Demonstration of SDN Enabled Dynamically Reconfigurable High Capacity Optical Access for Converged Services

Giuseppe Talli(1), Stefano Porto(1), Daniel Carey(1), Nicola Brandonisio(1), Alan Naughton(1), Peter Ossieur(1), Frank Slyne(2),Seamas McGettrick(2), Christian Blum(2), Marco Ruffini(2), David Payne(2), Rene Bonk(3), Thomas Pfeiffer(3), Nick Parsons(4), Paul Townsend(1)

(1)Tyndall National Institute, University College Cork, Cork, Ireland, (2)Trinity College, Dublin, Ireland,
(3)Nokia Bell Labs, Stuttgart, Germany, (4)Polatis Ltd., Cambridge, United Kingdom
giuseppe.talli@tyndall.ie

Abstract: A dynamically reconfigurable TDM-DWDM PON for converged multi-services at the physical layer (10G, 100G and wireless fronthaul) has been demonstrated for the first time implementing end-to-end SDN management of the access and core network elements.
OCIS codes: (060.2330) Fiber optics communications; (060.4510) Optical communication

1. Introduction

Dynamically reconfigurable time-division multiplexing (TDM) dense wavelength division multiplexing (DWDM) long-reach passive optical networks (LR-PONs) combined with a flat optical core offer the potential to support the ubiquitous delivery of high speed services to customers regardless of type or location [1]. Figure 1 illustrates the network concept and demonstrator, which incorporates a software defined networking (SDN) controlled flat optical core and a TDM-DWDM LR-PON connected to a primary core node, with a protection link to a secondary core node. 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 users, which share a 10G PON channel, to business users, with options to rent dedicated 10G PON channels or high capacity 100G point-to-point (P2P) 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 amplifier node (AN) and allocating a dedicated wavelength from the AN to the remote site (e.g. a CPRI channel). 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 [2]. This work substantially extends previous demonstrations [2, 3], presenting for the first time, to the best of our knowledge: 1) a dynamically reconfigurable TDM-DWDM PON physical layer for 512 split, 40+40 channels in up- and downstream with 100km reach using a 10Gb/s linear burst-mode receiver (LBMRx) and protocol implementation in the optical line terminals (OLTs) and ONUs using field programmable gate arrays (FPGAs); 2) burst-mode capable erbium doped fiber amplifier (EDFA) and semiconductor optical amplifier (SOA) ANs capable of carrying heterogeneous services and modulation formats; 3) physical integration of an SDN controlled access and core network; 4) SDN-enabled fast protection mechanism and end-to-end service restoration in case of a primary link failure; 5) SDN-enabled dynamic wavelength allocation (DWA) in response to an increased traffic demand.

2. Test-bed Description and Results

Figure 2 shows the details of the LR-PON physical layer, including the two amplifier technologies used in the AN. EDFAs, in Fig. 2(a), are readily available and a well-known platform for amplification in C-band, while SOAs, in Fig. 2(b), are a possible solution for integrated and compact ANs and also for extended operation outside the C-band [4]. The figure illustrates how the demonstration setup is constructed so that the two AN designs could be interchanged and also shows details such as amplifier gains and channel powers for both designs. Attenuators are added at various points in the system to emulate end of life standard single mode fiber attenuation (0.3dB/km) and realistic splitter losses in the optical distribution networks (ODNs), including excess loss [5]. Part of the total split (4×4 or 2×4) is located in the AN, and this also provides access to a redundancy path in the backhaul link for

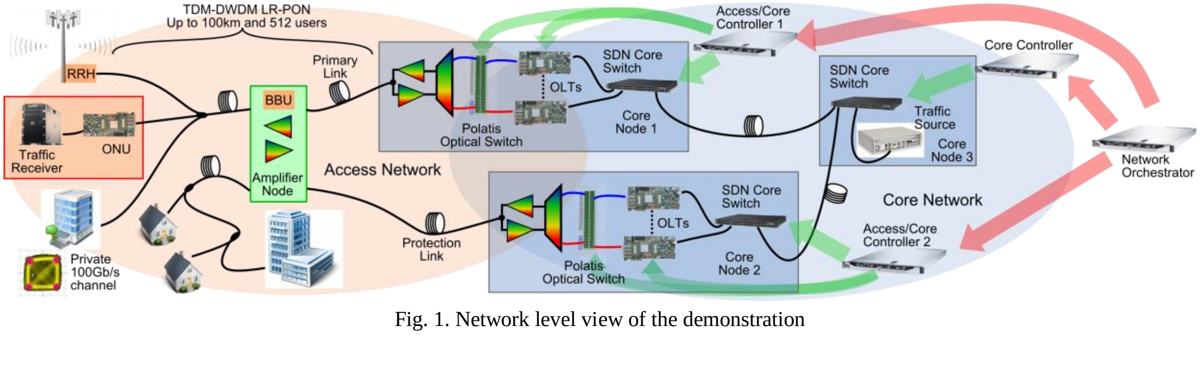

Fig. 1. Network level view of the demonstration

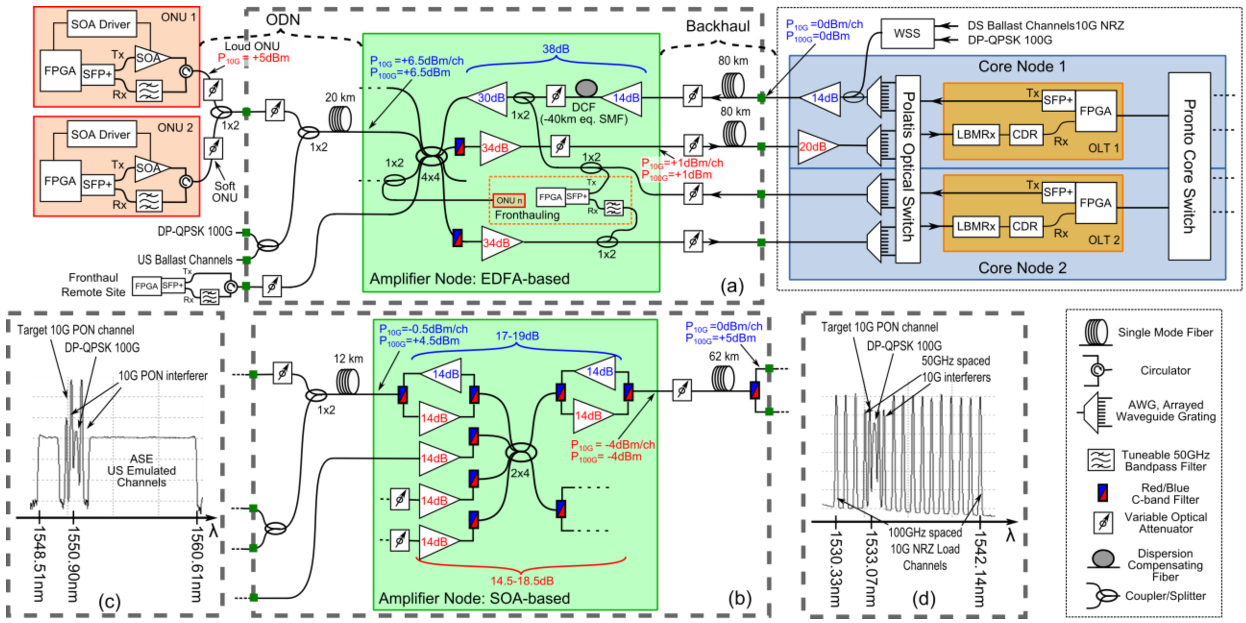

Fig. 2. Physical layer experimental setup with (a) EDFA-based and (b) SOA-based amplifier node

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. Two OLTs and two ONUs are fully implemented, and additional traffic equivalent to a power of 40 channels in both up- and down-stream is added to fully load the network. As indicated in Fig. 2(d) the additional downstream channels are emulated using 10Gb/s on-off keying (OOK) externally modulated DFB lasers. Figure 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 modulated to mimic the typical burst mode power variations of upstream channels. Similar entry points allowed the insertion of the 100G P2P channel, which is realized using a commercial transponder to generate and receive a dual-polarisation quadrature-phase-shift-keying signal. The OLT and ONU transmitters are commercial tunable 10G SFP+ transceivers, while the wavelength selection at the receivers is achieved by using a 50GHz tunable filter at the ONU and by routing the ports of a wavelength demux at the OLT using an Openflow-enabled Polatis optical switch, partitioned and shared by the two core nodes. The LBMRx at the OLT [6] is followed by a static electronic dispersion compensation-clock data recovery (CDR) module [3]. The fronthaul channel, demonstrated only for the EDFA case, is emulated by adding an SPF+ transceiver driven by an FPGA in the AN and in one ODN arm as shown in Fig. 2(a). The SDN controller architecture, 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 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 [2]. 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. A simplified core network is emulated using 2×40km fiber links and 3× Openflow bridges based on a Pronto 3780 switch with 10GbE interfaces.

Figure 3 (a) shows the bit error rate (BER) measured as a function of the ODN loss for the downstream direction for the two AN designs. Assuming a forward error correction (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 and to an overall split ratio of 512 (128×4 due to the additional AN split). The difference in the ONU's receiver sensitivities is ~2dB, while the difference in performance between the two designs of ~8dB roughly corresponds to the 7dB difference in power launched from the ANs in the ODN. Figure 3(b) presents the upstream BER measured at the OLT in burst mode operation on 2μs bursts generated by the ONUs. The ODN loss was varied from 16 to 35dB only for ONU 2 in order to vary the dynamic range (DR) of the burst powers from the two ONUs at the LBMRx. The ONU 1 burst power is maintained constant at close to the LBMRx overload power (loud ONU) with an ODN loss of 16dB. The BER is always below the FEC threshold for ONU 1 (below 1e-12 for most of the range in the EDFA case). 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 (up to 12dB), with the EDFA AN showing slightly better performance due to the lower noise figure (NF). The fronthaul channel can also operate error free (BER<1e-12) in both directions for an ODN loss of up to 28dB (Fig. 3). The performance of the 100G P2P link was characterized in downstream as a function of the ODN loss using an EDFA pre-amplifier at the receiver with 5.5dB NF. As shown in Fig. 4(a) both AN designs can support the 100G channel with ODN loss higher than 28dB (equivalent overall to 128×4=512 users). To account for the burst nature of the interfering OOK PON traffic in the upstream, the 100G channel, which has a power in the middle of the PON DR, was characterized 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. 4(b) show 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. The performance variation for the EDFA and SOA cases is due to different non-linear impairment mechanisms: cross-phase modulation in the backhaul and ODN fiber for the EDFA case, and cross gain modulation in the SOA.

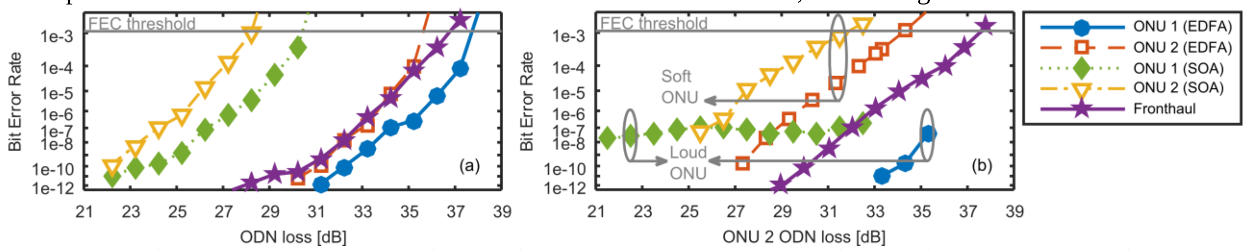

Fig. 3. BER of 10Gb/s PON downstream as a function of the ODN loss (a) and upstream as a function of ONU 2 ODN loss (b) for the two amplifier node designs and BER of the fronthaul channel as a function of the ODN loss.

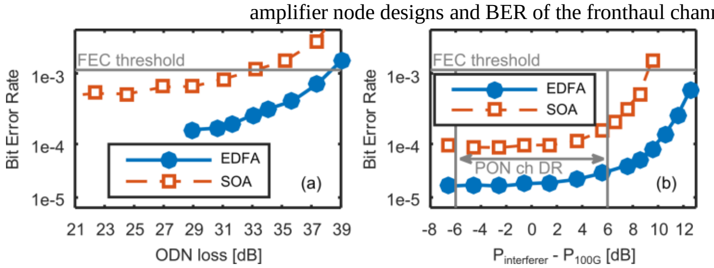

Fig. 4. Performance of the 100G link (a) in downstream as a function of the ODN loss and (b) in upstream as a function of the power of the neighboring OOK 10Gb/s PON channel for the two amplifier node designs.

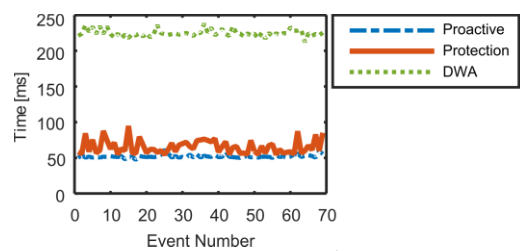

Fig. 5. Service restoration time for the protection mechanism and the DWA through the implemented SDN

Figure 5 shows the service restoration time for the protection mechanism, where backup OLTs are shared among PONs in an N:1 scheme. The baseline time between the two paths of ~50ms is given when the switchover is proactively triggered by the controller, without waiting for a failure event. In contrast, the protection results show the restoration time when a failure event is caused by a cut in the backhaul link between the primary OLT and the AN. Silence in the upstream activates a countdown timer in the primary OLT controller, which on expiry generates a failure detection and an in-band alarm to the Openflow access NC. The access NC alerts the NetO, which provisions the protection path. The average protection time is measured at 64ms, with variations between 50 and 100ms attributed to the random delay in the failure detection. The DWA results in Fig. 5 refer to the service restoration time when, in response to an increase in traffic demand, the NetO instructs the core and the access NCs to provision the new path, according to its knowledge of the full end-to-end network topology and the ONU traffic is moved to a different PON channel. Using a custom implemented physical layer operations, administration and maintenance (PLOAM) message, the primary OLT requests the ONU to tune to a wavelength provisioned by the secondary OLT. We believe that the measured provisioning time of ~225ms could be reduced by an optimized design of communication interfaces between the ONU FPGA and the tunable components.

3. Conclusions

A dynamically reconfigurable TDM-DWDM LR-PON for converged multi-service traffic (10Gb/s residential, 100Gb/s P2P business and wireless fronthaul) which also integrates an SDN control plane for the access and core network elements has been demonstrated for up to 100km reach, 512 users and 40 channels, showing a fast protection mechanism with service restoration and the DWA of an ONU in response to increased traffic demand.

Financial support from Science Foundation Ireland (SFI) (grants 12/IA/1270 and 10/CE/I1853), and EU FP7 project DISCUS (grant CNECT-ICT- 318137) is gratefully acknowledged. The Authors would also like to thank Oclaro for the loan of the 100G transponder.

4. References

[1] M. Ruffini et al., IEEE Comm. Mag., vol.52, pp.S24,S32, 2014.
[2] M. Ruffini et al.,Opt. Fiber Tech., Vol. 26, pp 31-41, Dec.2015
[3] C. Antony et al., OFC 2010, PDPD8
[4] R. Bonk et al., ECOC 2014, P.7.4
[5] ITU-T recommendation G989.2
[6] S. Porto et al., OFC 2014, M3I.5