End-to-end network design and experimentation in the DISCUS project

D.B. Payne, A. Arbelaez, R. Bonk, N. J. Doran, M. Furdek, R. Jensen, N. Parsons, T. Pfeiffer, L. Quesada, C. Raack, G. Talli, P. Townsend, R. Wessaly, L. Wosinska, X. Yin and M. Ruffini

***Abstract*— This paper reports on the overall outcome of the FP7 DISCUS project, which aimed at designing end to end architectures, protocols and a physical layer exploiting Long-Reach Passive Optical Network an integrated access and metro network interconnected with a low cost flat optical core network. Our architectural modelling results show that the benefits of node consolidation on the access side extends also to the core part, by enabling a flattening of the backbone network, bringing a substantial reduction in the number of router ports. In addition we report on testbed results demonstrating end-to-end service provisioning across multiple network layers, from application down to physical layer**

***Index Terms*— access metro, network convergence, fixed mobile convergence, Long-Reach PON, flat optical core, optical island, 5G architecture, next generation multi wavelength PON.**

I. Introduction

The use of optical amplifiers after the ODN split in passive optical networks can enable a distance reach from ONU to OLT of 100 or 125 km, with an optical split increase from ~32 ways to 512 or possibly 1024 ways and bit rates of 100Gb/s when using advanced modulation techniques. Such large distances and split sizes make the LR-PON solution an ideal candidate for access-metro network convergence.

The EU FP7 project DISCUS [1] exploited the LR-PON to the full to create a new end to end network architecture that could grow network capacity for future services by 1000 times while solving the three main problems that such network growth will incur, namely: long term financial viability, energy consumption scalability and the digital divide between densely populated and sparse rural areas.

The DISCUS architecture minimises equipment and nodes in the network while maximally sharing the remaining infrastructure, equipment and nodes over as many customers as possible. It also proposed an integrated services solution that could deliver all foreseen services with either fixed access customer premises equipment (CPE) or wireless and mobile devices at the network edge. The architecture uses LR-PONs to connect customer premises and wireless edge nodes to a small number of metro-core nodes or MC-nodes. These MC-nodes are interconnected by a set of optical light paths, where a light path was defined as a point to point optical connection between a pair of MC-nodes which has packet processing only at the ingress and egress nodes and none at any intermediate nodes. This set of MC-nodes and interconnecting light paths we call an "optical island" and when combined with the physical convergence of access and metro networks by use of the LR-PON, enables of the order 50:1 reduction in network traffic processing nodes in the end to end network delivering a truly cost-effective design across the entire architecture.

This paper describes the LR-PON driven converged access-metro network solution in combination with a flat optical core network as a complete end to end architecture. It describes results from the DISCUS project and testbed demonstrator network and the major conclusions from the project.

II. Optical access architecture

The LR-PON access solution that merges the access and the metro network enables the coexistence, management and control of different client systems on a shared network infrastructure. Customers can connect via both wireless and fixed network technologies while also addressing the different needs of residential, business and enterprise customers.

The LR-PON design targets a minimum split of 512 with logical split (defined by the protocol) up to 1024 ways to exploit improved component technologies as they become lower cost in the future. The target reach is up to 125 km and the basic bit rate is 10Gb/s symmetrical using OOK modulation. The long reach and large split enables massive consolidation of central office (CO)/local exchange (LE) sites on to a small number of MC-node locations. Our study shows that typically ~98% of existing CO/LE sites can be bypassed and closed down when this architecture is deployed, with massive savings in both capital and operational expenditure and corresponding reductions in power consumption.

To enable very high capacity growth without additional fibre infrastructure installation, DWDM over the LR-PON is used with up to 50 upstream and 50 downstream wavelengths, respectively, in the C-band, when using EDFA amplifiers in the LR-PON, and many more wavelength channels when other optical bands are opened via the use of linear SOAs [2]. A fully flexible dynamic wavelength assignment scheme was developed for assigning wavelength channels for different uses (depending on the business and ownership models) and customer demands. Such a design enables economic scaling for 1000 x bandwidth increase compared to the customer usage and bandwidths in 2012 when DISCUS started.

Two approaches for the design of the LR-PON were used; one was the traditional single amplifier node (AN) at the old LE/CO sites followed by the high split ODN and is suitable for the majority of the population living in denser areas, the other design uses an amplified chain that can further extend reach and split for the much sparser rural areas where larger capture areas are required to maximise infrastructure sharing to minimise costs. This combined design targets the financial, environmental and digital divide issues that arise with the ever increasing demand for more capacity.

Figure 1 shows a general overview of the DISCUS end-to-end network solution with a meshed flat optical core and the two LR-PON designs both of which combine the access and metro networks into a converged solution. The cable architecture used to connect the ANs to the pair of MC nodes is chosen to be predominantly a cable chain or open ring design which can re-use the configuration of earlier metro SDH network cabling structures. For densely populated urban areas two fibres are required in each direction of the cable chain for each LR-PON at the AN site and these fibres pass passively through any intermediate ANs on the cable chain. For sparsely populated rural areas an amplifier chain design is used where the intermediate ANs provide additional line amplification for the LR-PONs as the optical signals traverse the cable chain to the MC nodes [3]. Both LR-PON designs use dual homing to two MC nodes for resilience/protection purposes [4].

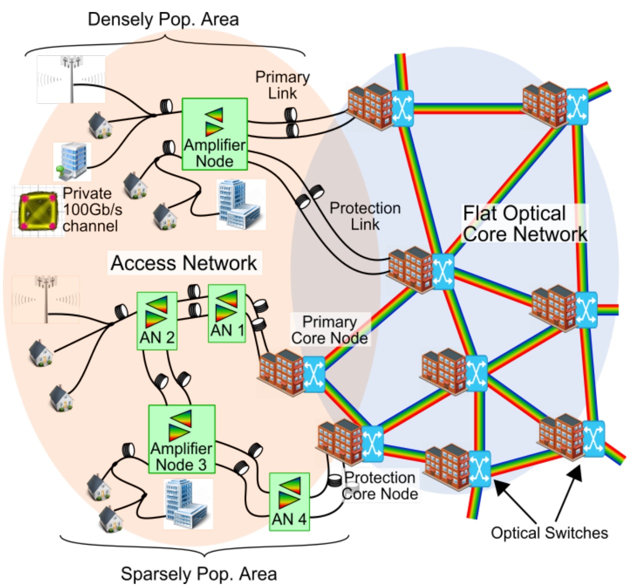

Figure 1 Access-metro converged network concept applicable for both urban and rural deployment

The use of cable chains maximises sharing of cable infrastructure and minimises cable costs compared to tree networks where each LE has its own cabling network to its pair of MCs. However the average MC-node to customer distance along the LR-PON does increase [6] and this has implication when considering current requirements of front-hauling radio systems.

The current delay requirements imposed on the fixed backhaul network are quite stringent and would limit fibre transmission reach to ~40km which is inadequate for the LR-PON access-metro convergence solution. To facilitate the use of current front-hauling protocols, some electronic processing would therefore need to be provided at a large number of LE sites to terminate the limited reach front hauling systems [7]. Depending on the size and power consumption of this equipment it may not be possible to close all the LEs that the proposed DISCUS architecture could enable. This will increase both capital and operational expenditure compared to the minimum cost network. How these additional costs are apportioned across the services requiring the front hauling technology is a question that was beyond the scope of the DISCUS project and is a question for a future project. Similarly the important question of whether the stringent delay constraints that current front-hauling places on the fixed network latency could be relaxed sufficiently to enable the increased reach required for the access-metro convergence proposed by the DISCUS architecture, will also need to be addressed in a future project. If the delay constraints can be relaxed then the lower cost and lower power consumption architecture as proposed by the DISCUS project can be realised.

The access network architecture infrastructure supports different client systems and services on different wavelengths that can carry different bit rates and modulation formats. The residential customers are addressed initially with a low cost on-off-keying (OOK) modulated TDM/TDMA protocol. As a future upgrade path, a 40Gbit/s TDM-PON using electrical-duobinary (E-DB) modulation has been studied and results were presented in [8]. Large enterprise service needs that require very large dedicated bandwidths are provided by transportation of point to point core-bandwidth light-paths over the LR-PON access network infrastructure reusing the PON fibre and ODN. Thus any dedicated bit rate from 10Gbit/s using simple OOK modulation to a coherent solution at bit rates of 100Gbit/s, in addition to shared PON protocol wavelengths, can be provided over the common LR-PON fibre infrastructure.

III. Metro-Core node design

Electronic switching, routing and aggregation are performed in the Metro-Core nodes. MC-nodes terminate the LR-PON access-metro converged networks, connected via cable chains to the ANs, which are typically located at what today are the CO/LE sites. The CO/LE site equipment and buildings can then be closed down. However Selected CO/LE sites can also be equipped with some level of wireless network local electronic processing functions if required for ultra-low latency services.

The DISCUS MC- node, architecture is presented in Figure 2. The architecture is centred on a large port count optical fibre switch which constitutes the main, flexible interconnecting layer of the node. It enables access, fibres, core fibres and any required electronic and optical equipment to be interconnected in completely reconfigurable arrangements.

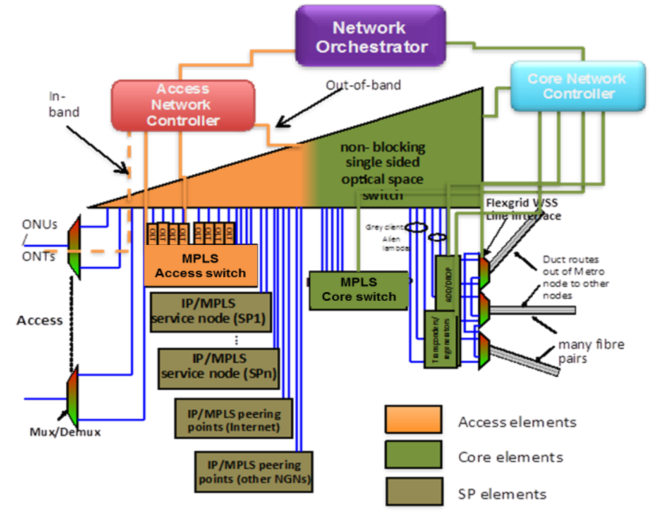

Figure 2 Metro-Core node architecture and control plane

The optical switch is implemented as a strictly non-blocking, hitless, scalable 3-stage optical switch fabric, capable of incrementally growing to over 24,000 ports with currently commercially available optical switching matrices [9]. After examining both dual-sided and single-sided Clos architectures, our analysis showed that the best choice is a single-sided folded 3-stage optical switch with an edge-fill strategy for graceful growth. This allows more flexibility with any-to-any port connectivity and simpler operations and management; it also allows some ports to only pass a single switch matrix if particularly low loss paths are required. The same single-sided reconfigurable switch is used in both the first and second stages. The single sided switch can also be partitioned so that functions that do not need full interconnection flexibility can be kept to a local partition, for example only a small fraction of access fibre require direct connection to core fibres. Partitioning can be a convenient way to either grow the switch to larger sizes or reduce the switch matrix count for a switch with a required number of ports.

On the access side, the optical switch connects the WDM multiplexers facing the PON fibre to a number of electronic termination ports. It is envisaged that a number of these will be OLTs implementing protocols such as G-PON, XGS-PON, NG-PON2 and a standardised LR-PON protocol developed in the future. Some of the wavelength channels will be used however to provide dedicated virtual point-to-point links, at various rates (10G, 100G and beyond). It is also envisaged that some access connection could be switched directly to the core, passing through ancillary equipment that could provide functions from simple amplification to complete signal regeneration and protocol adaptation.

Access traffic is aggregated through an MPLS switch in the access side, which links to the IP/MPLS nodes from service providers. Where the node is also required to carry out core grooming, i.e., where the core part of the network is not implemented through a fully flat architecture, a core MPLS switch would carry out any core aggregation required.

Figure 2 also shows the SDN control plane hierarchical architecture, comprising of:

- The access network controller, in charge of controlling the access network elements. On the south-bound interface the access controller acquires information and advises the optical line terminals (OLT's) to assign the necessary capacity. For example, user traffic flows can be prioritized based on the dynamic bandwidth assignment (DBA) protocol to establish an assured capacity per virtual local access network (VLAN) or to establish a dedicated wavelength for a particular customer demand based on a dynamic wavelength assignment (DWA) protocol.
- The core network controller, in charge of controlling the elements carrying out core transmission. The technologies considered in this architecture are Wavelength Switched Optical Network / Spectrum Switched Optical Network (WSON/SSON) networks which are based on the GMPLS distributed control plane. The core controller is in charge of receiving commands from the network orchestrator and transforming them in the southbound commands for the metro/core network elements.
- The network orchestrator, in charge of taking requests from the SP and translating them into high-level commands for the access and core network controllers.

The end-to-end control plane architecture was demonstrated across a multi-laboratory, multi-project testbed case study, presented in [10] and also in the testbed demonstrator describe later.

IV. Flat Optical Core Network

The optical switching layer within MC-nodes allows interconnection to all other MC-nodes via transparent optical light paths that form a flat optical core network which we call an "optical island". Figure 3 shows that as user bandwidth increases a threshold is reached beyond which the flat optical

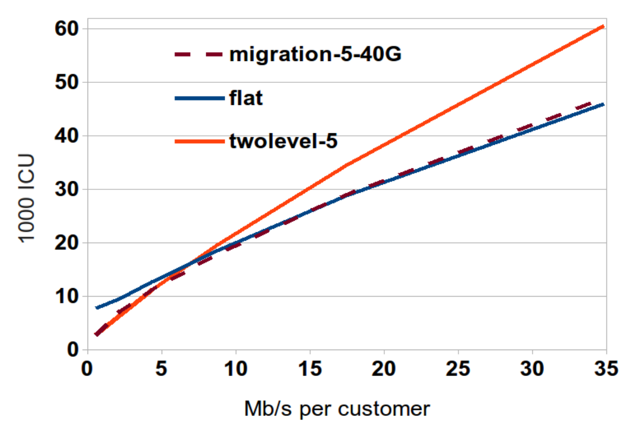

Figure 3 Relative costs of Flat and Hierarchical core networks as a function of user bandwidth.

core becomes lowest cost (and lowest power) and remains so for all subsequent bandwidth growth [11]. The results shown in **Error! Reference source not found.** are for a UK network containing 73 MC-nodes.

A strategy was also developed for evolving from the more conventional hierarchical core network, typical today, to the flat optical core network. The process is outlined in Figure 4 which shows a simple 7 node core network with initially a two tier hierarchical core network with four outer core nodes and three inner core nodes (a). As internode traffic increases some links between the outer core nodes have sufficient traffic to justify direct light-path interconnect without going through the inner core node (b).

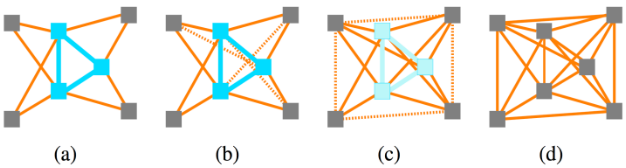

Figure 3 Migration from two level hierarchy to flat optical core

These light-paths are optically switched through the intermediate nodes. This is similar to providing optical “express “paths in today’s networks on high traffic routes. This continues with more direct light-paths between nodes as traffic grows (c) until all nodes are interconnected with transparent light-paths and the inner core nodes become part of the same hierarchical layer as the outer core nodes forming a single layer “optical island” (d).

The difference from today’s network is that small nodes that would not form the MC-nodes of the optical island would be bypassed via the LR-PON access network so that the total number of network nodes (both LE sites and small core nodes) is reduced as these nodes are closed down and bypassed leaving a small number of large MC-nodes fully light path interconnected with converged LR-PON access-metro networks dual parented onto those MC-nodes.

V. Multi-Service Testbed

To demonstrate the practical feasibility of the converged access-metro architecture as proposed in the DISCUS project, a testbed network was built that could showcase the main architectural principles and technologies. This included: 10Gb/s burst-mode receiver, tuneable filters and lasers, LR-PON ONU and OLT and SDN control plane. The demonstration testbed network was a multi-layer system integrating all network layers, from the user application down to the LR-PON protocol and the optical transmission layer.

Figure 5 illustrates the demonstrator network, it incorporates (SDN) controlled Metro-Core nodes and a TDM-DWDM LR-PON dual parented to a primary core node and a secondary MC-node. More details on this experimental test-bed setup can be found in [12]. By exploiting the dynamic allocation of DWDM channels, the LR-PON can support the convergence of a number of different user types and service demands, from residential and small business users, which share 10G PON channels, to larger business users, with options to rent dedicated 10G to high capacity 100G links.

The LR-PON demonstrator employed two different amplifier technologies in the AN: EDFAs, for use in the C-band; and linear SOAs, which are a possible solution for integrated and compact ANs and extend operation outside the C-band [2].

The three different transmission systems, 10G PON, 100G P2P and front-haul from a RRH were demonstrated simultaneously over a dynamically reconfigurable TDM-DWDM (40+40 DWDM channels in up and downstream directions) LR-PON with an ODN capable of supporting 512 users and a total reach of 100km. The 10G PON channels used a 10Gb/s linear burst-mode receiver (LBMRx) in the upstream link [13] and implemented the PON protocol in the OLTs and ONUs using field programmable gate arrays (FPGAs).

The example results in Figure 6 show bit error rate (BER) measurements obtained for the 10G PON channel in the upstream direction. The BER is measured at the OLT in burst mode operation on 2µs bursts generated by two different ONUs, with ODN path loss for the two ONUs adjusted to represent worst cast loud and soft adjacent bursts the most challenging situation for the burst-mode receiver, for more detail see [12]. The implemented SDN control plane implemented in Ryu, follows the open network foundation (ONF) architecture using three main interfaces: the application-controller plane interface (A-CPI) between the control plane and the application; the intermediate-controller

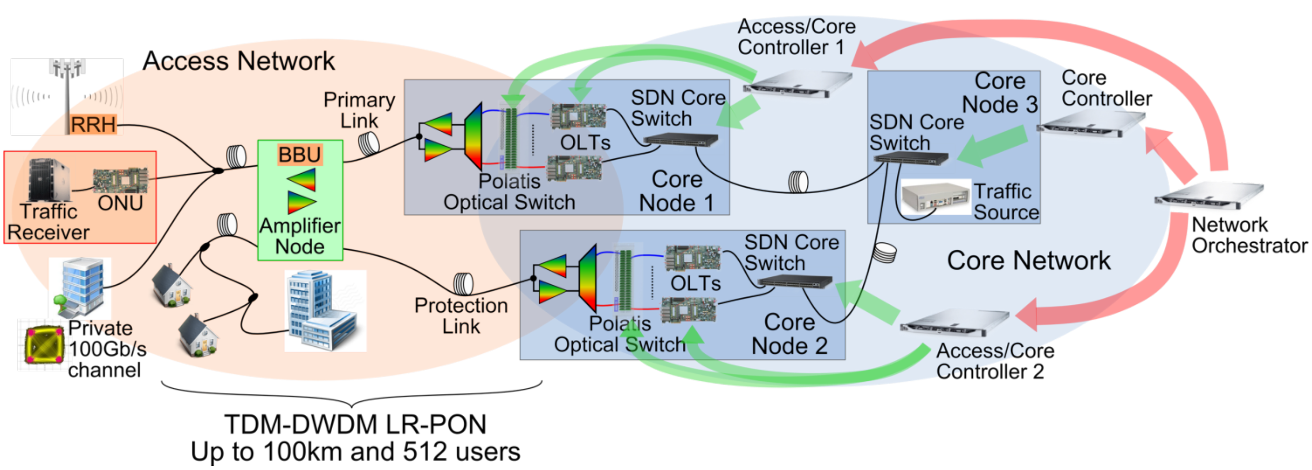

Figure 5 Network level view of the DISCUS demonstrator

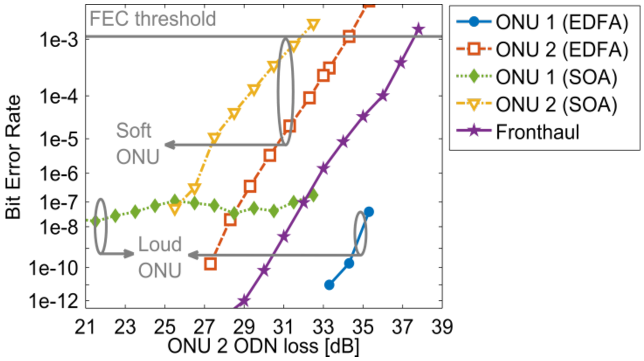

Figure 4 BER of 10Gb/s PON upstream as a function of the ODN loss for the two amplifier options.

plane interface (I-CPI) between the network orchestrator (NetO) and the access/core network controllers (NCs); and the device-controller plane interface (D-CPI) between the controllers and the physical devices [14].

The LR-PON protocol is a partial implementation of the XGPON standard, with the major differences being the longer distance and the higher split ratio supported.

Two different use cases have been demonstrated using the test bed in Figure 5 an SDN-enabled fast protection mechanism and end-to-end service restoration in case of a primary link failure [4]; and an SDN-enabled dynamic wavelength allocation (DWA) in response to an increased traffic demand.

A cost-effective 40-Gb/s single-carrier TDM-PON downstream channel was also demonstrated utilizing an integrated high-power DFB-EAM (distributed feedback laser-electro absorption modulator) in the OLT and a 3-level detection receiver with an avalanche photodiode (APD) in the optical network unit (ONU). More details on the demonstration setup can be found in [6] and [15]. The proposed low-cost 40 Gb/s downstream scheme is illustrated in Figure 7.

The 100Gb/s channel was implemented using a commercial DP-QPSK system and ran error free over the LR-PON infrastructure with adjacent DWDM wavelengths carrying 10Gb/s OOK modulated signals.

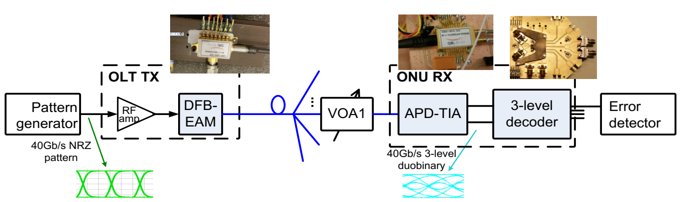

Figure 7 A cost-effective implementation of a 40-Gb/s TDM-PON downstream link: utilizing an integrated high-power DFB-EAM TOSA in the OLT and an APD-based 3-level detection receiver in the ONU

VI. Conclusion

This paper has presented the idea of using Long-Reach passive optical network for convergence of the access and metro networks into one integrated system that enables massive consolidation of CO/LEs. The work, carried out by the European DISCUS consortium, has targeted multiple areas, from the design of system level solutions and technology, to link, node and control plane architectures, to exploration of deployment optimisation over real country geographies and investigation of reliability options. Modelling results have shown the vast potential for cost and power savings, associated with the high level of node consolidation and reduction of electronic network cards enabled by the DISCUS architecture. The testbed results, which have been conducted end-to-end, across access-metro and core, and have targeted all the layers of the network, from the application down to the physical layer, have shown the practical feasibility of the concept.

Acknowledgments

Financial support from EU FP7 project DISCUS (grant CNECT- ICT- 318137) and Science Foundation Ireland (SFI) (grant 13/RC/2077), and is gratefully acknowledged.

References

[1] DISCUS web site http://www.discus-fp7.eu/

[2] R. Bonk et al., "Wavelength-transparent long-reach-high-split TWsDM-PON utilized by a non-gated parallel cascade of linear SOAs," European Conference on Optical Communication (ECOC), paper P7.4, 2014.

[3] M. Ruffini, et al., DISCUS: End-to-end network design for ubiquitous high speed broadband services. ICTON conference, 2013

[4] DISCUS Deliverable D4.13: Resiliency in heterogeneous long reach access networks. June 2015.

[5] DISCUS Deliverable D2.9: Final report on DISCUS architecture, October 2015.

[6] M. Ruffini, et al.,. Access and metro network convergence for flexible end to end network design [invited]. IEEE/OSA Journal of Communications and Networking, in press.

[7] M. Ruffini, Multi-Dimensional Convergence in Future 5G Networks. IEEE/OSA Journal of Lightwave technology, Vol. 35, No. 3, March 2017

[8] Recommendation G.989.2, “40-Gigabit-capable passive optical networks 2 (NG-PON2): Physical media dependent (PMD) layer specification”, ITU-T, December 2014.

[9] http://www.polatis.com/series-7000-384x384-port-software-controlled-optical-circuit-switch-sdn-enabled.asp

[10] S. McGettrick, et al., Experimental End-to-End Demonstration of Shared N:M Dual-Homed Protection in SDN-Controlled Long-Reach PON and Pan-European Core. IEEE/OSA Journal of Lightwave technology, vol. 34, No. 18, September 2016.

[11] C. Raack, et al., Hierarchical Versus Flat Optical Metro/Core Networks: A Systematic Cost and Migration Study. ONDM, May 2016

[12] G. Talli et al., Demonstration of SDN Enabled Dynamically Reconfigurable High Capacity Optical Access for Converged Services. OFC, PDPTh5B.1 2016.

[13] S. Porto et al., Demonstration of 10Gb/s burst-mode transmission using a linear burst-mode receiver and burst-mode electronic equalization, OFC 2014, paper M3I.5G.

[14] M. Ruffini et al., Software Defined Networking for next generation converged metro-access networks. Elsevier Opt. Fiber Tech., Vol. 26, pp 31-41, Dec.2015

[15] X. Yin, et al., 40-Gb/s TDM-PON downstream with low-cost EML transmitter and 3-level detection APD receiver. OFC 2016, paper Tu3C.1.