Experimental End-to-End Demonstration of Shared N:1 Dual Homed Protection in Long Reach PON and SDN-Controlled Core

Séamas McGettrick, Frank Slyne, Nattapong Kitsuwan, David B. Payne and Marco Ruffini
CTVR, The Telecommunications Research Centre, Trinity College Dublin,
smcgettrs@tcd.ie, fslyne@tcd.ie, kitsuwan@tcd.ie, david.b.payne@btinternet.com, marco.ruffini@tcd.ie.

Abstract: We demonstrate fast restoration of PON services with a dual-homed, shared-OLT protection mechanism. Using the SDN-controlled pan-European GEANT core network, we demonstrate PON protection and end-to-end service restoration times within 40ms and 155ms respectively.
OCIS codes: (060.4250) Networks; (060.4261) Networks, Protection and Restoration

1. Introduction

In this paper we present and evaluate a dual-home Long-Reach Passive Optical Network (LR-PON) protection mechanism where backup OLTs are shared among PONs in a N:1 scheme, and the service restoration is provided over an end-to-end Software Defined Network (SDN) controlled core network.

Protection mechanisms are essential to ensure adequate network availability. However, providing redundancy adds costs to the network and so protection mechanisms are mainly limited to the core and metro networks where the cost of adding redundancy to the network can be shared among many users. Future Long Reach Passive Optical Networks (LR-PON) have larger split ratios and longer reach compared with current PON systems which means that portions of the metro transmission network will be bypassed [1]. The likelihood of failure is thus higher due to higher probability of the long feeder fibers being cut. The LR-PON feeder fibers are replacing part of the current network that is usually protected and for this reason protection mechanisms become a requirement in LR-PON.

In our previous work [2], we looked at 1+1 protection mechanisms in the LR-PON and showed that the PON protocol together with hardware optical monitoring could re-establish control of all ONUs on the PON in as little as 4 - 5ms. However, this required downstream data to be replicated at both the primary and backup OLT. Although this method reduces protection times significantly, it also increases the cost of the network as extra capacity is needed to duplicate the downstream traffic. In [3], we removed the need for duplicated data by moving to a 1:1 protection scheme and found that this scheme could protect the network in approximately 125ms, although we also made propositions to reduce this further down to about 40 ms.

In the DISCUS architecture [4], we wish to provide cost effective end-to-end solutions for ubiquitous optical access networks. To that end, in this paper we investigate the more realistic N:1 protection schemes, where each backup OLT is shared among a number of primary OLTs. In the event of a failure the backup OLT is allocated to the failed PON and data is re-routed to the backup node. Furthermore, the backup OLT is located at a different location to the primary OLT (i.e., dual-homed to a different Metro-Core node) and utilises a backup feeder fiber path. Thus, unlike co-located backup OLTs which only protect the network from OLT board failure, this scheme protects against fiber dig-ups, OLT failure and metro node failure.

To test the scheme we have partially implemented the OLT and ONU units for LR-PON on Xilinx VC709 FPGA boards. We connect our PON system to the GEANT optical testbed [5] to act as the core network for our experiments. We use this setup to measure that end-to-end protection time after a failure occurs. In 2008 a similar experiment was carried out using commercial GPON hardware and the restoration time was found to be in the order of 30s [6]. Additional experiments results were published in 2013 [7], showing an automated protection mechanism based on VLANs that reduced access protection times to 4.5 s (with maximum values of 9.5 s). As we have implemented the OLT and ONU on FPGA and are using an open flow core network we have full access/control to all elements in our system. Furthermore, The XGPON standard – on which our LR-PON protocol is based - has greatly reduced the time required to register and range the PON [8]. The new protocol together with the accessibility and flexibility of our test bed means that we can show that end-to-end protection can be achieved in 155ms and we can give a breakdown of where this time is lost in the protection switch over.

2. Experimental Setup

This experiment combines the optical architecture testbed in Trinity College Dublin and the GEANT pan-european research network as shown in Figure 1. The testbeds are connected through two dedicated GE links. Although this link is well below the 10Gb capacity of the LR-PON, it does allow us to reliably evaluate latency effects between diverse network elements and the higher-level control layers.

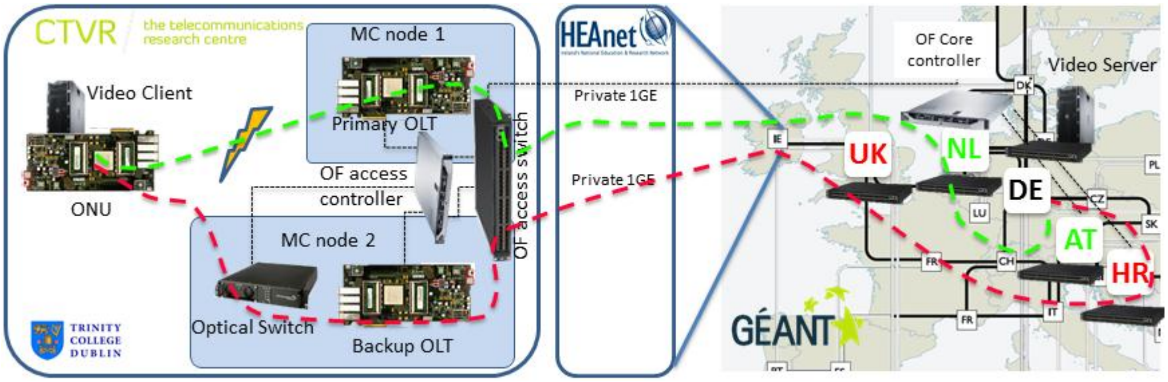

Figure 1 Logical view of combined LR-PON access and SDN Core network

The GEANT network is setup in a five-node network topology with nodes in Netherlands (NL), Germany (DE), Austria (AT), Croatia (HR) and United Kingdom(UK). A server colocated at the DE node acts as the source for data in this experiment. The primary data path in the core is between nodes DE , AT and NL (shown in green in Figure 1). The backup path between nodes DE, HR and UK (shown in red in Figure 1). We implement two paths to emulate dual-homed PON network where the primary and the backup OLT are in different locations. Data on the primary data path is routed to the Primary OLT and data on the backup link is routed to the backup OLT. The NL node also hosts the Openflow core network controller which can be used to control which of the two paths data takes to our access network.

The access network comprises of a pronto 3780 switch, running release 2.2 (Openflow v1.3 compatible firmware), three Xilinx vc709 development boards [9], acting as primary OLT, backup OLT and ONU, a glimmerglass optical switch, A Dell T620 with 10G SFP cards acting as client machine attached to the ONU and a separate Dell R320 Server acting as access Openflow controller. The Proto switch is configured as multiple virtual bridges to act as standalone bridges each with a seperate Openflow controller. These are connected to a gateway machine on the core side of the network and one of the Primary or backup OLTs on the access side.

The PON hardware is implemented over three Xilinx VC709 boards [9]. Two of the boards act as the Primary and backup OLT and the final board acts a single common ONU. The PON hardware is a partial implementation of the XGPON standard, modified where neccessary to allow for the longer reach and higher split ratios needed by the Discus architecture. The backup OLT is connected to the access network through the optical switches to allow switching between various PONs, as would be the case in the N:1 protection scenario. The optical switch is a MEMs-based Glimmerrglass 16x16 port with nominal switching time of 25 ms. In N:1 protection the backup OLT can be assigned any 1 of N PONs depending on where the failure occurs. In our tests we establish a connection to the PON when instructed to do so by the access Openflow controller.

3. Failure Detection and Data Redirection

In our test scenario the feeder fibre between the primary OLT and first stage splitter on the PON is cut. This stops all upstream and downstream data on the Primary link. A hardware unit in the primary OLT FPGA monitors the upstream data. On a LR-PON of 125 Km, like the one proposed by the DISCUS project, failure detection could take approximately 2.5ms in worst case conditions (i.e 1.25ms for fiber delay and 1.25 ms for quiet window).

When a failure is detected, the hardware detection unit alerts the OLT controller which sends an in-band upstream alarm to the Openflow access network controller. Since the Openflow bridge does not terminate the optical path, it is not possible for Openflow path switching rules to be invoked directly by the interception of this alarm, which remains solely in the data plane. Instead the alarm is forwarded to the higher layer control infrastructure by the OLT. The Openflow access network controller notifies the optical switch on the backup OLT path to tune to the failed PON and finally the management controller notifies the backup OLT to take control of the PON. The PON meanwhile has entered the Loss of Downstream Sync state and will remain there for 100ms or until the backup OLT begins to send synchronization words downstream. If the backup OLT does not take over before the 100ms time out the entire PON will have to be reactivated, and re-ranged to resume transmitting data. Once the backup OLT has taken control of the PON the PON is ready to start receiving data again.

The Openflow access network controller also passes a message to the core network Openflow controller. This causes the core network to redirect data from the primary path to the backup path to emulate more closely the dual homed nature of the Long-Reach PON.

4. Results

The restoration time of the network can be split into three distinct phases. These are the time it takes to detect a failure on the network, the access network recovery time and the core network recovery time. The hardware monitoring at the OLT can detect a failure in the network in about 2.5 ms. The access recovery time was measured by a counter at the ONU. The counter is started when the ONU enters the temporary loss of downstream synchronization state and stops when the ONU returns to the operational state. Figure 2 shows the results of 50 repetitions of the protection experiment. We found the access recovery time to be approximately 40 ms, which can be further broken down further into time required to tune the optical switch (30 ms) and time needed by protocol to re-establish downstream synchronization (2 ms). From our previous work [2] we know that some time may be needed to rerange the ONUs in addition to the synchronization time (between 2 and 4 ms), however in this work we assume that ranging to the backup OLT can be done during normal operation of the PON. The remaining time is needed by the local Openflow controller to communicate with the optical switch to connect the backup OLT.

The Optical switch is tuned to the failed PON in parallel with the core Openflow controller being notified of the failure. The core controller begins to reconfigure the network to reroute data to the backup path to the backup node. Once both the access and the core networks have been reconfigured data starts to arrive at the destination again. In these experiments we found end-to-end service restoration times of 155ms on average. This time can be approximately broken down to 10ms failure detection and access OF Controller, 30ms latency between access and core OF controllers, 30ms core OF controller processing time, 20-30ms latency to push backup path rules to core switches and 50ms latency over backup path.

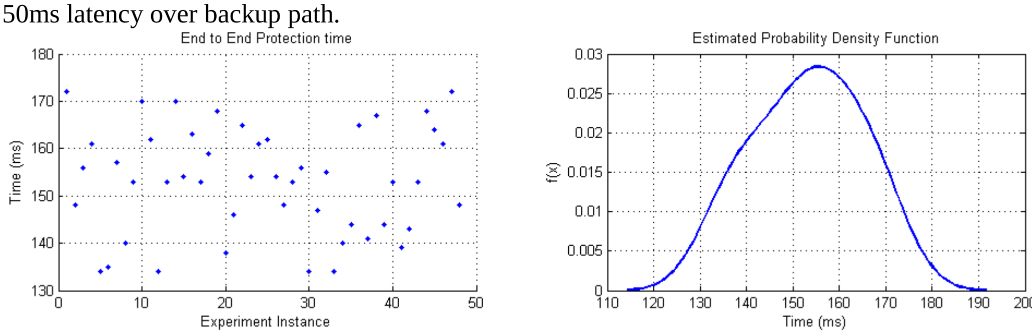

Figure 2 N:1 Protection time for dual homed LR-PON with SDN core

5. Conclusions

In this paper we have described and evaluated a dual-homed N:1 protection scheme for Long Reach PON architecture implemented on FPGA which we combined with SDN Openflow core to demonstrate end-to-end protection. Through our experimentation we have found that N:1 protection in the access network is limited by the speed of the Optical switch to tune to the failed PON. However, this time is generally masked by the longer and more variable time needed to reroute data in the core. We show that with our flexible SDN core and LR-PON access network we can re-establish access connectivity in 40ms and end-to-end connectivity in approximately 155ms.

Acknowledgements

This material is based upon works jointly supported by the Science Foundation Ireland under Grant No: 10/CE/I1853 and European Union seventh Framework Programmed (FP7/2007-2013) under grant agreement N. 318137 (Collaborative project “DISCUS”).

References

[1] M. Ruffini et al., Deployment Strategies for Protected Long-Reach PON, JOCN vol.4, No 2, Jan 2012.
[2] S. McGettrick et al., “Ultra-fast 1+1 Protection in 10Gb/s Symmetric Long Reach PON,” P.6.18, ECOC 2013..
[3] F. Slyne et al.,“Design and experimental test of 1:1 End to End Protection for LR-PON using an SDN mult-tier Control Plane”, Vol. 3 No. 9 ECOC 2014.
[4] M. Ruffini et al.,“Discus: An End-to-End Solution for Ubiquitous Broadband Optical Access”,IEEE Comms Mag., vol. 52, no. 2 Feb. 2014.
[5] GÉANT network. Online: http://geant3.archive.geant.net/.
[6] Kang et al., “Restoration of Ethernet Services over a Dual-Homed GPON System, NWD2, OFC 2008.
[7] A. Rafel et al., “Automatic Restoration over a Type B Dual Parented PON Using VLAN Switching”, We.2.F.1, ECOC 2013.
[8] S. McGettrick et al., “Improving Hardware Protection Switching in 10Gb/s Symmetric Long Reach PONs”, OW3G, OFC 2013.
[9] Xilinx VC709 board. Online: http://www.xilinx.com/products/boards-and-kits/EK-V7-VC709-CES-G.htm