Passive Optical Networks for Particle Physics Applications

S Papadopoulos†‡, I Darwazeh‡,Jan Troska†,I Papakonstantinou†,Csaba Soos†,Francois Vasey† and John Mitchell‡

† European Organisation for Nuclear Research - CERN ‡ University College London

Abstract: This paper explores the possibility of using Passive Optical Networks for particle physics applications. The information that needs to be transferred can be classified in two broad categories; data generated by the experiments and control information. Two different architectures are described; each designed to carry one of the two different types of information. A demonstrator of the network carrying the timing, trigger and control information is currently under test. Architectures for the upgraded data readout, which is going to be subject of future work, is also discussed.

1. Introduction.

In this paper we try to identify appropriate architectures for particle physics applications based on Passive Optical Network (PON) technologies, for possible future deployment at the European Organisation for Nuclear Research (CERN). The networks that are in use at CERN fall in two categories; those that transfer control information and those that transfer raw data, each with different set of requirements. We suggest a customised version of Time Division Multiplexing (TDM) PONs for Timing, Trigger and Control (TTC) information and we discuss the possibility of using Wavelength Division Multiplexing (WDM) PONs for data readout (DAQ) applications. Currently a demonstrator of the upgraded TTC network is under test, while identifying the appropriate architecture for the upgraded DAQ is going to be the subject of future work.

2. TDM PON for TTC applications

2.1. Motivation

CERN's largest accelerator, the LHC, accelerates two counter-rotating beams of protons in order to periodically produce collisions inside the four experiments that are located around its 27km circumference. The collision products are measured by a series of different detector types that effectively take snapshots at specific instants in time, synchronous with the particle collisions. In order to achieve high measurement resolution, accurate synchronization between the accelerator and the devices installed at the detector "front-ends" is required. At the same time the information collected by the detectors is stored in circular buffers and is available to be transmitted in order to be stored and further processed offline. The amount of raw information produced by the detectors is too large to be read-out in its entirety – e.g. the CMS detector front-ends can produce >300Tbps – so it is necessary to filter out data that do not contain useful information based on simple rules and provide Triggers on interesting events. In order to take this decision some feedback information is required. The network that provides the synchronization and triggering information to the detector front-ends is known as the Triggering Timing and Control (TTC) network. In addition, limited feedback is provided by a separate electrical network known variously as a throttling (TTS) or busy system.

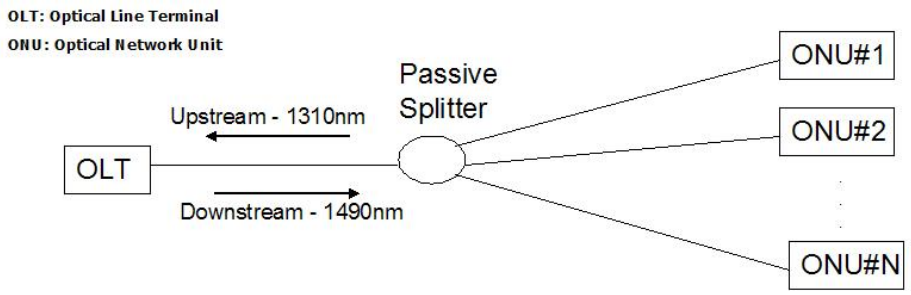

Figure 1 Generic architecture of a TDM PON

The development of a single point-to-multipoint architecture that can support bidirectional transmission over a single link, for the transfer of the triggering, timing and control signals while simultaneously providing a feedback path, is highly desirable in the context of the foreseen luminosity upgrade of the LHC. The thus upgraded accelerator (SLHC) will require upgraded experiments to handle increased numbers of particle collisions. We suggest a new TTC network architecture based on industrial standards for TDM PONs which can provide point-to-multipoint bidirectional connectivity (Figure 1).

2.2. Application requirements vs. TDM PON performance

One of our goals has been the construction of a flexible and cost-efficient network. The benefits that can be delivered by this architecture are directly linked to the number of Optical Network Units (ONUs) that can be served by a single Optical Line Terminal (OLT), thus the splitting ratio. A simple power budget calculation can give an estimation of the maximum splitting ratio that can be achieved. The power budget is affected by several aspects. Transmitted power and receiver sensitivity are two primary factors. Dispersion which is commonly a source of power penalty can be neglected in our case since the maximum length of the system is around 1000m. This also shows that fiber loss, although it should be taken into account is not expected to have a high impact on system performance (it is below 1db). The main sources of power penalty are expected to be the splitting and connector losses.

The following graphs (Fig.1 and Fig.2) are based on theoretical calculation of the power margin vs. the splitting ratio. The following equation has been used for the calculations:

$$M_S = \overline{P}_{tr} - \overline{P}_{rec} - (\alpha_f L + \alpha_{con} + \delta_d)$$

where M_S is the system margin, $\overline{P}_{tr}$ is the average transmitter power, $\overline{P}_{rec}$ is the receiver sensitivity, α_f is the fiber loss per km of fiber, L is the fiber length, α_{con} is the connector loss and δ_b is the dispersion power penalty. [1] The average transmitter power value that has been used in this analysis is the average transmitter power obtained from laboratory measurements of commercially available transmitters that are currently in use to construct the demonstrator. The value that has been used for the receiver sensitivity is the one provided by the manufacturer, while for the rest of the quantities, typical values have been used. [1]

For the downstream a splitting ratio of 1:256 can be achieved, while for the upstream the maximum splitting ratio is 1:512 (Figure 1). In both cases, we have calculated the maximum splitting ratio after leaving a 3db safety margin.

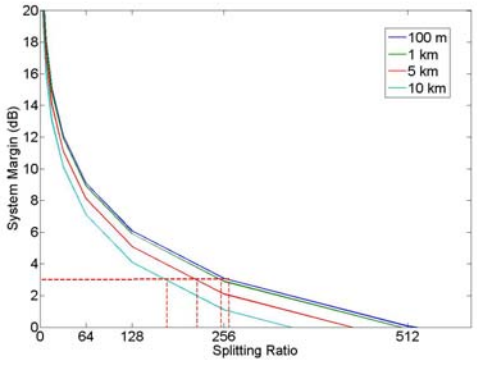

Figure 2 Power Margin vs. splitting ratio for various link lengths – Downstream

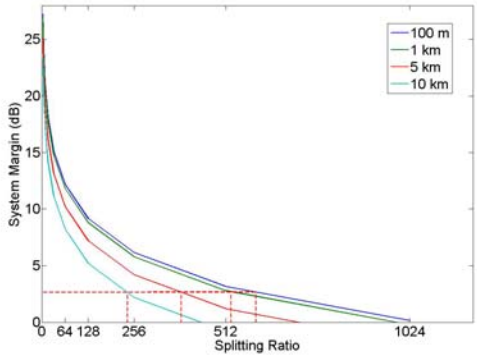

Figure 3 Power Margin vs. splitting ratio for various link lengths – Upstream

Another essential requirement is to transmit the 40MHz clock in the downstream. Therefore the data rate should be a multiple of 40MHz. We have used commercial EPON transceivers that are specified for a data rate of 1.25 Gbps. For this reason either 800Mbps either 1.6Gbps should be used. We have tested the transceivers at 1.6Gbps and their performance was adequate, so this rate has been selected. The data rate in the upstream is currently very low therefore the upstream data rate is not a limitation.

2.3. Development of a PON Demonstrator

A hardware PON demonstrator has been assembled in our laboratory to be able to measure network performance as a function of operational parameters. This demonstrator is based upon electrical transceivers embedded in a Xilinx Virtex5 FPGA mounted on a hardware development platform provided by Xilinx (ML523). These transceivers are connected to the commercial PON transceivers (1x OLT and 2x ONU) mentioned above. The optical network consists of a passive 1x8 optical splitter and a number of patch cables that allow us to vary the transmission distance. The FPGA firmware implements a basic PON protocol that would allow one OLT to communicate with up to 64 ONUs.

An additional system requirement is for the clock to be aligned to the time when collisions occur. Therefore the delay introduced by the PON (both its electrical and optical components) should be fixed and deterministic so that it can be compensated for. The work that has been carried out up to now to ensure that this requirement is going to be met is related to measurement of the timing properties of the transceivers embedded in the Xilinx Virtex5 FPGA.

The FPGA firmware that controls the clock and data recovery circuit of transceivers is still under development. The output clock frequency that is to be recovered from the higher bit-rate datastream is a factor of 20 times lower than the serial bit-rate. Thus we need to perform frequency division to recover the clock. Due to limitations of commercial FPGAs, frequency division introduces phase ambiguity. The rising edge of the lower frequency clock could be aligned to any of the rising edges of the high frequency clock (Figure 3). To resolve this problem a buffer has been used which starts filling with data at the edge of the low frequency clock. The position of the start of the frame with respect to the top of the stack can provide a measure of the phase offset of the recovered clock with respect the incoming datastream.

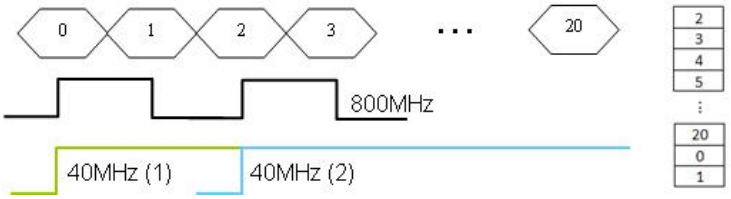

Figure 4 Phase ambiguity introduced by frequency division

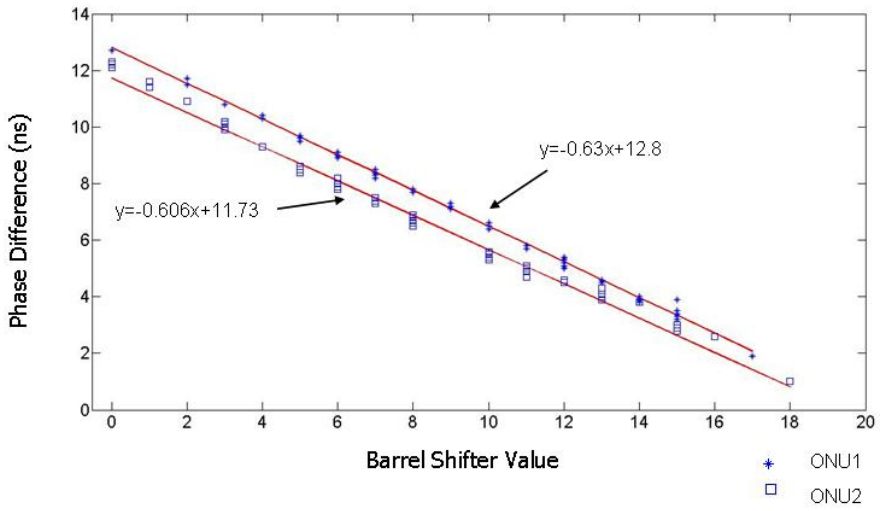

Figure 5 Barrel shifter value vs. phase difference

The first tests that have been conducted show that this technique can indeed be used to correct the ambiguity introduced by the frequency division. Figure 5 shows the relationship between the phase difference and the value of the barrel shifter (buffer) position of the MSB of the received data word for two different Optical Network Unit (ONU) receivers. The linear relationship observed indeed shows

that the required phase offset can be directly derived from the barrel shifter position, which in turn gives us confidence that automatic phase compensation can be implemented in firmware.

3. Future Direction – WDM PONs for DAQ Applications

3.1. Motivation

Although the current architectures of the data readout of the LHC experiments that comprised of multiple point-to-point links could meet the requirements for data readout at SLHC, point-to-multipoint architectures that would provide a more flexible architecture appear to be worth evaluating. The main problem that needs to be resolved is that data readout requires very high bandwidth in the upstream (approximately 5Gbps) and lower bandwidth (for control information only) in the downstream.

3.2. WDM PON architecture

The bandwidth requirement for the upstream is ~ 5Gbps. It becomes apparent that a TDM PON architecture is not an option for the upgraded DAQ network, as the bandwidth is shared in the upstream direction. However, a WDM PON architecture could be an option since each user can have a dedicated upstream wavelength and thus have the entire bandwidth is available per channel. Several WDM PON architectures can be found in the literature and they provide a wide choice of devices to be used for upstream transmission.

In the first instance the full requirements of the data readout system will be synthesized in order to be able to evaluate the relative merits of different components and architectures. Different available component types will be evaluated for their suitability for deployment inside a particle physics experiment, leading to a shortlist of candidates for detailed further study. Following this different network architectures will be similarly examined to arrive at a recommendation of the most suitable ones to be investigated further.

4. Conclusion

This paper investigated the possibility of using Passive Optical Networks for particle physics applications at CERN. Two types of applications have been considered; the distribution of the timing trigger and control signals (TTC) and the data acquisition. It seems that PONs can meet the application requirements in the first case (TTC); whether or not this is the most appropriate architecture is still an open question. In the second case the application requirements are quite challenging, therefore it is questionable whether PON architectures –even WDM PONs– will be able to meet the application requirements.

Acknowledgment. This work was supported by the ACEOLE FP7 EU framework.

References.

[1] Agrawal, Govind D., “Fiber-Optic Communication Systems”, New York: John Wiley & Sons, Inc., 2002