From central office cloudification to optical network disaggregation

Marco Ruffini
CONNECT research centre,
The University of Dublin,
Trinity College, Dublin, Ireland
marco.ruffini@tcd.ie

Daniel C. Kilper
College of Optical Sciences
University of Arizona
Tucson, AZ, USA
dkilper@optics.arizona.edu

***Abstract*—The Cloud-based Central Office blurs the distinction between access, metro and data centre networks. We further show how this scenario provides new use cases for dynamic and disaggregated optical networking and new challenges that may be amenable to machine learning methods.**

Keywords—virtual CO, optical disaggregation, machine learning

I. Introduction

Telecommunications networks are in continuous evolution. Initially developed to support a world-wide telephone network for voice transmission, they have evolved, over the past 40 years, into a crucial infrastructure supporting highly heterogeneous applications, spanning a wide range, from entertainment to business, health and financial markets. While the most salient changes to the network over time were the increase in overall capacity and the move from synchronous Time Division Multiplexed (TDM) networks (e.g., Sonet/SDH) to packet-switched systems, other architectural changes have also occurred, aimed at increasing network reliability, flexibility, and especially at reducing cost of ownership.

Such changes were driven by the growth of bandwidth-hungry applications, that started to put the network under strain, and push substantial network upgrades across all parts of the network: access, metro, core and data centres. Data centres started to be deployed in larger numbers becoming closer to the end users and providing a better quality of experience. Concurrent with this, end user traffic growth has shown signs of slowing in the long haul, whereas it is accelerating in the metro area.

In parallel, as end users required higher capacity, optical fibre started to be deployed deeper into the network, either as a partial replacement for copper in fibre-to-the-cabinet (FTTCab) architectures, or all the way into the user premises (Fibre-to-the-home or premises). Besides boosting the link capacity, the deployment of access fibre brought the advantage of allowing for longer distances between users and central offices (CO), which could be extended well beyond the last mile. New architectures thus started to emerge, exploiting long transmission (even above 100 km [1]) to reduce the number of central offices, a trend known as CO consolidation, in an effort to progressively reduce the cost of network ownership. This effectively meant merging the access with the metro transmission network (a trend called access-metro convergence [2]), as the access fibre could bypass part or all of the metro network, to directly reach COs located further into the network core.

The convergence between access and metro networks is also occurring from a transmission technology perspective. As access network rates continue to increase to 10Gb/s and above, WDM becomes a common multiplexing platform across the entire network. Cloud-RAN is among the main drivers for such fast growth in access rates, as it is expected that a large amount of mobile network cells will be deployed in urban areas to increase the capacity to the end users. We can thus expect scenarios where wavelength channels of 10 Gb/s and above, generated at remote radio heads (RRH) of mobile cells, can be transmitted over dynamic WDM channels, to be routed transparently towards the baseband unit (BBU) location. This effectively unifies the transmission technology across access, metro and data centre networks.

II. From Cloud CO to disaggregated optical networks

One of the latest developments for telecoms operators, which builds on the success of Software Defined Networking (SND) and Network Function Virtualisation (NFV), is the virtualization of the central office. The idea is to replace most of the dedicated hardware equipment in a CO with software running on general purpose servers. Such hardware typically includes optical line terminals (OLT) of Passive Optical Networks (PON), baseband units (BBU) of mobile networks, broadband network gateways (BNG), and firewalls, etc.

As such functions are replaced by their equivalent software implementation, the central office turns from a collection of proprietary hardware devices into a collection of servers interconnected by switches, effectively resembling a small data centre. Indeed, the key project that brought the cloud CO concept to life was named "central office re-architected as a data centre" (CORD) [3]. This core of servers and switches is then surrounded by optical transmission systems with ROADMs to reach both into the access and metro (/core) regions of the network. The reason behind this architectural shift towards virtual COs is manifold, the first being the reduction in cost of ownership, as proprietary hardware is replaced by commodity servers and switches. In addition, NFV allows for more efficient usage of resources, as memory and computational resources can be dynamically reassigned across different functions. Finally, such an architecture can lend itself better to multi-tenancy, as different virtual providers can be given control of virtual slices of network resources, which can be grown on demand.

If COs are converging towards DC architectures, we also see that traditional DCs have started to adopt more flexible external optical interconnection technologies, such as ROADMs and WDM PON. We thus see the progressive convergence of network nodes, being data centres, cloud COs, or edge cloud nodes, towards a common architecture made up of an inner core of computational and switching resources, surrounded by an outer core providing flexible optical interconnections.

The next section describes a use case of fixed-mobile convergence in a cloud CO environment.

A. *Fixed-mobile convergence in the cloud CO*

An example of fixed-mobile convergence developed in a virtual CO environment is that of a variable-rate fronthaul architecture, initially proposed in [4]. An SDN controller communicates with the BBU to assess the instantaneous cell throughput and, as this drops below its peak capacity, the controller triggers a change in the wireless bandwidth used by the cell. This produces a reduction of the fronthaul rate, as its sample rate is proportional to the wireless bandwidth, thus making the fronthaul rate not fixed, but proportional to the actual cell traffic. As multiple cells are interconnected by a shared medium, such as a PON, the variable-rate fronthaul technique restores the statistical multiplexing properties of a group of cells, enabling efficient fronthaul transport over shared access medium. The principle was demonstrated in a testbed using an LTE software radio [5] signal transmitted over a custom PON implementation [6], showing sub-second end-to-end re-configuration times [4].

As mobile cells move towards higher bandwidth (e.g., with techniques such as carrier aggregation) and multi-channel systems (e.g., MIMO and massive-MIMO), the rate of the fronthaul link will grow proportionally. This will require highly dynamic setup and tear down of high-rate optical channels, which will possibly run transparently from the RRH to the rack or server where the BBU is located, in order to minimize end-to-end latency. However, achieving highly dynamic provisioning of wavelength channels across ROADMs networks is still challenging. The next section will discuss possible solutions that make use of techniques derived from machine learning, and provide an insight into the concept of optical network disaggregation.

B. *Optical network disaggregation through machine learning*

Over the past 10 years, networks have seen the progressive use of SDN both vertically, across the different layers of the network stack (e.g., from the IP, to MPLS and Ethernet layers) and horizontally, across multiple network domains, often supported by higher layer orchestration techniques [7]. However, only in the past couple of years have vendors and operators started to consider the possibility of applying SDN to open up the optical transmission layers, which is today mostly custom engineered, closed systems. Managing the optical layer through a central SDN controller, necessitates standardizing the system component interfaces such as the ROADM switches and WDM transceivers. To achieve this, the closed optical systems are 'disaggregated' such that the constituent components can be separately addressed and controlled through SDN interfaces. Like the transformation of the CO into a DC, disaggregation of the optical systems allows the components to be separately sourced, resulting in large potential cost savings and increased scalability. Note that this approach effectively opens up the possibility of integrating systems developed from different vendors and systems within the different network domains of access, metro, and long haul.

However, while disaggregation opens up the optical system market for more affordable and scalable solutions, it removes the performance optimization of end-to-end optical transmission engineering achieved in single vendor systems. It also introduces a greater variability and uncertainty in the component composition along a given optical path. The system dimensions can be constrained so that transmission performance and hence these uncertainties are no longer a factor. However, this limitation may compromise the cost benefits of disaggregation. Thus, novel solutions are required to tackle the transmission related challenges brought about by the optical layer disaggregation. Furthermore, methods are needed to enable these optical systems to adapt effectively to dynamic bandwidth scenarios such as described in section II.A, when the data rates exceed PON system capacities.

Machine learning (ML) is a promising new approach to improve the transmission performance, reproducibility, and stability in these new metro edge cloud networks utilizing disaggregated optical systems. The following are two examples or recent application of ML to optical network disaggregation: 1) to estimate the transmission performance of optical signals [8, 9] and 2) to predict mobile traffic variations in a cloud radio access network to allow for the long wavelength reconfiguration times of transmission systems [10]. Transmission performance estimation is particularly problematic because there is limited opportunity to collect data on deployed systems. However, transfer learning significantly reduces this burden [9].

Acknowledgment

Financial support from SFI 14/IA/2527, 15/US-C2C/I3132, and 13/RC/2077 (CONNECT) and NSF Center for Integrated Access Networks (CIAN) EEC-0812072, NSF grant CNS-1650685 and CNS-1737453 is gratefully acknowledged.

References

[1] M. Ruffini, et al., Access and metro network convergence for flexible end to end network design. IEEE/OSA JOCN, Vol. 9, No. 6, June 2017

[2] M. Ruffini, Metro-Access Network Convergence, OFC 2016, Th4B.1

[3] Central office re-architected as a data centre: www.opencord.org

[4] P. Alvarez, et al., Experimental Demonstration of SDN-controlled Var-rate Fronthaul for Converged LTE-over-PON. OFC 2018, Th2A.49

[5] Software radio stsems (SRS): http://www.softwareradiosystems.com/

[6] S. McGettrick, et al., Ultra-fast 1+1 Protection in 10 Gb/s Symmetric Long Reach PON. ECOC 2013.

[7] V. Lopez, et al.,. End-to-end Service Orchestration From Access to Backbone. IEEE/OSA JOCN, Vol. 9, No. 6, June 2017.

[8] C.L. Gutterman, et. al. Neural Network based Wavelength Assignment in Optical Switching. ACM SIGCOMM Workshop on Big Data Analytics and Machine Learning for Data Communication, 2017.

[9] W. Mo, et. al. ANN-Based Transfer Learning for QoT Prediction in Real-Time Mixed Line-Rate Systems, OFC 2018, W4F.3

[10] W. Mo, et. al. Deep Neural Network Based Dynamic Resource Reallocation of BBU Pools in 5G C-RAN Networks" OFC 2018, Th1B.