CORA Cork Open Research Archive
Cartlann Taighde Oscailte Chorcaí

Title	SDN enabled dynamically reconfigurable high capacity optical access architecture for converged services
Authors	Talli, Giuseppe;Slyne, Frank;Porto, Stefano;Carey, Daniel;Brandonisio, Nicola;Naughton, Alan J.;Ossieur, Peter;McGettrick, Seamas;Blumm, Christian;Ruffini, Marco;Payne, David;Bonk, Rene;Pfeiffer, Thomas;Parsons, Nick;Townsend, Peter
Publication date	2016-08-31
Original Citation	Talli, G., Slyne, F., Porto, S., Carey, D., Brandonisio, N., Naughton, A., Ossieur, P., McGettrick, S., Blumm, C., Ruffini, M., Payne, D., Bonk, R., Pfeiffer, T., Parsons, N. and Townsend, P. (2016) 'SDN enabled dynamically reconfigurable high capacity optical access architecture for converged services', Journal of Lightwave Technology, 35(3), pp. 550-560. doi: 10.1109/JLT.2016.2604864.
Type of publication	Article (peer-reviewed)
Link to publisher's version	10.1109/JLT.2016.2604864
Rights	© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Download date	2024-04-20 12:35:03
Item downloaded from	https://hdl.handle.net/10468/3184

UCC

University College Cork, Ireland

Coláiste na hOllscoile Corcaigh

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

SDN Enabled Dynamically Reconfigurable High Capacity Optical Access Architecture for Converged Services

Giuseppe Talli, *Member, IEEE*, Frank Slyne, Stefano Porto, *Member, IEEE*, Daniel Carey, *Student Member, IEEE*, Nicola Brandonisio, *Member, IEEE*, Alan Naughton, *Member, IEEE*, Peter Ossieur, *Member, IEEE*, Seamas McGettrick, Christian Blümm, *Member, IEEE*, Marco Ruffini, *Senior Member, IEEE*, David Payne, Rene Bonk, Thomas Pfeiffer, Nick Parsons and Paul Townsend, *Member, IEEE*

***Abstract*— Dynamically reconfigurable time-division multiplexing (TDM) dense wavelength division multiplexing (DWDM) long-reach passive optical networks (PONs) can support the reduction of nodes and network interfaces by enabling a fully meshed flat optical core. In this paper we demonstrate the flexibility of the TDM-DWDM PON architecture, which can enable the convergence of multiple service types on a single physical layer. Heterogeneous services and modulation formats, i.e. residential 10G PON channels, business 100G dedicated channel and wireless fronthaul, are demonstrated co-existing on the same long reach TDM-DWDM PON system, with up to 100km reach, 512 users and emulated system load of 40 channels, employing amplifier nodes with either erbium doped fiber amplifiers (EDFAs) or semiconductor optical amplifiers (SOAs). For the first time end-to-end software defined networking (SDN) management of the access and core network elements is also implemented and integrated with the PON physical layer in order to demonstrate two service use cases: a fast protection mechanism with end-to-end service restoration in the case of a primary link failure; and dynamic wavelength allocation (DWA) in response to an increased traffic demand.**

***Index Terms*— Optical fiber communication, passive optical networks, dense wavelength division multiplexing, erbium doped fiber amplifiers, semiconductor optical amplifiers, access protocols, software defined networking, fast protection, dynamic wavelength allocation**

I. Introduction

An end-to-end network design, which considers access, metro and core networks, can provide a more economically scalable architecture solution that is better able to cope with the future data traffic growth and also able to support the ubiquitous delivery of high speed services to customers regardless of type or location [1]. The number of network interfaces and switches can be reduced by using a fully meshed flat optical core, thus reducing cost and power consumption. In order to maintain a flat all-optical structure the number of core nodes (CNs) will need to be around 100 for a typical large EU country. This sparser node distribution can be supported by a growth both in reach and capacity in the access network. Passive optical networks (PONs) represent a cost effective mechanism for deploying fibre-to-the-home, where protocols like G-PON and the 10G version XG-PON allow multiple users to share a time division multiplexed (TDM) channel. These were recently extended by the NG-PON2 standard [2] that introduces wavelength division multiplexing (WDM), allowing the use of multiple TDM and point-to-point wavelength channels over the same physical PON fiber infrastructure. In this paper we further enhance this solution by proposing dynamically reconfigurable time-division multiplexing dense wavelength division multiplexing (DWDM) long-reach passive optical networks (LR-PONs) to provide an efficient solution to enable the convergence of multiple service demands and different user types on a single architecture. For densely populated areas a simple and efficient architecture can use a single optical amplifier node (AN) from where a totally passive optical distribution network (ODN), with splits up to 512 or 1024 users and reaches of 20km, can be connected forming a tree-like structure [3]. By exploiting the dynamic allocation of DWDM channels, the LR-PON can support both residential users, which share a 10G PON channel, and business users, with options to rent dedicated 10G PON channels or high capacity 100G links. Low latency services, such as wireless front-hauling, could also be supported by placing an optical networking unit (ONU) and the necessary processing equipment in the AN and allocating a dedicated wavelength from the AN to the remote site (e.g. a CPRI channel).

Fig. 1 illustrates the network concept and demonstrator, which incorporates a TDM-DWDM LR-PON, connected to a primary CN with a protection link to a secondary CN, and a flat optical core, mutually controlled using software defined

Financial support from Science Foundation Ireland (SFI) (grants 12/IA/1270, 12/RC/2276 and 10/CE/I1853), and EU FP7 project DISCUS (grant CNECT-ICT- 318137) is gratefully acknowledged.

G. Talli, S. Porto, D. Carey, N. Brandonisio, A. Naughton, P. Ossieur and P. Townsend are with the Tyndall National Institute, University College Cork, Cork, Ireland (e-mail: giuseppe.talli@tyndall.ie).

F. Slyne, S. McGettrick, C. Blümm, M. Ruffini and D. Payne are with the CONNECT research center in Trinity College, Dublin, Ireland.

R. Bonk and T. Pfeiffer are with Nokia Bell Labs, Stuttgart, Germany.

Nick Parsons is with Polatis Ltd., Cambridge, United Kingdom.

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

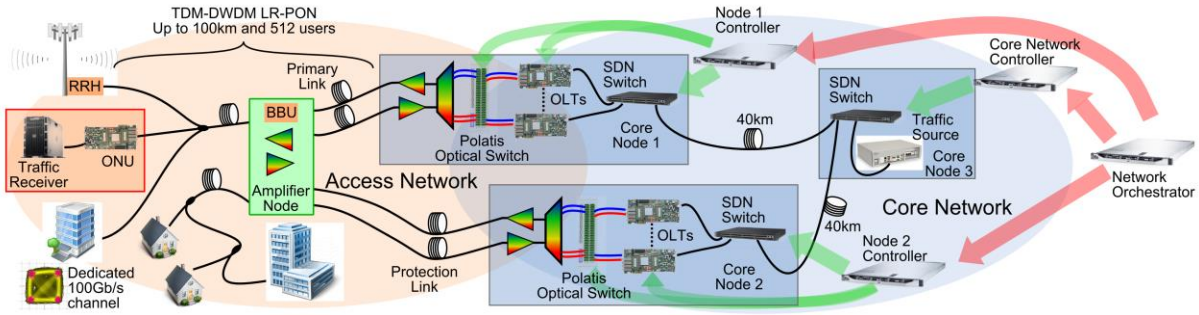

Fig. 1. Network level view of the demonstration.

networking (SDN). The SDN control plane can enable highly dynamic service and capacity provision over the LR-PON in response to changing demand by implementing agents in the network elements [4]. Since in a LR-PON architecture, much of the currently protected metro network is replaced by the backhaul links, the ANs, which are serving hundreds of residential and potentially business users, would need to be dual parented for protection with a secondary backhaul link to a geographically separated CN. Fast protection of these links is required in order to fulfill user requirements, particularly for enterprise and mobile backhaul applications. Fast PON protection enabled by SDN control also allows the implementation of protection load balancing schemes, which allow substantial cost reduction in both IP and PON backup resources [5], by increasing the ability to share protection equipment across the network.

This paper, based on the work presented in [6], substantially extends previous demonstrations of evolutionary access networks [4, 7-9]. For example the work presented in [4] focuses only on the SDN control plane aspects of a TDM-DWDM LR-PON and does not include integration with a realistic physical layer, while the demonstrations in [8-9] focus only on the physical layer and on statically allocated DWDM channels. A very advanced demonstration of a time and wavelength division multiplexing (TWDM)-PON is reported in [7], which supports up to 1024 users over 40km reach, using 4+4 dynamically tunable channels in the up- and down-stream direction and also integrates protocol functionalities. Compared to the previous examples, in this work, for the first time to the best of our knowledge, a dynamically reconfigurable TDM-DWDM PON physical layer is demonstrated with 100km reach and 512 split, which is capable of supporting an emulated system load of 40+40 channels in the up- and down-stream direction and incorporates protocol implementation in the optical line terminals (OLTs) and ONUs. The physical layer, operating in burst mode in the upstream direction, includes the use of a 10Gb/s linear burst-mode receiver (LBMRx) and ANs capable of carrying heterogeneous services and modulation formats employing burst-mode capable erbium doped fiber amplifiers (EDFAs) and semiconductor optical amplifiers (SOAs). The demonstrator also presents the physical integration of an SDN controlled access and core network implementing two use cases: an SDN-enabled fast protection mechanism with end-to-end service restoration in case of a primary link failure; and an SDN-enabled dynamic wavelength allocation (DWA) in response to an increased traffic demand. The SDN-enabled DWA is used for dynamic capacity provisioning and hence wavelength and service reconfiguration times of a few hundred ms are targeted. The following section presents a detailed description of the experimental test-bed, while Section III describes the SDN architecture and the PON protocol implementation. Section IV presents the experimental results subdivided between physical layer characterization of the 10Gb/s PON channels, the 100G private wavelength, and the two control plane use cases: fast protection and DWA.

II. Test-bed description

Fig. 2 shows the details of the LR-PON physical layer, including both amplifier technologies used in the AN. The EDFAs, in Fig. 2(a), are readily available and a well-known platform for amplification in the C-band, while the SOAs, in Fig. 2(b), are a possible solution for integrated and compact ANs and also for extended operation outside the C-band [10]. The same ONUs and CNs are connected to the two LR-PON test-beds using a Polatis optical switch used for experimental convenience to reliably interchange between the two AN designs. Fig. 2 also shows the details of the ODNs, of the backhaul links implemented and the details of the AN designs, such as amplifier gains and channel powers. Attenuators are added in the ODNs to emulate end of life standard single mode fiber attenuation (0.3dB/km) and realistic splitter losses including excess loss [2]. Throughout this work we consider in the ODN loss calculations a worst case loss of 3.5dB for each ×2 split. Part of the total ODN split is located in the AN, 4×4 or 2×4 respectively for EDFA and SOA ANs. This is used to combine up- and down-stream traffic and also provides access to a redundancy path in the backhaul link for resilience and protection. The protection link is implemented only for the EDFA-design and consisted of just a few meters of fiber to achieve the maximum differential reach with the primary path. Attenuators are used also in the backhaul link to emulate end of life fiber attenuation.

In the EDFA-based AN, the downstream signals coming from the primary backhaul link are amplified using a two stage EDFA configuration with an overall gain of 38dB. A dispersion compensating fiber (DCF) module, equivalent to 40km of standard single mode fiber (SMF), is used mid-stage to partially compensate the primary downstream link. The high power second stage EDFA, with a maximum output power of +30dBm, is shared between the primary and secondary downstream link, in order to reduce the number of

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

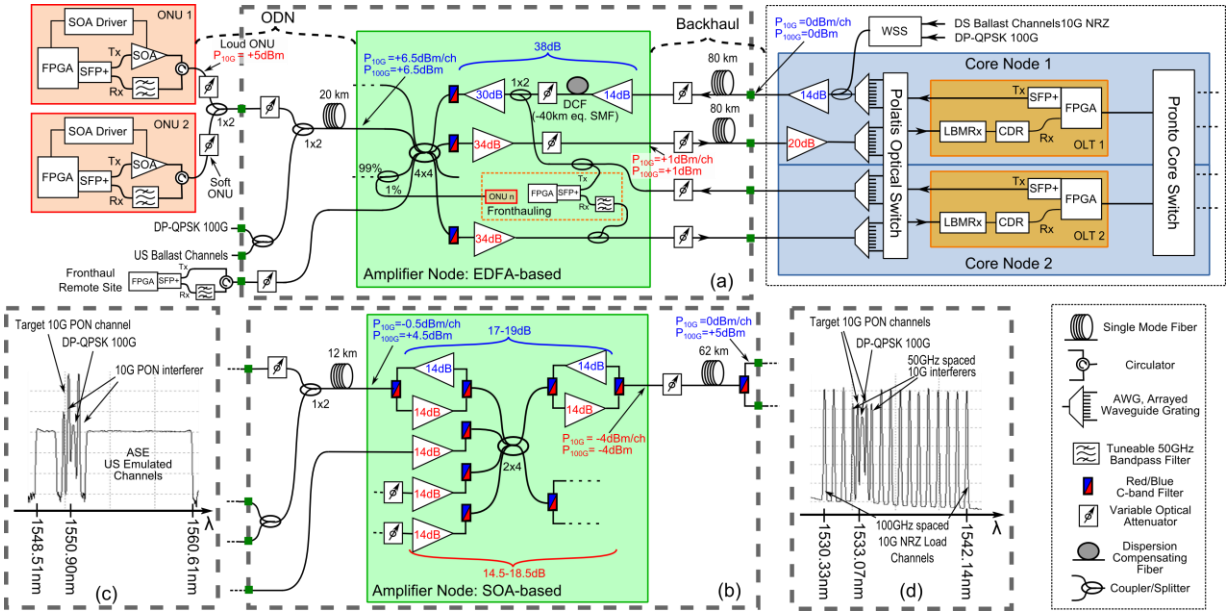

Fig. 2. Physical layer experimental setup with EDFA-based (a) and SOA-based (b) amplifier node; upstream (c) and downstream (d) spectra.

EDFAs in the laboratory demonstrator, but in a real deployment the secondary path should employ a separate device to increase resilience. A 3dB coupler is used mid-stage to provide access to the second stage EDFA from the secondary path. The upstream primary and secondary path EDFAs are each connected to a port of the 4×4 splitter, through a red/blue C-band filter in order to remove potential back-reflection from the high power downstream signals. All the EDFAs used in upstream are commercial devices with fast gain stabilization in order to reduce the impact of gain transients induced by the burst traffic [11].

The SOA-based AN comprises a chain of SOAs for individual amplification of downstream and upstream wavelength bands. Red/blue C-band filters are used to separate the upstream and downstream bands at each instance. The amplifier chain is established by equipping each branch of a 2×4 splitter with an SOA for the down- and up-stream direction, respectively. In the experiments, the downstream direction is equipped with an SOA at the primary backhaul connection as well as one of the 2×4 splitter branches. In the upstream direction, each of the 4 branches includes an SOA to consider the effect of noise funneling. The effective small-signal gain of the entire amplifier node is 18dB at 1550nm in upstream and 19dB at 1530nm for the downstream direction. The effective noise figure of the AN is around 13dB in upstream direction including the noise funneling contributions and about 9 dB in downstream direction. The 3dB saturation input power of the AN is +1dBm in upstream direction and –2dBm in downstream direction, while the 1dB saturation input power is –5dBm and –6dBm respectively.

Both ANs are designed to support 40 downstream and 40 upstream channels on a 50GHz spaced C-band wavelength plan. In the experimental demonstration the downstream channels are located between 1530.33 and 1542.14nm and the upstream channels are located between 1548.51 and 1560.61nm, leaving a guard band of approximately 6nm between down- and up-stream. It should be noted that, even though the system and in particular the optical amplifiers are designed to support 40 channels in both directions, the wavelength plan adopted in this demonstration can only support 32 channels in order to use available band splitting filters. Two OLTs and two ONUs are fully implemented and additional traffic equivalent to a power of 40 channels in both up- and down-stream directions is added to fully load the network. As can be seen from the spectrum in Fig. 2(d) the additional downstream channels are emulated using 15 distributed feedback (DFB) lasers, which are externally modulated at 10Gb/s with on-off keying (OOK) using a LiNbO3 Mach-Zehnder modulator. The power of these additional channels is set equal to nominal power for the two channels neighboring the target PON downstream channels, while the power of the remaining 13 channels is adjusted in order to obtain an overall downstream power equivalent to 40 channels at nominal power. For experimental convenience the additional downstream channels, coupled in the system at the CN, after the Polatis optical switch and before the booster amplifier, are always present, which allows the use of simpler EDFA controls in the downstream link.

Fig. 2(c) shows the upstream band, where the additional channels are emulated using amplified spontaneous emission (ASE), generated by an SOA which is filtered and flattened using a wavelength selective switch (WSS) and then modulated to mimic the typical burst mode power variations of upstream channels. The WSS also enables a band to be carved in the ASE for the 10G PON signals, the 100G signal and the fronthaul signal. The emulated upstream additional channels are inserted in the ODN using a spare splitter arm and the power is adjusted using an optical attenuator to emulate the power of 36 additional channels.

The OLT and ONU transmitters are implemented using commercial tunable 10G small form-factor pluggable (SFP+) transceivers driven by field programmable gate arrays (FPGAs) implementing the PON protocol. In addition to the

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

experimental results presented in [6] a forward error correction (FEC) algorithm has also been implemented for physical layer testing on FPGAs using standard Reed-Solomon (RS) encoders and decoders, based on RS(248,216) both in upstream and downstream [12]. Data is received by the FPGAs in blocks of 248 symbols and each symbol of 8 bits is processed by one decoder. A group of 8 interleaved decoders has been used for processing 8 symbols simultaneously with a bit rate of 10Gb/s.

In the ONUs the wavelength selection at the receivers is achieved by using a 50GHz tunable filter, while an external SOA is also used after the SFP+ transmitter to create the envelope of the bursts and blank the ONU output during the periods of time when data is not transmitted [13], with the driving signal generated by the FPGA. It should be noted that the ONU SFP+ transmitter with the external gating SOA is only an emulation, which is sufficient for the purpose of this system demonstration, but the transmitter used in a real ONU would have to be optimized for the specific application in terms of cost and performance. The ONU receiver and tunable filter assembly would also need to be optimized in terms of cost, while the FPGA would probably be replaced by an application specific integrated circuit (ASIC).

At the CN side the tunability is realized by routing the ports of a wavelength demux to the OLTs using a Polatis optical switch, partitioned and shared by the two CNs. At the OLT the upstream signal is detected using the LBMRx [14], which is followed by a static electronic dispersion compensation (EDC)-clock data recovery (CDR) module [9]. The use of a LBMRx, compared to a more commonly used limiting BMRx is also necessary in order to preserve the pulse-shape of the signal that is then processed by the EDC [15]. While the EDC is operated in static mode, the CDR of the module is able to recover the phase of the incoming bursts within a relatively short preamble of ~200ns. In this demonstration a static EDC is sufficient since both implemented ONUs are located at the same distance, but in a realistic deployment scenario a fast, burst-by-burst, adapting EDC would be required [15]. The recovered and re-timed signal at the output of the CDR module is then sent to the frame synchronizer, the first processing block of the PON protocol implemented in the FPGA.

The 100G channel is implemented using a commercial transponder to generate and receive a dual-polarization quadrature-phase-shift-keying (DP-QPSK) signal and is inserted using similar entry points as the down- and up-stream emulated channels. The fronthaul channel, demonstrated only for the EDFA case, is emulated by adding SFP+ transceivers driven by FPGAs located in the AN and at one ODN arm as shown in Fig. 2(a). Architecturally, the backhauling part of the wireless signal can be carried over one of the PON channels and terminated in the AN using an ONU connected to a 1% power tap on one of the ODN 4×4 splitter port.

III. SDN Architecture and PON Protocol

The overall SDN architecture, shown in Fig. 3, follows the open network foundation (ONF) architecture using three main

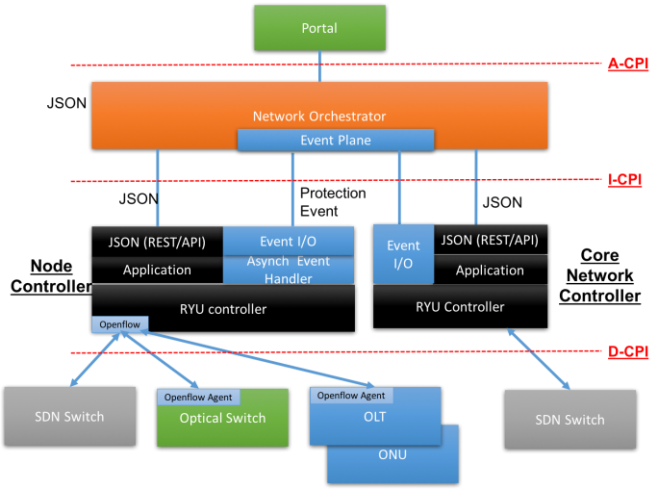

Fig. 3. Architecture of the overall SDN control plane.

interfaces [16]: the application-controller plane interface (A-CPI) between the control plane and the application (a user portal in this case); the intermediate-controller plane interface (I-CPI) between the network orchestrator (NetO) [17] and the node controller (NC) and core network controller (OF); and the device-controller plane interface (D-CPI) between the controllers and the physical devices [4].

The D-CPI operates between the physical devices, i.e. OLT, optical switch and OpenFlow (OF) SDN switch, and the node controller. The OF SDN switch, a Pronto 3780 48x10G ports running release 2.4 (OF v1.4 compatible firmware) implements a native OF interface. Thus the communication to the controller is carried out over the OF v1.4 protocol through the Ryu OF controller application programming interfaces (APIs). The OLT and the optical switch do not implement a native OF v1.4 protocol, thus we have developed an OF agent collocated with the controller that emulates the OF protocol link. The controller sends OF-compatible commands to the OF agent, which then translates the messages into appropriate commands for the physical optical switch and OLT devices. The physical link to the optical switch is then carried out through a transaction language 1 (TL1) session, while the OLT uses a proprietary messaging system over a UART serial interface. TL1 is a man-machine management protocol defined by Bellcore which is commonly used to manage optical broadband and access equipment. The type of messages implemented by our controllers and orchestrator are reported in Table I.

The D-CPI messages shown in Table I are those between the OF agent and the physical device. The 'Patch_Connect' and 'Patch_Disconnect' messages are those that invoke connection between optical switch ports. These are translated over the interface to the TL1 'ENT-PATCH::inport,outport:;' and 'DLT-PATCH::inport,outport:;' messages. The D-CPI messages to the SDN switch are not reported here as they follow the OF V1.4 syntax. Also the D-CPI messages from the NC to the OF agent for the optical switch and OLT follow the OF v1.4 standard. The only extension to the standard OF protocol is for the wavelength selection at the OLT. The

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

TABLE I
LIST OF THE MAIN CONTROL PLANE MESSAGES USED IN OUR IMPLEMENTATION FOR THE USE CASES PRESENTED

Interface Type	Command Type	Source	Dest.	Use Case	Main Parameters
D-CPI	Patch_Connect	NC	Opt. Switch	Protect., DWA	Input-output ports
D-CPI	Patch_Disconnect	NC	Opt. Switch	Protect., DWA	Input-output ports
D-CPI	Status_Report	OLT	NC	Protect., DWA	Status value
D-CPI	Failure_Detected	OLT	NC	Protec.	ID of pre-set failure_event
D-CPI	Create_Flow	NC	OLT	Protect., DWA	MAC, xgem_port, mpls_tag
D-CPI	Delete_Flow	NC	OLT	Protect., DWA	Flow_ID
D-CPI	Set_DS_lambda	NC	OLT	DWA	ONU_ID, channel
D-CPI	Set_US_lambda	NC	OLT	DWA	ONU_ID, channel
I-CPI	Failure_Detected	NC	Orch.	Protec.	---
I-CPI	Invoke_Failover	Orch.	NC, CNC	Protect.	---
I-CPI	Client_Failure_ Recovery	Orch.	NC	Protect.	---
I-CPI	Create_Path	Orch.	NC, CNC	DWA	Source, Destination, QoS params
A-CPI	Resource_Request	Portal	Orch.	DWA	Source, Destination, QoS params
A-CPI	Resource_Confirm ation	Orch.	Portal	DWA	Request ID

'ofp_port_mod_prop_optical' data structure and the 'OFPOPF Configure' method which are defined in the OF v1.4 standard have been enhanced to include an FPGA stream identifier used to differentiate between the upstream and downstream directions of the port considered. These are then mapped to the parameters of the 'Set_US_lambda' and 'Set_DS_lambda' D-CPI functions (described below) from the OF agent to the FPGA Microblaze controller interface in the OLT.

Most of the commands reported in the table are those between the OF_agent and the OLT. It should be noted that here we present an abstracted readable format, rather than the low-level messages sent to the Microblaze soft processor of the FPGA board. The 'Status_Report' is a general message to report the status of a given parameter. The 'Failure_Detected' message indicates that the OLT has identified a failure and it triggers the protection action at the NC. The 'Create_Flow' creates an entry on the OLT flow table, providing flow association between the device MAC address, the ONU XG-PON encapsulation method (XGEM) port and the multiprotocol label switching (MPLS) label (used to identify a Pseudo-Wire). The NC uses the 'Set_DS_lambda' and 'Set_US_lambda' for indicating the downstream and upstream transmission wavelengths to the ONU. This last command is sent to the OLT, which then generates a physical layer operation, administration and maintenance (PLOAM) message to communicate with the ONU. Once the ONU has received the corresponding XG-PON PLOAM message, it sends a 'Write WL<lambda>' over the UART interface to change the center wavelength of its tunable filter.

While our implementation of the NC-to-OLT is proprietary, very recently a standardization effort was initiated by the Broadband forum to define a D-CPI OF interface for PONs (e.g. see contribution [18]).

The I-CPI layer takes charge of the messaging between the NetO and NCs. The messages used in our case are: "Failure_Detected" reporting from the NC to the NetO that one of the OLT connections has failed; 'Invoke_Failover', used by the NetO to activate the pre-configured protection path on the NCs; "Client_Failure_Recovery", is used by the NetO to inform the NC that the protection was successfully established. Alternatively, the DWA use case employs a 'Create_Path" message with source, destination and quality of service (QoS) parameters (triggered by the 'Resource_Request' message described below).

In order to accelerate some of the time-critical messages at the I-CPI level, we have created an asynchronous module to pick up the failure event form the OLT. This is then passed to an Event Plane, implemented using the ZeroMQ libraries (as explained in [19]), which bypasses the slower Javascript object notation (JSON) interfaces. In this experiment we have used a proprietary interface for the I-CPI, although in [20] we have demonstrated the interoperability of our NC with the control orchestration protocol (COP) [21].

The A-CPI interface operates between a user portal and the NetO. In the DWA use case, a "Resource_Request" is sent by the portal to the NetO indicating the source and destination points as well as the relevant QoS parameters such as committed information rate (CIR) and peak information rate (PIR).

The LR-PON protocol is a partial implementation of the 10G symmetric XG-PON standard [12]. However, not all functions from the standard are implemented in our system. In addition we carried out some modifications to the XG-PON standard to adapt the protocol to the longer PON logical distance supported, up to 125km as opposed to 60km, and the higher split ratio supported, 512 versus 64. The reader can refer to [19] for more details on the OLT and ONU implementation.

The protocol is implemented over four Xilinx Virtex7 XC7VX690T FPGAs acting as primary and secondary OLTs and two ONUs. A Xilinx Microblaze soft processor, which is instantiated on the FPGAs, provides a (North Bound) UART management interface to the PON OLTs and ONUs hardware. This interface allows control of PON functionalities such as resetting the hardware, viewing hardware status, simulating hardware failure, loading the bandwidth map and setting XGEM mappings.

In the experimental test-bed a simplified core network is emulated using a Pronto 3780 switch with 48 10GbE interfaces, running Openflow v1.4 compatible firmware and

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

configured with three virtual OF bridges. The PON backplane connection between the OLTs and the core network is enabled by the Pronto switch 10GbE interfaces, which are also used to interconnect the virtual bridges (as shown in Fig. 1) using SFP+ transceivers and two 40km fiber links.

IV. Experimental Results

A. 10G PON and Fronthaul Channels

The overall performance of the downstream link is effectively power limited at ONU receiver and Fig. 4 shows the bit error rate (BER) measured as a function of the ODN loss for the two AN designs. The results for the SOA- and EDFA-based ANs and the fronthaul channel have also been reported in [6] while additional results for the EDFA-based obtained using the FPGA code implementing the RS(248,216) FEC are added here. Assuming an FEC threshold of 1.1e-3, both designs are able to support an ODN loss of at least 28dB, which corresponds to a 128 split plus 12km of fiber. Considering the additional split in the AN, the overall split ratio supported by the system is 128×4=512. The post-FEC BER measurements of the EDFA case confirm that a BER of <1e-12 is obtained for pre-FEC BER below the FEC threshold of 1.1e-3.

The difference in the performance of ~2dB between the two ONUs is due to the difference in their receiver sensitivities, which is within the spread of the performance of the commercial SFP+ receivers. The difference in performance between the AN designs of ~8dB roughly corresponds to the 7dB difference in power launched from the ANs in the ODN. The launched power is limited in the SOA case by the saturation power of available devices. The overall output power of the last SOA is around +16dBm, which is beyond the 3dB saturation power for the devices employed (+14dBm). Under these conditions in single channel operation we should expect severe patterning that would impair the 10G channels.

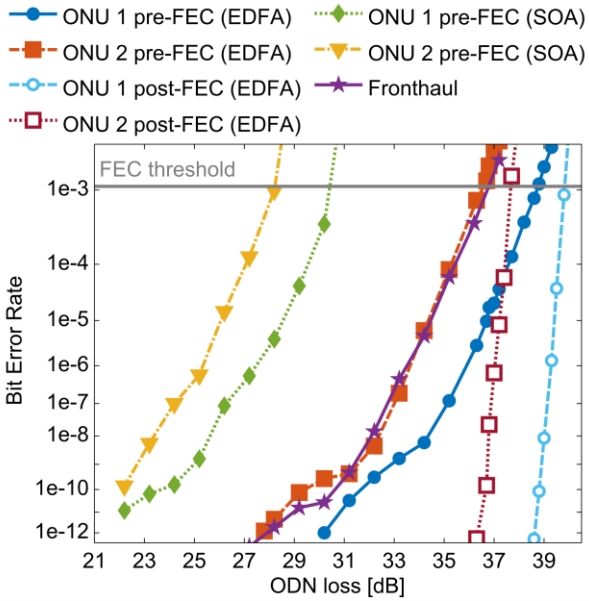

Fig. 4. BER of 10Gb/s PON downstream as a function of the ODN loss for the two amplifier node designs and BER of the fronthaul channel as a function of the ODN loss.

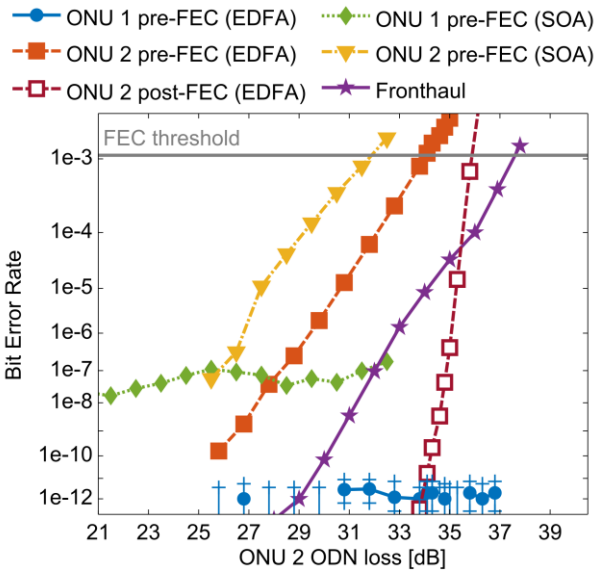

Fig. 5. BER of 10Gb/s PON upstream as a function of ONU 2 ODN loss for the two amplifier node designs and BER of the fronthaul channel as a function of the ODN loss.

However, due to the fact that the channels are de-correlated and that their individual powers are lower than the saturation power, the patterning causes only a small penalty. In the EDFA case the power launched from the AN is sufficient to support an even a higher split ratio. Hence, by optimizing the system design for a specific split ratio the power could be reduced, potentially reducing the cost of the EDFAs and the overall AN.

The performance of the upstream link was evaluated as a function of the ONU 2 burst power, while maintaining ONU 1 burst power constant close to the LBMRx overload power, effectively covering the entire dynamic range (DR). The loud burst from ONU 1 is hence acting as a worst case interferer for ONU 2 in terms of the LBMRx operation. Fig. 5 presents the upstream BER measured by the OLT in burst mode operation on 2µs bursts generated by the two ONUs. The ODN loss was varied from 16 to 35dB only for ONU 2 in order to vary the DR of the burst powers from the two ONUs at the LBMRx. The burst power from ONU 1 is maintained close to the LBMRx overload power (loud ONU) with an ODN loss of 16dB. Fig. 6, shows as an example, two traces recorded at the entry to the ODN, when the signals of the two ONUs are coupled together with different dynamic ranges. We can clearly see the two bursts with 5dB power dynamic range, where the loud packet is from ONU 1 and the soft packet is generated by ONU 2, while with a DR of 17dB the soft packet is barely visible above the noise floor.

Similarly to the downstream case, the results for the EDFA-

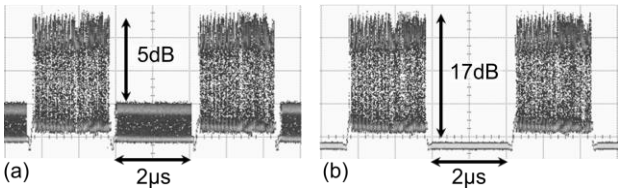

Fig. 6. ONU1 and ONU2 bursts launched in the ODN with 5dB (a) and 17dB (b) power DR.

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

based AN are obtained implementing FEC in the FPGA, but in this case operating in burst mode. The BER is always below the FEC threshold for ONU 1 (Fig. 5) and more precisely between 1e-7 and 1e-8 for the SOA case. The ONU 1 pre-FEC BER in the EDFA case is around 1e-12, with some data points showing error free operation within the time interval of the measurement and with the error bar indicating the 95% confidence interval [22]. The small variation of ONU 1 pre-FEC BER with the variation of ONU 2 ODN loss, and hence DR, is due to the non-ideal residual error in the fast gain setting of the burst mode receiver transimpedance- and post-amplifier. This small effect is more evident because ONU 1 is working close to the receiver overload. The post-FEC BER of ONU 1 is error free over the entire DR and hence it is not presented in Fig. 5.

The comparison of the ONU 2 pre- and post-FEC BER in the EDFA case confirms that the FEC threshold of 1.1e-3 is applicable also in burst mode operation for the upstream link. The ONU 2 BER shows that the AN designs can support ODN losses of up to 31.5dB and 34dB, corresponding to dynamic ranges of 15.5dB and 18dB respectively for the SOA and EDFA designs. Both designs can support a larger DR than that introduced by the non-uniform loss of the ODN splitters, which can be up to 12dB considering a 512 split ratio and realistic splitter losses [2]. The performance of the upstream link is optical signal to noise ratio (OSNR) limited and hence the EDFA AN shows slightly better performance compared to the SOA-based design due to the lower noise figure (NF).

Fig. 4 and Fig. 5 also show the BER measured for the fronthaul channel, which can also operate error free, with a BER of <1e-12, in both directions for an ODN loss of up to 28dB. The performance in downstream follows closely the performance of the downstream PON channel and it is essentially limited by the power launched from the AN and by the receiver sensitivity. Since the fronthaul is a dedicated channel terminated in the AN, a slightly higher power than the PON channels could also be launched in order to include some system margin. In the upstream direction the receiver in the AN is optically preamplified by one of the upstream EDFAs and hence the performance is limited by the power launched into the ODN and by the EDFA NF. The fronthaul channel shows better performance than the PON upstream channel as a function of the ODN loss because it is not transmitted through the link towards the CN and hence maintains a higher OSNR. Moreover it should be noted that the fronthaul channel is a continuous mode link and hence does not require the use of a burst-mode receiver.

B. 100G Link

The performance of the 100G link was characterized in downstream as a function of the ODN loss using an EDFA pre-amplifier at the receiver with 5.5dB NF. The BER as a function of the ODN loss is shown in Fig. 7 for both AN configurations. The performance of the 100G downstream link is impaired by the OSNR degradation and also by the non-linear impairments introduced by the fiber propagation and in the SOAs. The OSNR degradation is dependent on the ODN

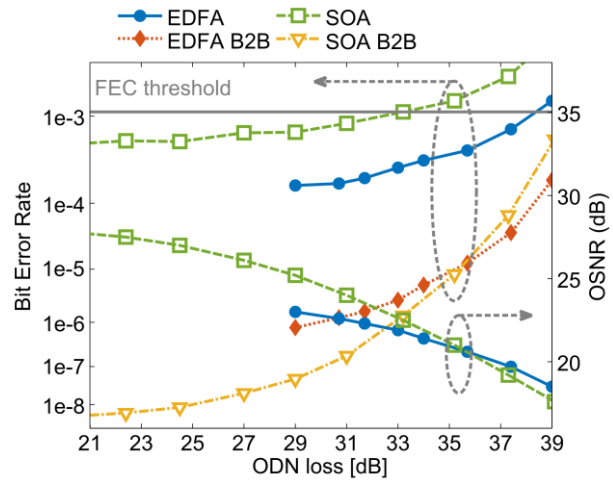

Fig. 7. Performance of the 100G link in downstream as a function of the ODN loss for the two amplifier node designs.

loss and Fig. 7 also presents the measured OSNR of the 100G signal before the downstream receiver. For high ODN loss the OSNR is limited by the ASE introduced by EDFA pre-amplifier in front of the 100G receiver. The small difference between the two AN configurations is because the EDFA-based AN launches a higher 100G signal power into the ODN (2dB higher). For low ODN loss the limitation in OSNR comes from the AN and the CN amplifiers. The SOA design has a better OSNR because by design the 100G is +5dB higher than the other nominal channels and also because it presents lower losses in the backhaul section due to the shorter backhaul link. As shown in Fig. 7 both AN designs can support the 100G channel with ODN loss higher than 28dB, which is equivalent overall to 128×4=512 users. Despite the higher OSNR the SOA-based AN has poorer performance due to the non-linear cross-talk with the 10Gb/s downstream channels mostly caused by cross-gain and coupled phase modulation inside the SOA. The EDFA-based design is also impaired by non-linear cross-talk with the 10Gb/s downstream channel, in this case caused by cross-phase modulation (XPM) in the ODN and backhaul fibers.

In order to evaluate the effect of the non-linearities for both AN designs Fig. 7 also shows the BER measured as a function of the ODN loss in back-to-back (B2B). In this case the 100G transmitter is connected directly to the EDFA pre-amplifier at the receiver without going through the LR-PON. The ODN loss is emulated by a variable optical attenuator at the input of the EDFA pre-amplifier, which is used to change the OSNR of the signal before the 100G coherent receiver. The value of the ODN loss is estimated by comparing the B2B OSNR (measured after the EDFA pre-amplifier) with the one measured for the two LR-PON architectures. Assuming a negligible impact from other sources of impairments, such as for example chromatic dispersion, the effect of the non-linearities in the ODN and backhaul fibers can be clearly seen by comparing the B2B curve with the one obtained with the complete LR-PON for the EDFA case. Similarly by comparing the two curves for the SOA-based AN, the penalty introduced by the non-linear cross-gain and coupled phase modulation inside the SOA is clearly evident.

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

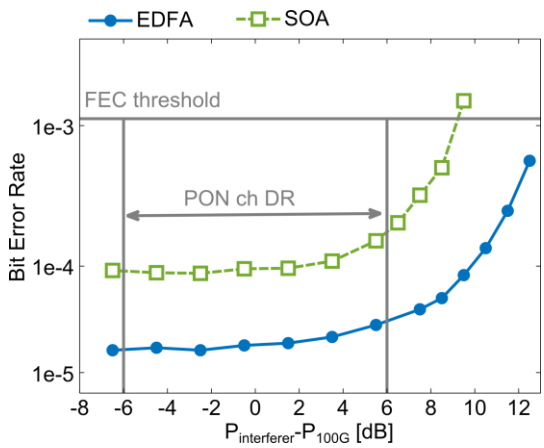

Fig. 8. Performance of the 100G link in upstream as a function of the power of the neighboring OOK 10Gb/s PON channel for the two amplifier node designs.

In the upstream direction the OOK PON traffic, which is interfering with 100G channel, presents bursts with different power. In this specific design, the non-uniform ODN loss can cause up to 12dB of dynamic range between burst powers. By system design the 100G channel power is set in the middle of the PON channels DR. For this reason the 100G channel was characterized in terms of impairments caused by the non-linear cross talk from two neighboring 10G OOK PON channels bursting with high power. The BER of the 100G upstream link was measured as a function of the power of two 50GHz-spaced interfering channels operated with 2µs bursts and 2µs gaps overlapped in time to provide a worst case for the non-linearity. The results in Fig. 8 show that the 100G channel can work below FEC threshold with bursts in neighboring channels with DRs larger than the 12dB caused by the non-uniform ODN loss, for both AN designs. Similarly to the downstream case, the non-linear impairment mechanism is different for the EDFA and SOA cases. For the EDFA-based design, XPM in the backhaul fiber is the major source of impairment. The BER is almost an order of magnitude lower due to the 5.5dB lower power launched in the backhaul fiber: +1dBm per channel, compared to +6.5dBm per channel launched power in the ODN fiber in the downstream case. On the other hand, the performance of the SOA-based design is mostly degraded by the cross-gain and coupled phase modulation inside the SOAs. The nominal SOA output power of the OOK interfering channels, which we set in the middle of their DR, is also ~3.5dB lower compared to downstream case, which reduces the impact of the non-linear cross-talk.

We believe that the case analyzed with two neighboring high power interfering channels is a realistic worst case. While in the XPM case channels spaced by 100GHz or more provide only a negligible contribution, in the SOA case the effect would be roughly the same within the entire gain bandwidth of the SOA. For this reason the absolute worst case for the SOA-design would be if all the 39 remaining OOK PON channels are bursting at the highest power in the DR. However, this is an extremely unlikely scenario due to the statistical distribution of the ODN losses and due to the randomness of the traffic distribution in the various PON channels. The influence of the bursting PON traffic from all the upstream channels on the amplifier gain was also characterized by ON-OFF modulating the ASE emulating the additional PON upstream channels with a repetition period of 20µs. The impact is negligible for both AN designs and we have not measured any appreciable penalty.

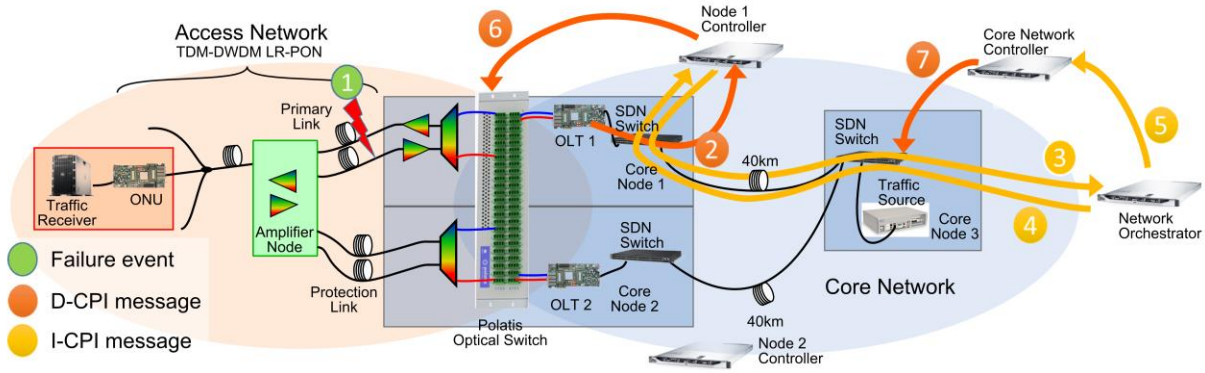

Protection Event
Report Receive Time (in seconds) Event Time (in seconds) Event <Interface:Source:Destination:message>
timestamp=1469191497.680592, msg=Logger 1469191497.680018 ExtEvent:AppController:OptiSwitch:status:{BreakingFibreNow}[1]
timestamp=1469191497.703878, msg=Logger 1469191497.703548 D-CPI:OLT:NodeController:FailureDetected[2]
timestamp=1469191497.705945, msg=Logger 1469191497.705533 I-CPI:NodeController:Orchestrator:FailureDetected[3]
timestamp=1469191497.709295, msg=Logger 1469191497.708944 I-CPI:Orchestrator:NodeController:Invoke_failover[4]
timestamp=1469191497.709388, msg=Logger 1469191497.708990 I-CPI:Orchestrator:CoreNetworkController:Invoke_failover[5]
timestamp=1469191497.714975, msg=Logger 1469191497.714589 D-CPI:NodeController:OpticalSwitch:Patch_Connect[6]
timestamp=1469191497.715073, msg=Logger 1469191497.714651 D-CPI:CoreNetworkController:SDNSwitch:add-flow[7]
timestamp=1469191497.733523, msg=Logger 1469191497.733162 D-CPI:CoreNode:CoreController:Patch_Confirm Data path recovered

Client Recovery
timestamp=1469191497.753103, msg=Logger 1469191497.747648 I-CPI:Orchestrator:NodeController

ONU status report
timestamp=1469191499.402748, msg=Logger 1469191499.402222 D-CPI:SecondaryOLT:NodeController:status Final status report

Fig. 9. Message exchange overview (top) and caption of control plane traces (bottom) for the protection experiment using UNIX Epoch standard timestamps.

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

C. Protection Experiment

The protection experiment demonstrates a dual-homed LR-PON protection mechanism where backup OLTs are shared among PONs in an N:1 scheme [19], and the service restoration is provided over an end-to-end SDN. The system carried out an initial phase of path-precomputation, where it sets up a backup path associated to the failure of a specific PON. The pre-calculation considers the input and output ports at the optical switch, the flow table configuration of the OF SDN switch (both access and core) and the configuration of the OLT flow table. For additional details of the PON activation mechanism operated by the OLTs and ONU, the reader can refer to [19].

The first test of the protection experiment is the failure event emulated by using the optical switch to simulate a fiber cut in the backhaul fiber link between the primary OLT and the AN. Silence in the upstream activates a countdown timer in the primary OLT, which on expiry generates a failure detection and an in-band alarm to the controller of node 1. The duration of this timer takes into account all normal silences on the PON due to the 1.25ms quiet windows and 1.25ms roundtrip time over maximum distance supported by the protocol of 125km, for a total of 2.5ms. The node 1 controller then alerts the overarching NetO which calculates a path to restore services to the ONUs according to its knowledge of the full end-to-end topology covering the core and access networks. The NCs are then instructed by the NetO to provision the protection path through the backup OLT. Fig. 9 gives a logical view of the messages exchanged by the control plane during protection, showing also a capture of the message flow for one of the protection experiment runs.

The results of the service restoration time for the SDN control plane based protection mechanism are shown in Fig. 10. The average restoration time over 70 measurements was 41ms. In Fig. 11 we show the breakdown of the various timings that comprise the 41ms protection figure. The hardware monitoring at the OLT can detect a failure in the network in about 2.5ms. A further 1ms is taken for the alarm packet to be created and sent to the CN switch. The time needed by the protocol to re-establish downstream

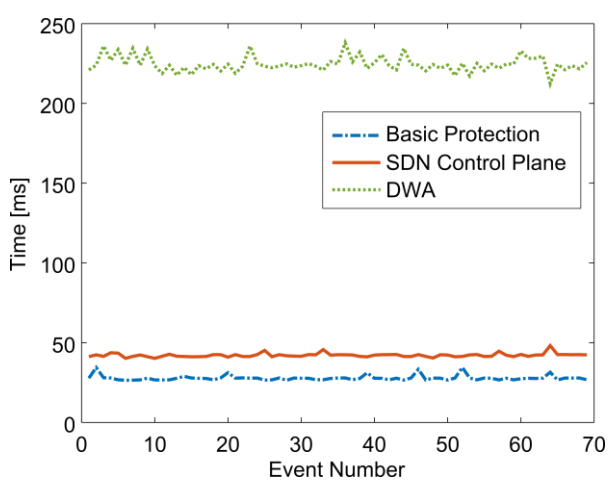

Fig. 10. Service restoration time for the protection mechanism and the DWA through the implemented SDN control plane.

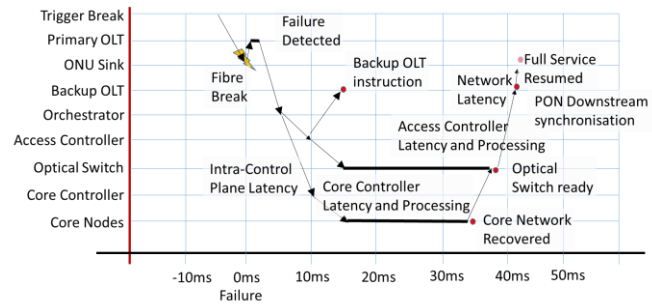

Fig. 11. Breakdown of SDN control plane protection timings.

synchronization is between 2-3ms. From our previous work [23] we know that some time may be needed to re-range the ONUs in addition to the synchronization time (between 2 and 4ms), however in this work we assume that ranging to the backup OLT can be done during normal operation of the PON. Intra-control plane communication is done through a dedicated network with typical latencies. The network latencies between both the OLT and NetO and the NetO and the NCs are emulated in the test-bed and set at 4ms each. The latency and the processing times for both the NCs is also emulated as 5ms each. The core network recovery happens in parallel to the access network recovery time. Accordingly, within 15ms of the failure, the optical and electronic switch components and the backup OLT have been instructed to reconfigure their protection paths. Within 33ms after the failure, the electronic switch components within the core and access are configured, and by 38ms, the optical switch component is configured.

In order to understand the effect of centralising both the NetO and the NCs, we compared the above results with the case where orchestrator and controllers are collocated within the CN. This was accomplished by setting the emulated intra-control plane latencies at zero. The results are shown in Fig. 10 as the basic protection line. On average, basic protection can be accomplished within 27.8ms.

D. DWA Experiment

The DWA use case exemplifies how capacity constraints in one PON channel may be overcome by re-allocating dynamically one or more end user ONUs to a different channel in order to assure quality of service. This could also be used for the opportunistic provision of high bandwidth services (on-demand video and big data transfers), to specific PON users on a dynamic basic. Since the DWA use case is aimed at capacity provision, the wavelength and service reconfiguration times targeted are in the region of a few hundred ms. It should be noted that although the wavelength assignment is not carried out at the granularity of individual burst transmission, we still refer to the system as a dynamic wavelength assignment as the change in wavelength is dynamically and automatically allocated by the controller as a response to an increase in user capacity, rather than being statically assigned by the network management plane.

The experiment setup and a capture for the control plane message exchange is reported in Fig. 12. As in the protection use case, the NetO orchestrates the provisioning of the new path, according to its knowledge of the full end-to-end topology, by instructing the NC. The NC instructs both the

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

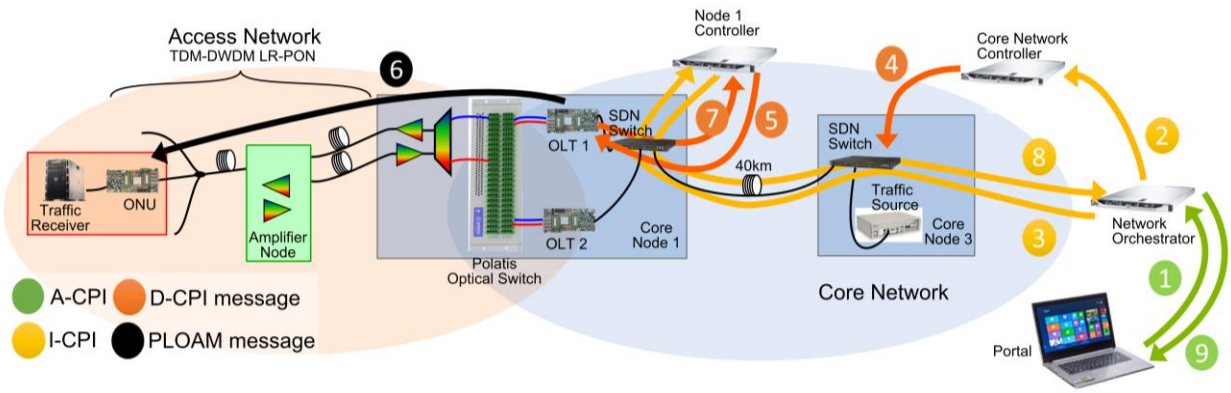

Timestamp	Type	Source	Destination	Message	Event Number
1458130894.036	A-CPI	Portal	Orch.	Resouerce_Request	1
1458130894.243	I-CPI	Orch.	CNC	Create_Path	2
1458130894.251	I-CPI	Orch.	NC	Create_Path	3
1458130894.265	D-CPI	CNC	SDN_switch	add-flow	4
1458130894.300	**MEAS**	**---**	**---**	**Traffic_Measurement_Break**	**4.a**
1458130894.334	D-CPI	NC	OLT	Set_DS_lambda	5
1458130894.402	PLOAM	OLT	ONU	Set_Wavelength	6
1458130894.521	**MEAS**	**---**	**---**	**Traffic Restored**	**6.a**
1458130894.702	D-CPI	OLT	NC	Status_Report: onu_registration	7
1458130894.704	I-CPI	NC	Orch.	Create_Path_Confirmation	8
1458130895.005	A-CPI	Orch.	Portal	Resource_Confirmation	9

Fig. 12. Message exchange overview (top) and caption of control plane traces (bottom) for the DWA experiment using UNIX Epoch standard timestamps.

Polatis optical switch and the secondary OLT to provision a new wavelength. The OLT provides a North Bound interface, through which the control plane can tune the OLT transceivers to a given wavelength. Since the ONU is remote from the control plane, the request to tune to a different channel is also performed through the OLT interface by the invocation of a custom LR-PON protocol message. Using a custom implemented PLOAM message, the primary OLT requests the ONU tune to the wavelength provisioned by the secondary OLT. The NC then acknowledges that it has completed the provisioning of the path through the access portion of the network. In the ONU, DWA is implemented by controlling the SFP+ transceiver tunable laser using an i2c bus and controlling the tunable filter through a UART interface. The lower part of Fig. 12 also shows the timing of some of the main events associated with the DWA experiment.

The DWA results in Fig. 8 refer to the service restoration time when, in response to an increase in traffic demand, the NetO instructs the core and the NC to provision the new path and the ONU traffic is moved to a different PON channel , i.e. from the probing event 4.a to the probing event 6.a. We believe that the average measured provisioning time of 225ms could be reduced by an optimized design of communication interfaces between the ONU FPGA and the tunable components.

V. Conclusion

A dynamically reconfigurable TDM-DWDM LR-PON that can support the convergence of multi-service traffic (10Gb/s residential, dedicated 100Gb/s business and wireless fronthaul) has been demonstrated for up to 100km reach, 512 users and emulated system load of 40 channels. A 100G dedicated channel and a wireless fronthaul channel were demonstrated co-existing alongside 10G PON channels. Two different amplifier node designs are presented employing either burst mode capable EDFAs or SOAs. The demonstrated architecture also integrates an SDN control plane for the access and core network elements, showing a fast protection mechanism, in the case of primary backhaul link failure, with service restoration and the dynamic reassignment of an ONU wavelength in response to increased traffic demand.

Acknowledgment

The authors would like to thank Oclaro for the loan of the 100G transponder.

References

[1] M. Ruffini et al., "DISCUS: an end-to-end solution for ubiquitous broadband optical access," IEEE Commun. Mag., vol. 52, pp. S24-S32, Feb. 2014.

[2] ITU-T recommendation G989.2

[3] G. Talli and P. D. Townsend, "Hybrid DWDM-TDM long-reach PON for next-generation optical access," J. Lightw. Technol., vol. 24, pp. 2827-2834, July 2006.

[4] M. Ruffini, et al., "Software defined networking for next generation converged metro-access networks," Optical Fiber Technology, Volume 26, Part A, Pages 31-41, Dec. 2015.

[5] A. Nag, et al., "N:1 protection design for minimising OLTs in resilient dual-homed long-reach passive optical network," Optical Fiber Communication Conf. (OFC), 2014, Paper Tu2F.7.

[6] G. Talli et al., "Demonstration of SDN Enabled Dynamically Reconfigurable High Capacity Optical Access for Converged Services," Optical Fiber Communication Conf. (OFC), 2016, Postdeadline Paper Th5B.1.

[7] K. Taguchi et al., "Field Trial of Long-Reach and High-Splitting λ-Tunable TWDM-PON," J. Lightw. Technol., vol. 34, pp. 213-221, Jan. 2016.

[8] B. Schrenk et al., "Demonstration of a Remotely Dual-Pumped Long-Reach PON for Flexible Deployment," J. Lightw. Technol., vol. 30, pp. 953-961, April 2012.

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2016.2604864, Journal of Lightwave Technology

[9] P. Ossieur et al., "A 135-km 8192-split carrier distributed DWDM-TDMA PON with 2×32×10 Gb/s capacity," J. Lightw. Technol., vol. 29, pp. 463-474, Feb. 2011.
[10] R. Bonk et al., "Wavelength-transparent long-reach-high-split TWsDM-PON utilized by a non-gated parallel cascade of linear SOAs," European Conf. on Optical Communication (ECOC), 2014, Paper P.7.4
[11] A. Kaszubowska-Anandarajah et al., "EDFA transient suppression in optical burst switching systems," 2012 14th International Conference on Transparent Optical Networks (ICTON), Coventry, 2012, pp. 1-4.
[12] ITU-T recommendation G987.3
[13] C. Antony et al., "High extinction switching of SOAs for in-band crosstalk reduction in PON," Electron. Lett., vol. 44, pp. 872-873, July 2008.
[14] P. Ossieur, et al., "A 10 Gb/s Linear Burst-Mode Receiver in 0.25µm SiGe:C BiCMOS," IEEE J. Solid-State Circuits, vol. 48, pp. 381-390, Feb. 2013.
[15] S. Porto, et al., "Demonstration of 10 Gbit/s burst-mode transmission using a linear burst-mode receiver and burst-mode electronic equalization [invited]," J. Opt. Commun. Netw, vol. 7, pp. A118-A125, Jan. 2015.
[16] ONF, SDN architecture, TR-502, Issue 1, June 2014.
[17] R. Munoz, et al., "Transport Network Orchestration for end-to-end Multi-layer Provisioning Across heterogeneous SDN/OpenFlow and GMPLS/PCE Control Domains," J. Lightw. Technol., vol. 33, pp. 1540-1548, April 2015.
[18] BroadBandForum, Proposed OpenFlow modification in WT-358, document bbf2016.639.02, June 2016.
[19] 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. To appear in J. Lightw. Technol., pre-publication version available at **http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7517316&filter%3DAND%28p_IS_Number%3A4357488%29**
[20] J.M. Gran Josa et al., "End-to-end Service Orchestration From Access to Backbone," Optical Network Design and Modeling Conference (ONDM), May 2016
[21] R. Vilalta et al., "The Need for a Control Orchestration Protocol in Research Projects on Optical Networking," European Conference on Networks and Communications (EuCNC), July 2015
[22] M. C. Jeruchim, P. Balaban, and K. Sam Shanmugan, "Simulation of Communication Systems," Second Edition, New York, Kluwer Academic/Plenum, 2000.
[23] S. McGettrick et al., "Improving hardware protection switching in 10Gb/s symmetric long reach PONs," Optical Fiber Communication Conf. (OFC), 2013, March 2013, pp. 1–3.

0733-8724 (c) 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.