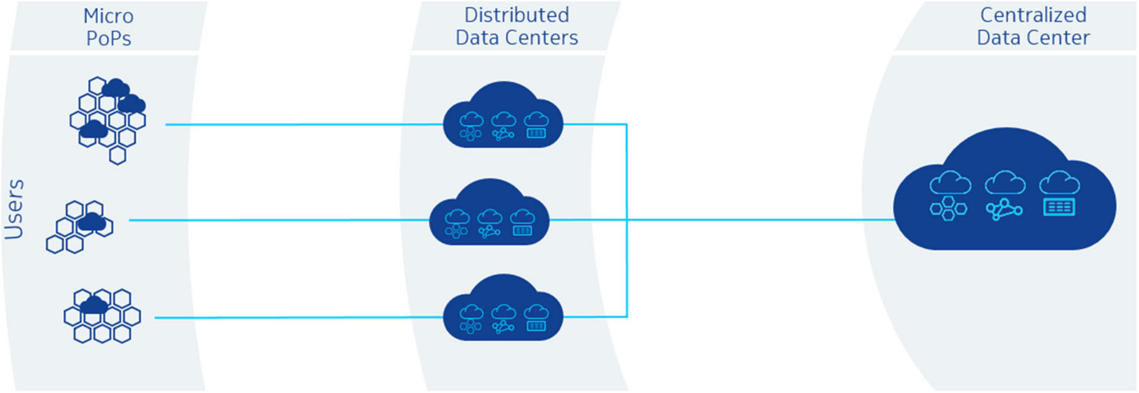

Fig. 3 A Future use case of three tier distributed cloud

Starting from the left, we have the Micro Points of Presence (PoPs). These are small computing entities dedicated to running virtual network functions as close as possible to the deployed equipment in the field ensuring low processing latency. These entities are part of the overarching cloud management system, but are not used for general computing purposes. Instead, they run a platform service for network functions, incorporating programmable networking hardware to ensure maximum performance. The second tier consists of small-to-medium size data centers. These data centers are located at strategic locations within the network, e.g., at peering points. Apart from virtual network functions, some general-purpose computing functionalities can be located here, e.g., CDN content servers. While the former run on the same VNF platform service as in the Micro PoPs, the latter run on a complementary platform modelled after current IT data centers. Finally, we have a few centralized IT data centers located at central hubs of the operator. These house essential IT and management functions that run in the cloud already today, but are not exposed to the actual user traffic, e.g., a user database.

The key is that the network cloud management system handles both the IT platform and the NFV platform identically, but is aware of the different purposes and locations, i.e., deploying a low-latency network function within a Micro PoP instead of a centralized data center. The underlying hardware can differ, depending on the requirements of the use cases, which is also the biggest challenge of NFV and SDN. On one hand, as much COTS hardware as possible should be used to reduce the costs and to minimize the administration overhead. On the other hand, some functionality cannot be handled using normal COTS hardware such as complex ciphering mechanisms or low-latency demanding Cloud RAN functions. In Table 1, we show different requirements of some Cloud RAN network functions. It is obvious that these delay, processing and memory requirements cannot be handled using normal COTS hardware. The challenge is to evaluate, how far we can get with COTS hardware and which form of specialized hardware can be used. For some parts, it might already be sufficient to use FPGA boards inside standardized hardware. For other functions, there might be the need for specialized hardware. In addition, the network function placement algorithms must be able to not only place different functions or function chains at a special location (MicroPoP), but also to place them on a special server or hardware.

4.1 Automated network operation

The automation of network operation through softwarization is a key trend within the ICT industry. The need of network operators to reduce costs while dealing with increasing demand drives this development. SDN and NFV have emerged as key enablers for this transformation process. However, the current model for the softwarization of networks mostly relies on adapting the concept of cloud computing from IT players like Amazon or Google. This concept uses multiple strategically placed data centers to consolidate operations and management for online services. Currently, this is a promising way to achieve a high degree of automation. Yet, given the challenges discussed in the previous section and the increasing demand through advanced services, this approach will likely only be able to serve as a stepping stone towards a fully-fledged Telco Cloud that encompasses not only IT, but also draws on concepts from the Telco domain.

This is one of the reasons why operators are still reluctant to large-scale SDN and NFV deployments. Further reasons include that not all open issues and challenges have been solved yet. While development speed of software is faster than development of hardware, research work is still away from reaching the scalability of carefully crafted ASIC-based network elements. Further, as with all architectural changes, security considerations hinder early adaptors. Finally, a seamless migration path is needed, as no operator is willing to replace a complete network infrastructure all at once.

4.2 Performance and scalability challenges

The performance of virtualized network functions, i.e., the number of processed packets per second, the packet

Table 1 Requirements of different Cloud RAN network functions

Network function	Timing	Processing	Memory	Scaling driver
PDCP	< 1 ms per PDU	high (en/de-cryption)	high (u-plane buffer)	throughput
RRC	< 5 ms per message	medium	medium (per UE and DRB state)	connections
RLC	< 1 ms per PDU	medium	medium (re-transmission buffer for ARQ, uplink Tx buffer in case of fronthaul split)	throughput
MAC (protocol)	TTI (< 0.2 ms)	medium	medium (frame construction buffer)	throughput
MAC (scheduler)	TTI (< 0.2 ms)	high	medium (depends on state/implementation)	throughput

(Packet Data Convergence Protocol (PDCP); Radio Resource Control (RRC); Radio Link Control (RLC); Medium Access Control (MAC); Transmission Time Interval (TTI); Packet Data Unit (PDU); User Equipment (UE); Data Radio Bearer (DRB); Automatic Repeat Request (ARQ))

processing delay and delay variation, or the number of supported connections, can degrade significantly compared to the performance of highly optimized middle boxes. For our first use case that is the number of GTP tunnels (bearers) that can be handled by the SDN SGW and PGW as well as the forwarding delay imposed by the retagging and untagging of GTP packets. This, however, can be balanced by reduced expenditures and the possibility to scale out the virtualized network function tailored to the actual demand.

The first contributor to the variation of execution speed of a VNF is the operating system's separation of memory into kernel and user space. To provide an application running in user space, data received at the network card has to be copied between kernel and user space. This not only limits the throughput that a server can achieve, but also increases and results in variations of the processing delay.

Further, multi-tasking, the execution of multiple processes inside an operating system, additionally accounts for possible interference. Besides the available CPU time, also other resources including CPU caches and network I/O are shared between multiple processes. Additional sharing happens, when virtual machines of multiple tenants run on the same server hardware. In such scenarios, prediction of resource availability becomes even harder.

A multitude of optimization techniques are available to improve the speed of packet processing in software. Among these are optimized software components like kernel modules, zero-copy techniques, and kernel bypasses, or hardware-based acceleration mechanisms.

Table 2 gives an overview of different solutions. A user space application provides the lowest performance. By running in a virtual machine, however, it can easily be scaled horizontally and vertically, and does not require any specialized features. A dedicated middle box provides high performance through its ASIC implementation at high cost and limited scalability. Increased demands typically lead to high acquisition costs for new devices, and vertical scalability is limited by the availability of expansion slots.

The other approaches are located between these extremes. The FPGA-based approach does not support virtual machines and requires the use of a hardware description language. However, it allows appending and programming of FPGA-NICs on demand. Software-based optimization techniques, like zero-copy or kernel bypass, often lead to complex implementations and often require support by the operating system. Intel DPDK requires a reimplementation of the specific function against the DPDK-API. Further, it depends on the underlying hardware possibly limiting the available resources.

A performance comparison between a DPDK-enabled and a pure user-space implementation of an SGW is presented in [10].The results indicate a performance boost up to a nine-fold number of packets per second for the DPDK-enabled implementation compared to the standard user-space implementation. To support DPDK, the SGW has to make use of DPDK APIs, potentially increasing the implementation's complexity compared to a traditional user-space implementation. Further, specific hardware components must be available to enable DPDK packet processing. This may reduce horizontal scalability, since hardware constraints must be considered.

For the time being, it is not clear which implementation solution provides comparable performance characteristics to the middle box approach while reducing the costs significantly. Additionally, it remains an open question how larger processing delays influence service chains. User-space implementations increase the processing delay by several orders of magnitude [10] and thus might be less appropriate for such a scenario than hardware or software accelerated solutions, which have other drawbacks like limited scalability, higher complexity, or missing virtualization support.

4.3 Security considerations

SDN brings new security challenges to networks. First, the separation of forwarding and control plane in SDN introduces an additional interface that increases the attack surface of the

Table 2 Classification of implementation approaches for Network Functions with respect to virtualization support, performance, scalability and complexity

	Virtualization Support	Performance	Vertical Scalability	Horizontal Scalability	Complexity
User-space	++	−	++	++	++
Zero-copy / kernel bypass	+	0	+	+	+
DPDK	+	0	+	+	+
FPGA	−	+	−	0	−
Middle box	−	++	-- (0)*	−	−

The grade of support of each feature ranges from very bad (−-) to very good (++) with (0) indicating borderline

*applies if expansion slots are available

overall system. It could allow attacks on the integrity and confidentiality of the controller-switch communication, DoS attacks, or attacks aiming at gaining some control over switches and controllers by exploiting vulnerabilities in the protocol software or the interface configuration. However, securing such an interface is a well-known task and suitable means are available, such as usage of IPsec or TLS.

Second, SDN introduces the northbound interface, where applications can access SDN controllers to control the network. The concept of possibly several applications, including third party applications, executing control over a network raises many security issues. They include authentication of applications, authorization of requests, resolving conflicting requests, and prevent malicious applications from compromising a controller at this interface and subsequently exhibit unauthorized control over network resources. We assume these issues are resolvable, but solutions must be carefully designed and implemented.

On the positive side, unified, centralized control has the potential to make networks more secure. Controlled suitably, all the SDN switches in a network can contribute to the execution of network security functions such as traffic filtering – an approach that may be very valuable against distributed denial of service attacks, one of the most dangerous attack pattern in today's networks. SDN provides new ways to implement security solutions in networks. It may enable more flexible and efficient deployment of security solutions if those solutions can be implemented as applications running on controllers without relying on traditional security devices.

The separation also allows implementing controllers in NFV environments. Besides the general advantages of the NFV approach, there is also a possible security gain: NFV environments can help to overcome DoS attacks at least temporarily by dynamically allocating additional resources to controllers.

However, SDN controllers in NFV environments will be exposed to significant new threats applicable to all virtual network functions in the environment. The NFV infrastructure may fail to provide a 100% secure isolation between different VNFs. Malicious VNFs may exploit this to attack other VNFs, for example by reading or even modifying the target VNF's memory.

Securing an NFV environment is a complex task and largely out of scope of these brief security considerations. A robust, sound implementation of the "cloud stack" minimizing exploitable vulnerabilities is clearly the basic security requirement. Moreover, the environment must offer suitable security features, including means to enforce even physical isolation of certain critical functions. Finally, a high degree of automation must be provided for the security management and orchestration of complex VNF setups.

4.4 Network technology migration

The migration to any new network technology is typically a gradual transition over time. In case of SDN and NFV, not only technical novelties but also economic factors decide the pace and the extent of the network operator's migration to these two emergent technologies. In general, SDN and NFV will need to co-exist and be interoperable with the existing network systems. In fact, ISPs are increasingly considering the concept of hybrid SDN control plane, referring to an Internet control plane architecture, where the new and old control plane paradigms co-exist – the centralized SDN controller and the distributed OSPF/IS-IS routing protocol. This way, a hybrid control plane architecture can use the legacy routing protocol for packet forwarding, while SDN controllers can inject high priority rules on top [14].

It is expected that the migration to SDN is generally less of a challenge in the scenarios where the administrative domains are highly controlled. That is likely the reason why we see most SDN deployments in data centers today. Furthermore, addressing the challenge of migration to SDN in combination with NFV would likely be more beneficial than studying migration to each technology in isolation [15]. NFV alone requires other migration challenges to be addressed, such as the orchestration strategy between the legacy services and NFV, as well as mobility and portability between different sites and vendors [16].

Finally, the challenge of network technology migration is also the focus of various industry-led initiatives, such as OpenDayLight, attempting to integrate SDN and non-SDN technologies (e.g., PCEP, SNMP) in the same framework. Some OpenDayLight controller implementations, such as ClosedFlow [17], present a system where SDN can control the existing proprietary hardware thus mimicking the fine grain control as it is typically enabled by OpenFlow.

5 Conclusion

In this paper, we presented three SDN use cases in the context of various aspects of network evolution, especially for mobile and the future 5G networks, towards cost efficient IT-based solutions using standardized hardware and software-based concepts, such as SDN and NFV. The goal was to highlight open issues and challenges for integrating IT concepts in telecommunication networks. To take full advantage of a future network cloud environment, we conclude that it is clearly not sufficient to simply port the current SDN and NFV functions and applications for running on “bare metal.” Network service and provisioning automation through softwarization and cloudification is generally a grand challenge. Solving it requires a more granular and customizable network architecture than we have today, which paradoxically is also expected to inherit the cost benefits and simplicity of commodity hardware and IT processes. Realizing the new ideas in this space needs to be gradual, but also require an agile and adaptive network cloud that does not follow the principle of “one-size-fits-all”, but rather tailors to specific network services in terms of functionality and locality.

Acknowledgements This work has been carried out in the framework of the CELTIC EUREKA project SENDATE-PLANETS and it is partly funded by the Federal Ministry of Education and Research (BMBF).

References

1. Kreutz D, Ramos F, Verissimo PE, Rothenberg CE, Azodolmolky S, Uhlig S (2015) Software-defined networking: a comprehensive survey. Proc IEEE 103:14–76
2. Sezer S, Scott-Hayward S, Chouhan PK, Fraser B, Lake D, Finnegan J, Viljoen N, Miller M, Rao N (2013) Are we ready for SDN? Implementation challenges for software-defined networks. IEEE Commun Mag 51:36–43
3. An X, Kiess W, Varga J, Prade J, Morper H-J, Hoffmann K (2016) SDN-based vs. software-only EPC gateways: a cost analysis. In NetSoft Conference and Workshops (NetSoft)
4. Alcatel-Lucent Technology Whitepaper (2014) The journey to packet core virtualization. Available: http://www.tmcnet.com/redir/?u=1011494. Accessed 20 September 2017
5. Weldon MK (2016) The future X network: a bell labs perspective. CRC Press, Boca Raton
6. ETSI NFV (2012) Introductory white paper. Available: https://portal.etsi.org/nfv/nfv_white_paper.pdf. Accessed 20 September 2017
7. Han B, Vijay G, Ji L, Lee S (2015) Network functions virtualization: challenges and opportunities for innovations. IEEE Commun Mag 53:90–97
8. Mijumbi R, Serrat J, Gorricho J-L, Bouten N, de Turk F, Boutaba R (2016) Network function virtualization: state-of-the-art and research challenges. IEEE Communications Surveys & Tutorials 18(1):236–262
9. Masip X, Marín E, Jukan A, Ren GJ, Tashakor G (2016) Foggy clouds and cloudy fogs: a real need for coordinated management of fog-to-cloud computing systems. IEEEWirel Commun 23(5): 120–128.
10. Lange S, Nguyen-Ngoc A, Gebert S, Zinner T, Koepsel A, Sune M, Raumer D, Gallenmüller S, Carle G, Tran-Gia P (2015) Performance benchmarking of a software-based LTE SGW. In 11th International Conference on Network and Service Management (CNSM)
11. Gebert S, Hock D, Zinner T, Tran-Gia P, Hoffmann M, Jarschel M, Schmidt E-D, Braun R-P, Banse C, Köpsel A (2014) Demonstrating the optimal placement of virtualized cellular network functions in case of large crowd events. In ACM SIGCOMM Computer Communication Review
12. BISDN (2016) eXtensible DataPath Deamon (xdpd). Available: https://www.Bisdn.De/?page_id=15. Accessed 21 June 2016
13. Khatri V, Abendroth J (2015) Mobile guard demo: network based malware detection. In IEEE Trustcom/BigDataSE/ISPA
14. Caria M, Das T, Jukan A, Hoffmann M (2015) Divide and conquer: partitioning OSPF networks with SDN. In IFIP/IEEE Integrated Network Management Symposium
15. Das T, Drogon M, Jukan A, Hoffmann M (2015) Study of network migration to new technologies using agent-based modeling techniques. J Netw Syst Manag 23(4):920–949
16. López LIB, Valdivieso AL, Villalba LJG, López D (2015) Trends on virtualisation with software defined networking and network function virtualisation. IET Networks 4:255–263
17. Hand R, Keller E (2014) ClosedFlow: OpenFlow-like control over proprietary devices. In 3rd Workshop on Hot topics in Software Defined Networking