Ultra-fast 1+1 Protection in 10 Gb/s Symmetric Long Reach PON

Séamas McGettrick, Lei Guan, Alan Hill, David B. Payne, Marco Ruffini

CTVR, University of Dublin, Trinity College, Ireland, mcgettrs@tcd.ie (S. McGettrick)

Abstract *In this paper we evaluate PON protection switching times using our FPGA-based hardware implementation of a Long-Reach PON protocol based on the XGPON standard. We first compare reactivation times of different 1+1 protection scenarios, and then we propose hardware modifications to reduce PON protection times.*

Introduction

Protection is an essential feature in modern telecommunications networks. It is utilized to ensure adequate network availability in the event of equipment failures or physical layer damage in the infrastructure. However, since protection requires additional redundant infrastructure in the network, it inevitably adds an extra cost. Thus historically protection mechanisms are mainly utilized in the core or metro networks where the cost can be shared among a large number of users. In traditional copper access networks the cost of implementing protection mechanisms, in the access network, is prohibitively high for the mass of access connections and is therefore unjustified for most users.

However Passive Optical Networks (PONs) have introduced a passive tree structure in the access, where fibres are progressively shared towards the central office among a larger number of users. This concept is further exploited in Long Reach PON (LR-PON), where the longer fibre reach allows an increase in the number of users per PON to over a thousand, in some scenarios. The increased reach also allows the bypass of the metro transmission network¹. The likelihood of a failure is thus higher, due to the higher probability of fibre cable cut in the long feeder fibre. The feeder fibre is also replacing the backhaul or metro access network which is usually protected. For these reasons protection mechanisms become a requirement in LR-PON.

Fig. 1 shows the structure of the proposed LR-PON. The LR-PON implementation used in this work is based on the XGPON² protocol with two major modifications: higher split ratio (up to 1024 users) and longer reach (up to 125 km). The XGPON standard supports two different protection mechanisms: recovery from intermittent loss of downstream synchronization (LODS) and recovery by cold start reactivation. Upon detection of downstream loss of signal, the ONUs enter a LODS state where they remain for a period of up to 100ms. If within this

interval the backup OLT node can re-establish the downstream synchronization with the ONUs, the ONUs can return to the operational state and the PON will remain active with minimal disruption. However, if the backup OLT does not start sending downstream frames before the 100ms LODS state times-out, the ONUs return to an initial state and the backup OLT will be forced to reactivate all ONUs individually.

Furthermore, to ensure that connections of POTS, T1 and E1 legacy services are not dropped during protection switchover, some POTS exchanges require <120 ms frame loss period, and T1/E1 leased lines require 50 ms protection times. It is therefore necessary to achieve protection within a few tens of ms^{4,5}.

Previous tests⁶ based on commercial GPON equipment showed that the software overhead of the alarm monitoring system for protection introduced a delay in the order of 9 seconds. Although this figure seems pessimistic, as the authors did not have access to the inner working of the PON, it gives us an indication of the time it would take to trigger a protection event via software signalling between ONUs and OLTs. Thus if the backup OLT relies on a software alarm to initiate a protection event the response time will be well outside the 100ms window needed for fast protection, and far too slow to hold POTS and leased line connections.

Our previous work³ showed that XGPON has greatly improved cold start protection time (i.e. the second of these two protection scenarios) when compared to GPON. However these protection times could be reduced even further if the backup OLT could take over the PON during the 100ms LODS state.

In this paper we propose a hardware-based 1+1 protection mechanism which operates on the backup OLT to further reduce the protection time. The LR-PON ONUs and OLTs used for the experimental test were implemented on the NetFPGA 10G board⁷. Protection time results are presented for both XGPON protection scenarios. In addition to these results we propose two further more realistic protection

Fig. 1: Block Diagram of OLT and ONU prototypes

scenarios and show how even with less than perfect knowledge of the PON we can still reactivate the PON very quickly.

Experimental Setup

The block diagram showing all the functions required to implement the activation process in the OLT and ONU are presented in Fig. 1. These include frame synchronization and alignment, scrambling, framing, physical layer messaging engines (PLOAM), static bandwidth map allocation and controlling finite state machines. The current hardware does not include an optical channel. The optical channel is emulated using a long delay line implemented in the FPGA fabric. User data is simulated using random data generators. The experiment models a 1+1 protected system where the PON is dual parented at the last splitter level at the optical amplifier node (i.e. feeder fibre protection). Two ONUs are used to represent the users on the PON.

Activation Models

The four protection scenarios investigated in this work are the following.

Cold Start Activation is the most basic form of protection available to the PON. If the protection event takes longer than 100ms to start, the whole PON must be reactivated as if from a cold start. The protocol instructs all ONUs to dump ONU-ID information and ranging information, and to return to the initial state, if the OLT does not start broadcasting valid downstream frames within 100ms of a failure occurring. All ONUs are re-ranged, and receive their individual equalisation delays in unicast Ranging_Time PLOAM messages from the backup OLT. These results are taken from our previous work³ and represent a baseline time for protection after the protection event is initiated.

Ideal Protection is recommended in the XGPON standard. The standard assumes that the backup OLT has up to date information about the PON including active ONUs and ranging times. It also assumes that the backup OLT has prior knowledge of the differential distance between the primary and backup OLT to the first splitter. Since all the ONUs branch after the first splitter, the only information required to range the PON is the difference in

distance between the first splitter and the primary and backup OLT. The entire PON can thus be re-ranged with a single broadcast PLOAM message.

Range One protection scenario is a more pragmatic version of the ideal protection and is the first of the two new protection scenarios introduced in this paper. It is assumed that the backup OLT has up to date information regarding active ONUs; however the backup OLT no longer needs to know the differential distance between the primary and backup feeder fibres. Instead the backup OLT sends a ranging request to a single ONU and uses this information to calculate the differential distance and then provides the delay offset adjustment for the entire PON with a single broadcast PLOAM message.

Flexible OLT protection is the second new protection scenario introduced in this paper. This scheme calculates the differential distance between the primary and backup OLTs and the first splitter during normal PON operation. It can do this because it knows that all the ONUs are synchronized to each other, i.e. the distance from the first splitter to the ONUs remain unchanged and thus the only difference in round trip time is caused by the differential distance between the primary and backup OLT to the first splitter. When a protection event is initiated all ONUs are already in the LODS state and so the OLT knows that no upstream traffic is being sent on the PON. The backup OLT therefore can send a bandwidth map to the ONUs and monitor when the first burst returns. The OLT can then compare this arrival time to the expected arrival time of the first burst. In this way it can calculate the differential distance while broadcasting data.

Hardware Protection monitoring

As discussed in the introduction, software monitoring can be slow and prevent protection switchovers within the required 100ms LODS timer duration, let alone the few tens of ms required for leased lines. One method to ensure that the backup OLT responds quickly to a failure is to allow it to monitor the upstream traffic. The backup OLT will therefore listens to upstream packets on the PON. A fault in the primary OLT can be detected if the backup OLT

Tab. 1: PON protocol settings overview

Name	Value
Max Users	1024
Downstream BW	10 Gb/s
Upstream BW	10 Gb/s
Max Random Delay (MRD)	125 μ s
Reach	100 Km
Fibre Delay (FD)	500 μ s
Response Time (RT)	35 $\pm$ 1 μ s

notices that the upstream data channel has not transported any frame for a time interval longer than the set quiet window. Once some ONUs are active on the PON (i.e. registered, ranged and not in sleep mode), the longest quiet time (QT) on the upstream data path can be given by the equation below; where FD is the fibre delay, RT is the response time and MRD is the max random delay used by the ONUs during activation:

$$QT = 2 * FD + RT + MRD$$

This time corresponds to the quiet window used during the activation process. In a 100 Km PON with a processing delay of 35 $\pm$ 1 μ s this corresponds to a maximum of 1161 μ s. If the upstream data path is idle for longer than this time one of two things has happened. The ONUs have moved to the LODS state because the primary OLT is no longer functional or the upstream fibre for the backup OLT has been cut. In the first case the backup OLT needs to initiate a protection event. In the second case the backup OLT will initiate a protection event but the optical amplifier at the first splitter (which is controlled by a management unit) is turned off and will stop the backup OLT from interfering with the primary OLT signal and an alarm is raised indicating backup OLT or secondary path failure. Using this method the backup OLT can automatically respond to failure events before the software monitoring system detects a failure.

Results

Table 1 shows the parameters used for testing the LR-PON protection mechanism. The primary OLT activates the PON and then resets to mimic a fault in the hardware. This causes the two attached ONUs to enter the LODS state. After 2.6 ms the backup OLT detects the failure and attempts to take over from the primary OLT. The ONUs regain downstream synchronization, return to the operation state, and then wait for ranging and bandwidth grant messages. Fig. 2 shows the time required for the backup OLT to return all ONUs on the PON back to operational state, correctly re-ranged, after the protection event. Cold start protection offers a base line measurement in this diagram as it is the time it would take the PON to activate all the ONUs if the backup OLT had no prior information about

Fig. 2: Reactivation of a PON after protection event.

the PON. In this case, activation on the PON can take approximately 7ms. As expected the fastest method to reactivate the PON after a protection event is the ideal case where the backup OLT has access to up to date information about the state of the PON and knows the differential delay between itself and the primary OLT. In this case reactivation of the PON takes as little as 1.5ms. The range one method assumes that it knows all the ONU IDs on the PON but does not have the differential distance between OLTs. In this more realistic case the reactivation takes in the region of 2.5ms. Finally the flexible method allows the PON to return to normal operation without re-ranging, and thus has the same return to functioning as the ideal case however re-ranging is completed in parallel and can take 2.5ms. These numbers ignore the time taken to raise the alarm. For cold start, the last ONU is out of action for a minimum of 107 ms, whereas for all other methods using our HW monitor are out of action for an additional 2.6 ms, waiting for an alarm from the hardware monitor.

Conclusions

Our results show, that with upstream monitoring, a PON can switch over and re-range to a backup OLT in less than 4ms.

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 Programme (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] ITU-T, 10-Gigabit-capable passive optical network; TC layer specifications, G.987.3 (10/2010)
- [3] S. McGettrick et al., Improving Hardware Protection switching in 10Gb/s Symmetric LR-PONs, OFC 2013.
- [4] ITU-T, 10-Gigabit-capable PON ,G987.1, (10/2010)
- [5] ITU-T, Gigabit-capable PON ,G984.1, (03/2008)
- [6] J. Kang et al., Restoration of Ethernet Services over a Dual-Homed GPON SYSTEM, OFC'08
- [7] NetFPGA consortium, available at: www.netfpga.org