

Moving the Network to the Cloud: Multi-Tenant and Multi-Service Cloud Central Office

Marco Ruffini

The University of Dublin, Trinity College, Ireland. marco.ruffini@tcd.ie

Abstract *This tutorial provides an overview of the various frameworks and architectures outlining current network disaggregation trends that are leading to the virtualisation/cloudification of central offices. The discussion will include the optical layer disaggregation and provide an overview on future challenges.*

Introduction

The Software Defined Networking (SDN) and Network Function Virtualisation (NFV) concepts have recently changed the way we operate networks. While the concept of decoupling hardware forwarding and software control plane operations had been discussed since the beginning of the previous decade (the reader should refer to¹ for a comprehensive survey), the start of the SDN revolution in telecommunications networks typically dates back to the introduction of OpenFlow². By providing an API that could interface with off-the-shelf hardware switches, OpenFlow enabled the possibility to develop and test new ideas over real networks, paving the way to many research projects across the globe, for the development of new network protocols and functionalities. Within only a few years SDN managed to grasp the attention of the data centre and telecomms industry.

In the mean time, in parallel with the control plane programmability offered by SDN, the concept of data plane programmability and virtualisation also continued to develop, resulting in the emergence of NFV frameworks. While the SDN and NFV concepts are in principle independent, they are typically considered highly synergistic. Together, by decoupling hardware and software operations, and virtualising their components, they have opened the way to new frontiers, to reduce network costs and improve network usability and efficiency. While the early adopters were data centre operators, recently their applicability has moved towards public telecommunications networks, with projects like the Central Office Rearchitected as a Data centre (CORD)³, that have pioneered its use in access and metro networks. This idea has quickly attracted the interest of the networking industry, gaining in only a few years the support of several operators across the world, many of which have started carrying

out network trials. This has also led to novel standardisation activities, with the BroadBand Forum (BBF), for example, leading the Cloud Central Office (Cloud-CO) working task⁴.

The concept of central office virtualisation brings together different network technologies, providing functional convergence for mobile and optical access networks, and paving the way for its integration with disaggregated ROADM networks at the metro level. By merging mobile, residential and enterprise services into a common framework, built around a data centre architectural core, the cloud CO can achieve significant capital and operational cost savings. However many challenges remain in order to guarantee the quality of service required to run upcoming 5G applications, while multiplexing network and processing resources across multiple operators and services.

In the reminder of the paper we will borrow the term Cloud-CO to refer not just to the specific BBF architecture and working task, but to mean the general concept of virtualisation of the central office. In the next section this tutorial paper will introduce a number of different development frameworks that implement the Cloud-CO concept. We then extend the disaggregation to the optical layer, briefly mentioning some of its pros and cons and their importance towards a fully virtualised network. Finally we explore some of the outstanding challenges we believe should be addressed in the near future.

Cloud Central Office Architectures

A Cloud-CO is a framework for bringing NFV into a telecommunications central office, where functions that typically run on dedicated hardware are moved to software frameworks running on commodity hardware (e.g., servers). This moves the CO architecture towards that of a data centre. Its

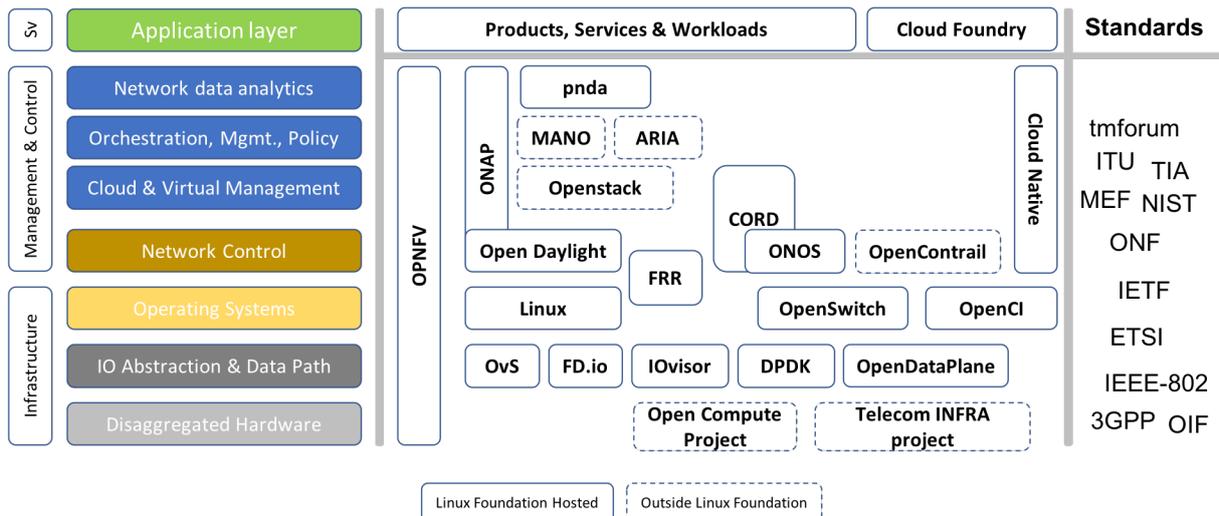


Fig. 1: Classification of NFV-related development frameworks⁸

implementation relies on the development of several software components that closely interoperate to deliver an end-to-end solution. There are today a number of software implementations for NFV/SDN, the most popular being ONAP⁵, OPNFV⁷, MANO⁶, CORD³, only to mention a few. The diagram in Fig. 1 (re-drawn from source⁸), is an attempt to map their functionality with respect to a management and control plane stack. The main observation to make is that although many of these are hosted by the Linux foundation, they originated as independent projects, from different organisations, following different standardisation efforts and thus there is substantial overlap in functionality across them. However, in some cases, there have been liaison efforts, for example between ONAP and OPNFV.

An interesting aspect of NFV and SDN is that they enable convergence across multiple network domains in an unprecedented fashion. A representative use case is the integration of fixed and mobile networks, where for example the scheduler of a Passive Optical Network (PON) and that of a mobile BaseBand Unit (BBU) can be synchronised to reduce the transmission latency over the PON in a cloud RAN implementation^{9,10}.

In summary, all frameworks mentioned above are currently under development, and it is difficult to anticipate how their interoperation will evolve and whether some will dominate over others or else they will keep evolving and be deployed in parallel.

Optical Layer Disaggregation

As separation and virtualisation of both data and control plane has opened up the control plane of protocol stacks at layer 2 and above, recently

the possibility of opening up the optical layer is also being evaluated. The main reason this has lagged behind, while the rest of the network was undergoing disaggregation, is that there is a fundamental difference between virtualising layers operating in the digital domain (e.g., above L2) and the optical transmission layer that operates in the analogue domain and needs to address optical transmission impairments. Indeed, optical communications is typically based on closed systems where the same vendor provides transponders, in line amplifiers and ROADM nodes, so that the variability and unpredictability of the system is minimised. This is important especially for longer links (e.g., in the long-haul) where the available optical margins are squeezed to a minimum. Thus, when we attempt to open up the optical systems, making use of components from different vendors, some argue that the variability and uncertainty of performance on an end-to-end path increases, reducing the optical margin and making the network potentially less efficient. As of today there are ongoing discussions on the feasibility and benefit of disaggregating optical networks, and most recognise a trade-off between the need for increasing the amount of transmission monitoring components, in order to reduce the system uncertainty, and the cost they add to the network. Studies in¹¹ for example suggest that the benefits of optical layer disaggregation are more likely to occur in the metro area, where the optical margins are less strict than in the regional and core networks. On the other hand, there are several benefits in opening up the optical layer. The main is the ability to source components separately, which in turns allows to avoid vendor lock-in and increases competition, which

can drive down prices and improve component performance.

Work on this area is ongoing, with several consortia involved in the definition of interfaces or interoperability specification (e.g., the Open ROADM¹², the TIP Open Optical & Packet Transport¹³ and the ONF Open and Disaggregated Transport Network - ODTN¹⁴).

Upcoming Challenges

Besides the definition of frameworks, architectures and interfaces, disaggregating network functions brings up challenges related to the performance of the network and computational resources of a distributed and shared cloud infrastructure. One of the most common examples is Cloud-RAN, where the separation of remote radio head (RRH) and baseband unit (BBU) has brought up issues of tight latency constraints across the network due to bounded round trip time of acknowledgment messages and issues of RRH-BBU synchronisation. Another example is the Optical Line Terminal (OLT) virtualisation in a PON, which can provide multi-tenancy to a number of Virtual Network Operators (VNOs) sharing the same physical infrastructure. While projects like CORD have started implementing OLT disaggregation¹⁵, this mostly involve management functions, while a full MAC virtualisation, e.g.. including the Dynamic Bandwidth Allocation (DBA) is still under investigation^{16,17}. If a cloud CO based on common data centre architectures needs to disaggregate and virtualise hardware components, such CO will need to deliver bounded network and processing performance for some of the VNFs. In addition, upcoming applications linked to virtual and augmented reality will only exacerbate such requirements.

As of today, how to assure bounded latency and jitter to VNF chains across a cloud CO is still an unsolved issue, especially when scaled to several million flows. More work is thus required on end-to-end network and processing performance differentiation in the access/metro area. While some recent work has focused on NFV orchestration across domains and VNF placement optimisation¹⁸, the research community still needs to develop frameworks to scale QoS tools to cope with several million flows, within an automated framework that spans multiple network domains.

Finally, as the cloud CO becomes more and more integrated with access and metro optical networks, we envisage that dynamic, poten-

tially disaggregated, optical networking will become part of the solution, providing dynamic, low-latency links for high-capacity 5G mobile access over a multi-service, multi-tenant, statistically multiplexed network infrastructure.

Acknowledgements

Financial support from SFI 14/IA/2527 (O'SHARE), 13/RC/2077 (CONNECT) and 15/US-C2C/I3132 is gratefully acknowledged.

References

- [1] N. Feamster, et al., The Road to SDN. ACM Networks, 11(12), Dec. 2013.
- [2] N. McKeown, et al., OpenFlow: enabling innovation in campus networks. ACM SIGCOMM Computer Communication Review, 38(2), Feb. 2018.
- [3] L. Peterson, et al., Central office re-architected as a data center. Comms Mag. 54(10), Oct 2016.
- [4] Broadband Forum WT-384, 2017, <http://www.broadband-forum.org>.
- [5] <https://onap.readthedocs.io/en/latest/>
- [6] ETSI, Network Functions Virtualisation (NFV); Management and Orchestration. GS NFV-MAN V1.1.1, Dec. 2014.
- [7] A. Kamadia and N. Chase, Understanding OPNFV: Accelerate NFV Transformation using OPNFV, Mirantis editor, 2017.
- [8] www.linuxfoundation.org
- [9] H. Uzawa et al., Practical Mobile-DBA Scheme Considering Data Arrival Period for 5G Mobile Fronthaul with TDM-PON, ECOC '17.
- [10] S. Zhou et al., Mobile-PON: A high-efficiency low-latency mobile fronthaul based on functional split and TDM-PON with a unified scheduler. Th3A.3, OFC 2017.
- [11] M. Bleanger, et al., Margin requirement of disaggregating the DWDM transport system and its consequence on application economics. M1E.2, OFC '18.
- [12] Metro Open ROADM Network Model V. 1.1. White paper, July 2016, source: <http://www.openroadm.org/download.html>.
- [13] <https://telecominfraproject.com/open-optical-packet-transport/>
- [14] <https://www.opennetworking.org/solutions/odtn/>
- [15] Ali AlShabibi et al., Virtual OLT (vOLT), CORD whitepaper, Feb. 2016, source: <http://opencord.org/wp-content/uploads/2016/03/Virtual-OLT.pdf>
- [16] A. Elasad, et al., Virtual Dynamic Bandwidth Allocation Enabling True PON Multi-Tenancy. M3I.3, OFC 2017.
- [17] F. Slyne, et al., Demonstration of Real Time VNF Implementation of OLT with Virtual DBA for Sliceable Multi-Tenant PONs. Tu3D.4, OFC 2018.
- [18] L. Askari, et al., Virtual-Network-Function Placement For Dynamic Service Chaining In Metro-Area Networks. ONDM 2018.