A Timing, Trigger, and Control System With Picosecond Precision Based on 10 Gbit/s Passive Optical Networks for High-Energy Physics

Eduardo Mendes, Sophie Baron, Csaba Soos, and Francois Vasey

***Abstract*— The Large Hadron Collider (LHC) experiments will have their timing, trigger, and control (TTC) system upgraded as a consequence of the need for higher bandwidth and components which are obsolete. In this article, we present a TTC based on passive optical networks (PONs). TTC-PON is a point-to-multipoint bidirectional TTC system based on the 10 Gbit/s International Telecommunications Union (ITU) XG-PON technology and modern field-programmable gate array (FPGA) devices. Each master can handle up to 64 slaves through a fully passive network, delivering a fixed-phase recovered clock to all the destinations with less than 5-ps jitter. TTC-PON pushes the limits of the PON technology by exploiting cutting-edge custom protocols on top of the commercially available XG-PON optical modules. It can potentially reuse the current optical fiber infrastructure already installed in the experiments and allows for high flexibility in terms of partitioning, which can ease future upgrades of the TTC network. The system features a picosecond-level on-the-fly phase monitoring for each slave's recovered clock by exploiting the bidirectionality of the network. In addition, a full set of link-quality monitoring tools was developed, allowing real-time performance monitoring. An overview of the tailored protocols will be given together with the details on the system implementation, operation, and performance. A discussion on the characterization campaign of more than 1000 optical modules delivered to the first implementation of the TTC-PON system will be drawn.**

***Index Terms*— Fast timing, field-programmable gate arrays (FPGAs), high-energy physics (HEP) instrumentation, optical links, timing circuits.**

I. Introduction

THE timing, trigger, and control (TTC) system is responsible for the transmission of synchronization, data-taking, and control signals within the Large Hadron Collider (LHC) experiments [1]. It interfaces a central trigger unit (CTU) to detector subpartitions located in the front-end (FE) of the experiments (where radiation might be present).

The original implementation [1], developed in the 1990s, is based on a unidirectional point-to-multipoint (P2M) optical network with a data rate of 80 Mbit/s. The transmission is synchronous to the LHC Bunch Clock, a 40.079-MHz timing signal derived from the radio frequency driving the particle beams. This allows the detector electronics to be synchronized to the particle bunches circulating in the accelerator rings [2]. An electrical network from the FEs to the CTU is employed to feedback, with the shortest possible delay (μs), a "busy" signal indicating that an FE buffer is in an overflow or warning state.

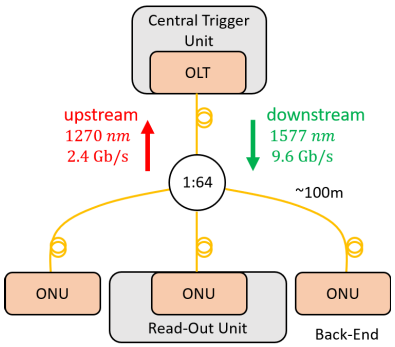

Fig. 1. Overview of future TTC upgrades for experiment back-end.

Even though the various LHC experiments have slightly different TTC implementations, the main requirements are similar. The TTC system must have a fixed and deterministic latency from the CTU to the FE after the system restarts [3]. The time interval error (TIE) jitter of the recovered bunch clock at the FE must be as low as 10 ps rms for the most timing stringent subdetectors in the experiments.

After the upgrades of the LHC experiments planned in 2022 and 2027 [4], the TTC system will also need an upgrade to cope with the expected higher particle collision rate and component obsolescence.

Inside the back-end zone (zone without radiation) of an experiment, the use of field-programmable gate array (FPGA) devices equipped with multigigabit transceivers (MGTs) is a natural choice for future TTC implementations [3]. In terms of optical network topology, a passive optical network (PON) commercial technology can be exploited [5]. Similar to the current TTC, the PON technology is P2M and uses single-mode fibers. This allows the potential reusability of existing TTC optical fiber infrastructure, while enabling a bidirectional communication scheme with higher data rates, requiring no electrical feedback used in the legacy TTC system. In addition, a PON-based TTC allows for high flexibility in terms of partitioning, which can ease future upgrades of the TTC network. An overview of a PON-based TTC system is shown in Fig. 1.

Manuscript received December 16, 2020; revised February 26, 2021; accepted March 2, 2021. Date of publication March 10, 2021; date of current version April 16, 2021. This work was supported by CERN.

The authors are with the Experimental Physics Department, CERN, 1211 Geneva, Switzerland (e-mail: eduardo.brandao.de.souza.mendes@cern.ch).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TNS.2021.3065128.

Digital Object Identifier 10.1109/TNS.2021.3065128

A proof-of-concept exploitation of commercial PON technologies and FPGAs for a TTC upgrade (TTC-PON) was first demonstrated in 2011 [6]. This article presents a TTC system based on modern FPGAs and XG-PON technology, where higher data rates are reached, and protocols were tailored to fulfill the requirements of the LHC high energy physics (HEP) experiments. In addition, a full set of online monitoring tools was developed for TTC-PON. They allow an evaluation of link quality and timing performance in real-time with no need for additional hardware. TTC-PON is currently being commissioned for the upgrade of the A Large Ion Collider Experiment (ALICE) and LHCb TTC system [7].

II. PON Technology

PON is a passive bidirectional P2M optical communication technology employed in commercial telecommunication access networks [5] in fiber to the home (FTTH) premises. In PON, two directions of propagation are distinguished: in the downstream, a master node called optical line terminal (OLT) broadcasts information to several slave nodes called the optical network unit (ONU) with a given wavelength; in the upstream, the ONUs send information back to the OLT using another wavelength. Different multiplexing schemes among the ONUs for the upstream data communication direction have been reported in the literature, such as code division multiplexing (CDM) [8], wavelength division multiplexing (WDM), and time division multiplexing (TDM) [5]. While WDM could be attractive for TTC applications, TDM-based PON networks are widely used in commercial applications and provide the basis of this work.

The proof of concept of a PON-based TTC [6] was built upon the 1G-EPON technology specified by IEEE, which is capable of handling 1 Gbit/s bidirectional communications. Since then, more modern PON technologies capable of handling 10 Gbit/s communications have become mature in the market. The 10G-EPON, from IEEE, supports a symmetric 10 Gbit/s communication scheme.

The International Telecommunications Union (ITU) also proposes a 10-Gbit/s PON technology recommendation [9]. This technology is called XG-PON and specifies asymmetric data rates of 10 and 2.5 Gbit/s in the downstream and upstream directions, respectively. Compared to the 10G-EPON technology, this technology has less overhead in the upstream using TDM, which significantly reduces the dead time and improves the upstream bandwidth efficiency. Less overhead for the upstream TDM is very attractive for our application as the "busy" waiting time is proportional to the time slot allocated to each ONU.

Future PON networks will take advantage of WDM and TDM to achieve an even higher bandwidth. The NG-PON2 technology, specified by ITU [10], offers a data rate of 40 Gbit/s in the downstream direction using four different wavelengths. These future technologies could be potentially exploited for a TTC-PON upgrade in a future study.

III. Data Layer Protocols

The physical layer of the TTC-PON system is based on the XG-PON technology. In order to have a transmission synchronous to the 40.079-MHz Bunch Clock frequency of the LHC, the line rates adopted in TTC-PON are multiples thereof and are 9.619 and 2.405 Gbit/s in the downstream and upstream directions, respectively. Those are the closest possible to the ITU recommendation of 9.953 and 2.489 Gbit/s. Henceforth, for simplification, 40.079 MHz will be referred as 40 MHz (and similarly for the multiple frequencies).

TABLE I
Downstream Frame Fields

Field	Size	Function
HDR	8 bit	Header field used by the ONU receiver for frame alignment and time-division multiplexing time reference
SC	4 bit	Slow-Control field used for internal system control and monitoring
USER	200 bit	User payload field for user data-transmission
FEC	28 bit	Forward Error Correction field to protect the frame against line-errors

An important aspect, inherited from the standard, is that the ONU uses its downstream recovered clock to synchronize the upstream transmission. The system can comfortably work with 64 slaves.

The communication protocols for the downstream and upstream directions are tailored to meet the specifications of the HEP experiments, namely low and fixed latency in the downstream direction (in the order of 100 ns), low jitter (in the order of picosecond) for the downstream recovered Bunch Clock, and small waiting time per ONU in the upstream data-transmission scheme (in the order of micrometers).

A. Downstream Framing

The downstream frame is constructed on a Bunch Clock period basis of 25 ns. A total of 240 bit are divided into four main fields, detailed in Table I.

The slow-control (SC) and user payload (USER) fields are scrambled using a self-synchronous scrambling technique [11] to ensure the dc-balance. The header and scrambled data are then systematically encoded using two Bose Chaudhuri Hocquenghem (BCH) codes [12], denoted BCH(120,106) to indicate that a 106 bit word forms a 120 bit frame. Each of these codes is capable of correcting up to two random errors per word. The motivation for this type of code was that mainly random errors were observed in the system at high attenuations.

The higher bandwidth in the downstream direction offered by TTC-PON, compared to the current TTC system, allows for a more complex trigger scheme implementation and more flexibility in the control of different subpartitions of an experiment.

B. Upstream Framing

The upstream channel is shared among ONUs based on a TDM scheme consisting of 125-ns time slots per ONU which are arbitrated in a round-robin fashion by the OLT. Each ONU transmits during a 100-ns burst, resulting into an extremely short gap of only 25 ns between ONUs. The upstream framing is depicted in Fig. 2.

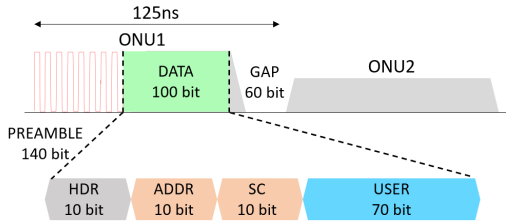

Fig. 2. Upstream framing.

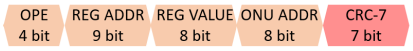

Fig. 3. Downstream SC framing.

The protocol was optimized to minimize the upstream waiting time. The burst size is reduced to 125 ns with respect to typical ITU XG-PON specifications (in the order of micrometers) while still respecting the minimum physical recommendations for the XG-PON standard in terms of gap and preamble. These are the lower limit recommendations of the XG-PON specifications which may not be implemented for all transceivers. The system was fully characterized with optical devices from Go!Foton and HiSense.

An 8b10b encoding scheme [13] is used to provide the dc-balance and error detection capability. Therefore, the effective user payload is 56 bit per burst.

It is worth noting that the nature of the upstream protocol is such that the bandwidth and the waiting time are dependent on the number of ONUs present in the system. For 32 ONUs, the effective user bandwidth is 14 Mbit/s, and the waiting time is 4 μs. Furthermore, for 64 ONUs, these values become 7 Mbit/s and 8 μs, respectively. These "latency" values are acceptable for the experiments that adopted TTC-PON.

In the current TTC system, the separate electrical network transmitting a single bit which indicates the "busy" state has a propagation delay of around 3 μs. In contrast, TTC-PON offers a fully integrated TTC network with the flexibility of optimizing the bandwidth and waiting time requirements for the different subdetectors.

C. SC Protocol

Each ONU has a byte-addressed 512 $\times$ 8b memory. The OLT can read and write from/to the ONU memory through the SC protocol. This memory contains important parameters for the upstream TDM and a full set of monitoring functions.

In the downstream direction, SC fields from nine consecutive frames are assembled to form 36 bit commands to a given ONU, as shown in Fig. 3. Through these commands, the OLT can address each ONU in the network individually, or use the broadcast reserved address (0xFF), to address all ONUs simultaneously.

The SC frame is made of 29 bits of control and 7 bits for a cyclic redundancy check (CRC) [11], namely CRC-7 error detecting code. This CRC code is capable of detecting three random errors and up to seven consecutive errors. The ONU recovers the SC frame position, thanks to a CRC-based headerless frame alignment technique [14].

Fig. 4. Upstream SC framing.

TABLE II
DESCRIPTION OF SC OPERATIONS

Operation (OPE)	Value	Description
IDLE	0000	Idle operation
WR	1111	OLT writes to ONU memory without asking for acknowledgment
RD	0001	OLT reads from ONU memory ONU answers a read command
WR_ACK	1110	OLT writes to ONU memory and asks for acknowledgment ONU answers with an acknowledgment command
RCVD_ERR	1000	ONU answers in case of detected error in the downstream CRC check

If requested, an ONU answers back to the OLT using its upstream SC field. The SC fields of three 8 bit bursts from the same ONU are assembled in order to form an SC answer as shown in Fig. 4.

The possible SC operations (OPE field) are displayed in Table II. The set of possible unused operation values are reserved for future implementations.

IV. UPSTREAM TDM ARBITRATION

In TTC-PON, the upstream data transmission occurs in a round-robin manner. This requires a careful calibration of the system in order to avoid having two ONU bursts overlapping with each other. The specificities of the TTC-PON upstream TDM arbitration are detailed in this section.

A. Round-Robin Arbitration

In order to implement a round-robin arbitration, each ONU in the system disposes of an internal timer and a time slot given by the OLT to start transmission. The OLT continuously sends a special type of header (referred to as heartbeat) every 125 ns multiplied by the number of ONUs in the system, which is used by the ONU to reset its internal timer.

B. Calibration

The different optical branches connecting an ONU to a TTC-PON system can have different lengths (the different branches can have a length difference in the order of kilometers). Therefore, at system startup, before all ONUs are allowed to transmit, a fully automatized calibration procedure requiring no external hardware is performed to compensate for latency asymmetries in the system. The calibration procedure is also used to find the position of PON devices' timing critical signals [15]. The functions are further detailed here in the order in which they happen during calibration.

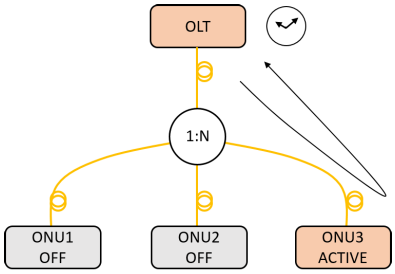

Fig. 5. Roundtrip time measurement principle.

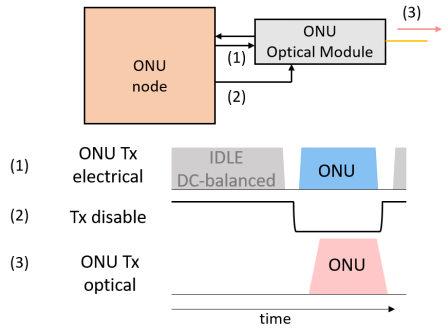

Fig. 6. ONU transmitter disable signal illustration.

1) Compensation for Latency Asymmetries: In order to compensate for different latencies of ONUs, the roundtrip time of each ONU is measured. During this measurement, a single ONU is asked by the OLT to transmit in continuous mode, while the others are off. Every time the ONU receives a heartbeat header, it will send it back to the OLT. By measuring the time elapsed between the heartbeat transmission and reception, the OLT can measure the roundtrip time for each ONU as shown in Fig. 5.

The roundtrip time is then taken into account in order to calculate the time slot allocated to the ONU for its upstream transmission. At the end of the calibration, the OLT sends to each ONU their individual timing information. The roundtrip time measurement has a bin size of 104 ps, equivalent to one downstream unit interval (UI).

2) ONU Transmitter Disable Positioning: In normal system operation, the ONU is only allowed to transmit during its own time slot. Therefore, while one ONU is transmitting, the laser of all the other ONU modules shall be turned off. This is achieved, thanks to the *transmitter disable* signal as shown in Fig. 6. This signal has to be applied to the ONU optical module and its timing is critical to ensure the ONU transmitter is deactivated in the right position without interfering with others.

During this second stage of the calibration, the ideal timing to deactivate the *transmitter disable* signal is found. The technique we propose here is an online adjustment orchestrated by the OLT module during calibration. A single ONU transmits a pseudorandom bit-sequence (PRBS) data in the user field during the adjustment (refer to Fig. 2). The delay of its *transmitter disable* signal is increased until a good position is found. The means of finding the good position is a bit-error-ratio (BER) test performed by the OLT receiver. The

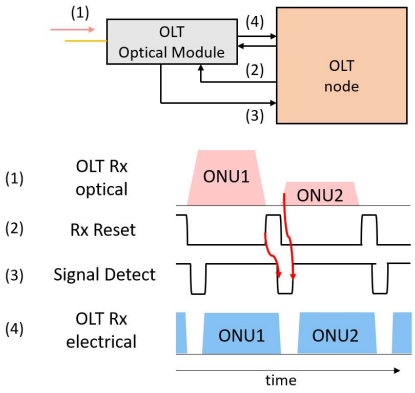

Fig. 7. OLT receiver reset signal illustration.

BER threshold for the test can be adjusted (the default is around 10^{-7}).

3) OLT Burst-Mode Receiver Reset Positioning: One final step is then required for proper operation of the system. The OLT burst-mode optical receiver has to deal with different receiver power levels from the many ONUs connected in the network. In order to reach the most demanding dynamic range specifications of the XG-PON protocol, a *reset* signal between ONU bursts is specified in the standard. This signal has to be applied to the OLT optical device precisely during the gaps as shown in Fig. 7. To ensure the position of the *reset* signal is exactly between the ONU optical bursts, during the third step of the calibration, we present an online position finding technique, which requires no external measurement circuitry.

The technique makes use of the *signal detect* output from the OLT optical module defined in the standard. This signal falls to zero at the falling edge of a *reset* pulse and moves to one once enough optical power is detected by the module. By sweeping the *reset* position and measuring the timing between the falling and rising edge of the *signal-detect* signal, we can establish the position of the *reset* signal with respect to the incoming bursts. The *reset* pulse has a duration of 25 ns and is repeated every 125 ns, thanks to the round-robin transmission nature of the upstream TDM scheme.

V. Node Implementation

The back-end boards that will be installed during the LHC experiments upgrades are equipped with modern FPGA devices [7]. The implementation of the TTC-PON OLT and ONU nodes is, therefore, provided as device-agnostic soft-cores which can be integrated in the FPGA firmware running on these boards. The cores are fully controlled via software, providing a high-level of flexibility and a low resource usage in the FPGA.

The designs are publicly available on GitLab [16]. The soft-cores are provided with example designs targeting commercial FPGA evaluation boards (KCU105 from Xilinx and A10GX from Intel). A control software runs in Python in a computer. They provide a departure point for a designer to evaluate the TTC-PON system and are generic to all LHC experiments.

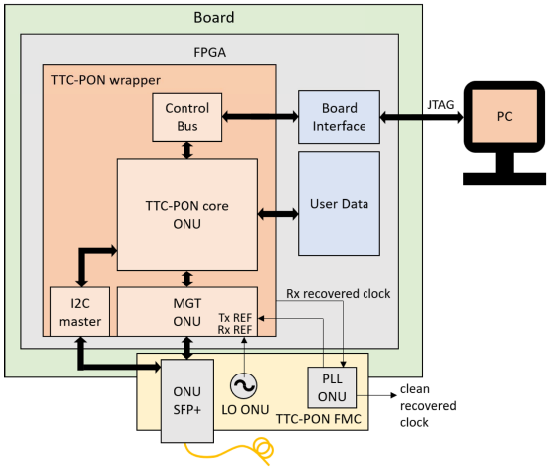

Fig. 8. Simplified block diagram of a TTC-PON ONU node.

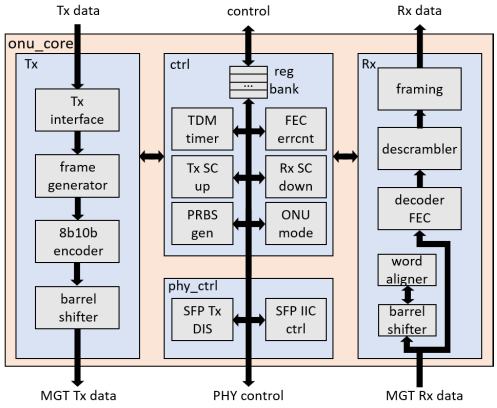

Fig. 9. Simplified block diagram of the ONU core.

The example designs make use of a custom-designed FPGA Mezannine Carrier (FMC) board, the TTC-PON FMC. The TTC-PON FMC contains all the required external components for a TTC-PON node implementation: XG-PON SFP+ optical module, phase-locked-loop (PLL), and local oscillator (LO).

The design structure was kept as similar as possible for OLT and ONU nodes. An illustrative block diagram is shown in Fig. 8.

In the FPGA core, the TTC-PON wrapper encapsulates FPGA-agnostic functions (implemented in the TTC-PON core) and device-specific interfaces (MGT and control bus). For an implementation in the Kintex Ultrascale family, the TTC-PON wrapper occupancy is around 2300 configurable logic block lookup table (CLB LUT) and 2800 CLB REG (Register) for an ONU. For an OLT, it occupies 2900 CLB LUT and 4600 CLB REG.

In the example design, a Joint Test Action Group (JTAG) interface is used to communicate between the TTC-PON wrapper and the computer. The control bus used to communicate with the TTC-PON core is an AXI-4Lite for Xilinx devices and Avalon for Intel devices. In addition, an example logic of how to handle the data interfaces is also provided. This example logic allows the interested user to quickly perform BER tests and latency measurements.

In Fig. 8, the TTC-PON FMC connections are shown for an ONU node in yellow. The ONU MGT receiver is fed with a free-running clock coming from an LO as a reference for the clock and data recovery (CDR) circuitry. The recovered clock is then used to feed an external PLL to generate a clean MGT transmitter reference clock. In the OLT node, a single reference clock is needed for the MGT. In this case, the external PLL is used to clean an external Bunch Clock input which then serves as reference to the MGT.

VI. FPGA Cores

The TTC-PON FPGA cores share a similar top-level architecture as shown in Figs. 9 and 10. The architecture is composed by four main blocks: transmitter (*Tx*), receiver (*Rx*), control (*ctrl*), and physical control (*phy_ctrl*). In the *Tx* and *Rx* blocks, the high-speed encoding/decoding protocols are implemented. The *ctrl* block is used for low-level functions related to the SC protocol, TDM arbitration, and error monitoring. The *phy_ctrl* block is used to implement SFP- and MGT-related functions. The actual implementation of those blocks is different for OLT and ONU devices and is detailed below. The main challenges in the system implementation are fully discussed in [17].

A. ONU Core

A block diagram of the ONU core is shown in Fig. 9. The *Tx* and *Rx* datapaths are clocked by a unique 240 MHz clock recovered from the downstream serial stream by the MGT. The clock-domain crossing (CDC) between the transmitter and the receiver recovered clock are dealt in the MGT level using a mesochronous fixed-phase first-in-first-out (FIFO) memory. For Xilinx UltraScale devices, the high precision timing distribution (HPTD) IP core [3] is used for the CDC to ensure a higher level of phase stability.

In the *Tx interface*, the data captured from the user is fed to the *frame generator* where it is assembled with the preamble, header, SC, and address fields. The header, address, SC, and user data are 8b10b encoded (*8b10b encoder*). A *barrel shifter* is employed to provide a UI-level (416 ps) phase-shifting capability to the transmitter. This is used in the fine alignment of bursts for the upstream TDM roundtrip time calibration.

In the *Rx* datapath, a *barrel shifter* and *word aligner* are used for the header alignment to the correct word boundary. The technique used for the frame alignment ensures a fixed phase between the header in the serial stream and the recovered clock edge after resets. This technique is detailed in [3] and [18]. Once the frame is aligned, the data is forward error correction (FEC) decoded (*FEC decoder*), descrambled (*descrambler*), and assembled in frames sent to the user interface (*framing*).

The *ctrl* block has a register bank which stocks all the control registers necessary for the core operation and monitoring. It also has the logic necessary for the transmitter SC upstream encoding (*Tx SC up*) and the receiver downstream SC decoding (*Rx SC down*). In addition, a *PRBS generator* for upstream fast BER tests and error logging (*FEC errcnt*) are also available. The *ctrl* block also contains the *TDM timer* used for the TDM

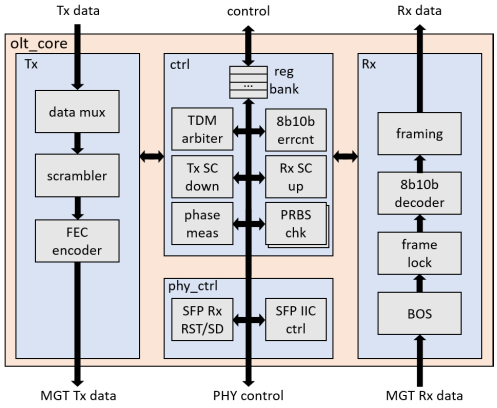

Fig. 10. Simplified block diagram of the OLT core.

arbitration. The ONU mode logic allows the ONU to run in different modes (calibration or normal operation).

B. OLT Core

The OLT core is shown in Fig. 10. In the *Tx* block, the user data is complemented by the SC and the header, multiplexed into words (*data mux*) to be then scrambled (*scrambler*) and FEC encoded (*FEC encoder*). It is important to note that the *scrambler* and *FEC encoder* run word by word instead of a direct encoding of the whole frame with a fully combinatorial logic. This choice is natural to save resources as the scrambler and BCH encoder are made of an unfolded architecture of a linear-feedback shift register (LFSR).

In the *Rx* path of the OLT, a known difficulty in burst-mode data reception is that the bursts coming from different ONUs can be out of phase, unlocking the traditional continuous MGT CDR circuitry. In TTC-PON, the new burst is so short that relocking the CDR for every new burst is impossible. Therefore, we avoided traditional CDR circuitry in the design and an alternative solution had to be developed [19].

The solution adopted in our system relies on an oversampling technique with a feed-forward phase selection. The CDR circuitry of the OLT MGT is put in a hold mode which locks it to the reference clock. The data arriving from the ONUs runs at 2.4 Gbit/s and is captured at 9.6 Gbit/s, and therefore, yielding four received samples for a single-bit transmitted by the ONU.

The raw data coming from the MGT goes to the blind-oversampler (BOS) block further detailed in Fig. 11. The BOS continuously downsamples the four captured phases of the ONU burst (*downsample*) and each phase is fed to a *word aligner* looking for the upstream header pattern. Once headers are found, the *frame data multiplexer* votes for the best phase among all phases which have identified a header. The selection criterion is based on which phase is the closest to the middle of the eye (i.e., no edges are observed). The resolution of the phase detection is 104 ps. This phase information is used for the precise timing measurement in the upstream calibration.

The selected phase feeds the *frame lock* circuitry. This circuit makes use of the cycle repeatability of the TTC-PON upstream protocol by ensuring that only headers detected in

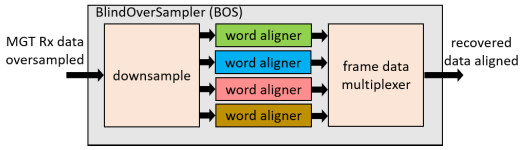

Fig. 11. Simplified block diagram of OLT BOS.

a certain time window are valid. This increases the system reliability by rejecting potential fake headers arriving during the preamble or gap periods where the electrical data coming from the optical module can be unpredictable. The data is then 8b10b decoded (*8b10b decoder*) and feeds the *framing* interface.

The *ctrl* block has the logic necessary for the transmitter SC downstream encoding (*Tx SC down*) and receiver SC upstream decoding (*Rx SC up*). In addition, two *PRBS checkers* for upstream fast BER tests and error logging (*8b10b errcnt*) are also available. The *ctrl* block also contains the *TDM arbiter* used for the TDM arbitration and the required logic for phase measurement during the roundtrip calibration.

VII. Monitoring Software

Several online monitoring features via software are available in the system to help users troubleshoot problems and gain confidence in its performance during operation. Those are divided into two categories, further explained below.

A. Link Quality

The link quality can be continuously monitored both in the downstream and upstream directions. An estimation for the BER can be calculated, thanks to the error correcting and detecting codes available in both directions.

XG-PON OLT and ONU optical transceivers are also compliant with the SFF-8472 standard for digital diagnostic monitoring. The TTC-PON software fully exploits this standard and provides monitoring of all parameters from the optical transceivers such as voltage, temperature, receiver power level, transmitter power level, and transmitter bias current. The continuous monitoring of those parameters over time, proposed by the software, can be useful to diagnose a potential degradation of the modules.

In addition to the previous link quality monitoring functions, modern FPGA transceivers are equipped with on-chip noninvasive eye diagram functions. This feature is also implemented for the ONU receivers.

B. Timing

Slow phase variations in the system can be monitored, thanks to the bidirectionality and synchronization of the network. Several other timing distribution systems employ a similar timing monitoring approach [20], [21]. The downstream phase variations are estimated from a continuous roundtrip (downstream and upstream) phase variation measurement.

The roundtrip phase is the position in which the ONU burst is detected. The bin size for this phase measurement is 104 ps.

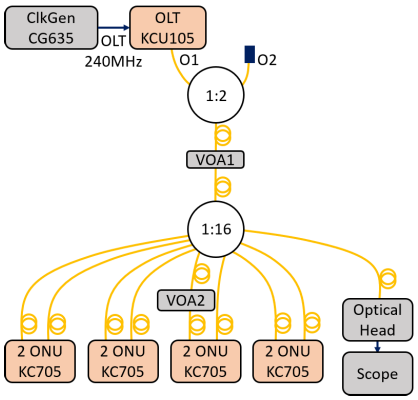

Fig. 12. Setup for quality assurance (shown for OLT).

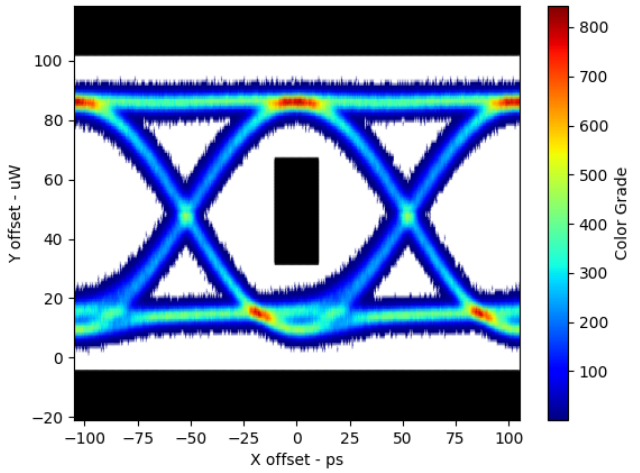

Fig. 13. Example of an eye diagram of an OLT transmitter at 9.6 Gbit/s. The transmitter power is estimated based on the eye OMA and the total loss measured in the optical path. The black area denotes the mask limits of the standard.

An assumed symmetric phase variation in the downstream and upstream paths yields a downstream phase variation which is half the one measured in the roundtrip.

The resolution of the measurement can be further improved with averaging. In particular, in the Xilinx UltraScale architecture, it is possible to shift the OLT high-speed receiver clock when the CDR is in hold mode. The phase of the OLT receiver clock is swept back and forth over 832 ps (two upstream UIs) in steps of 1.63 ps yielding 1024 measurements which are then averaged.

VIII. Characterization Results

The typical figures of merit used to characterized the TTC-PON system are discussed in [22]. In 2018, more than 1000 ONU and 100 OLT optical modules were received, which were delivered to the ALICE and LHCb experiments. A system test was performed for all the modules to verify their performance is maintained within the specified line rates and tailored protocols, to ensure confidence in the scalability of the system.

A. Link Performance

The quality assurance setup of OLT modules is shown in Fig. 12. The OLT is implemented on a KCU105 evaluation board and equipped with a TTC-PON FMC. The OLT receives an external clock from a commercial clock generator CG635. Eight ONUs are implemented in this system using four KC705 evaluation boards. Each KC705 is used to implement two ONUs: one ONU is implemented using a TTC-PON FMC, and another using the on-board small form-factor pluggable (SFP) connector and an external Si5344 PLL evaluation board. Special care is taken in the FPGA design so that the logic and clocking among ONUs are fully independent.

A variable optical attenuator (VOA) JDSU-mVOA, here denoted VOA1, is used to attenuate the common optical branch in the system for BER measurements. A second optical attenuator, JDSU-mVOA (VOA2), is used to perform dynamic range measurements (used to test the OLT receiver capacity to adapt to bursts of varying power levels). Note that VOA1 and VOA2 attenuate light in both upstream and downstream directions. The ONU modules are connected through a wideband splitter to the OLT. The splitters used in these setups are from Huhber + Suhner. An optical head 81495A is used to measure the power received by the ONUs and to convert the downstream optical signal to an electrical signal, which is fed into a high-speed oscilloscope DSA91204A, to perform eye-diagram measurements of the OLT Tx. The power of each branch of the system is calibrated beforehand to take into account the network uniformity of the optical tree. The total loss of the optical network from the OLT to the power-meter is also calibrated to allow an indirect measurement of the OLT transmitter power.

In the quality assurance for the ONU modules, the setup is similar to the one used for the OLT modules, except for the positioning of the optical head at the output O2 of the 1 : 2 optical splitter, in order to perform eye diagrams of the ONUs transmitted optical signals.

We summarize below the module performance in the upstream and downstream directions, including a discussion on the system power budget.

1) Downstream: A typical eye diagram for an OLT module is shown in Fig. 13. As the line rate is operated differently from the specified standard, eye diagram of the OLT is performed at three different data rates, 9.3, 9.6, 10.2 Gbit/s, respectively, to ensure that the module works comfortably within the data-rate margin.

In the downstream direction, the optical power budget depends on the performance of both the OLT transmitter and ONU receiver. The OLT Tx optical modulation amplitude (OMA) minimum specification is 5.68 dBm and the maximum sensitivity for the ONU receiver is −22 dBm (for a BER of 10^{-11}).

The sensitivity of ONU modules is measured by performing a BER versus OMA test. The line BER, with no error correction, here referred to as pre-FEC, is estimated from the FEC error counters as shown in Fig. 14 for any given ONU module. The system BER (here referred as post-FEC) is calculated using a PRBS error checker over the downstream

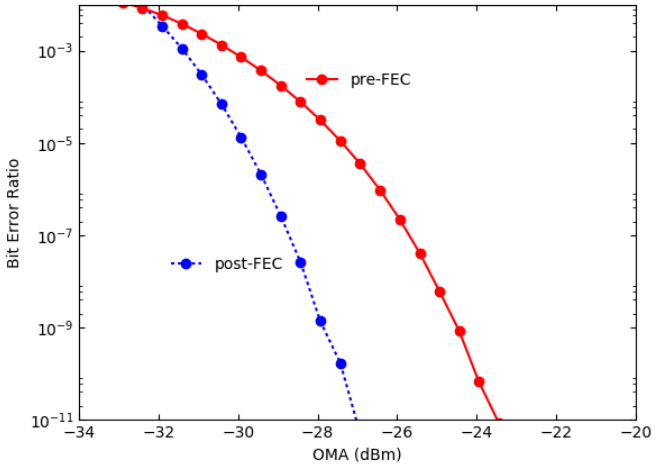

Fig. 14. Example of a BER in the downstream direction (ONU Rx sensitivity).

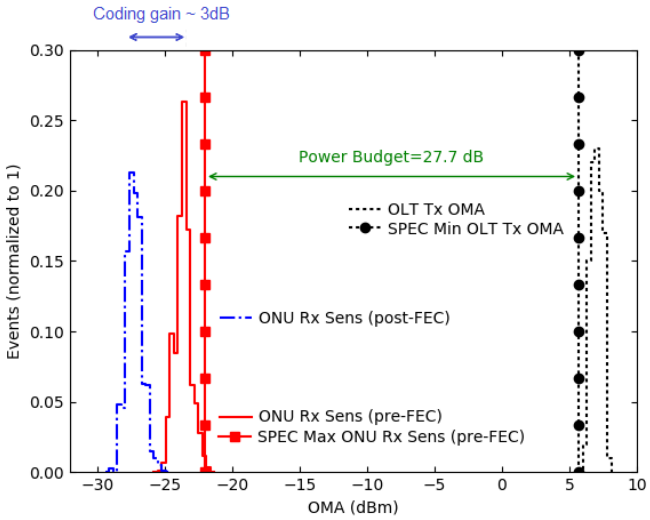

Fig. 15. Histogram power budget downstream. All sensitivity values are displayed for BER $= 10^{-11}$. The power budget shown does not include the FEC margin.

user payload. A coding gain of around 3 dB at BER of 10^{-11} is obtained in the downstream direction as a result of the FEC scheme.

A histogram summarizing the downstream performance measured for all modules in the quality assurance process is shown in Fig. 15, where it is shown that all OLT and ONU modules meet the specifications.

2) Upstream: A measured eye diagram in the burst mode for an ONU module is shown in Fig. 16. The burst eye diagram is captured from the last bits of the ONU preamble up until the final bit of the payload.

In the upstream direction, the optical power budget depends on the performance of the ONU transmitter and OLT receiver. The ONU Tx OMA minimum specification is 3.68 dBm and the maximum receiver sensitivity for the OLT receiver is -25 dBm (for a BER of 10^{-9}). The different BER values specified for the downstream and upstream sensitivity are constrained by the total duration time of the test, as the

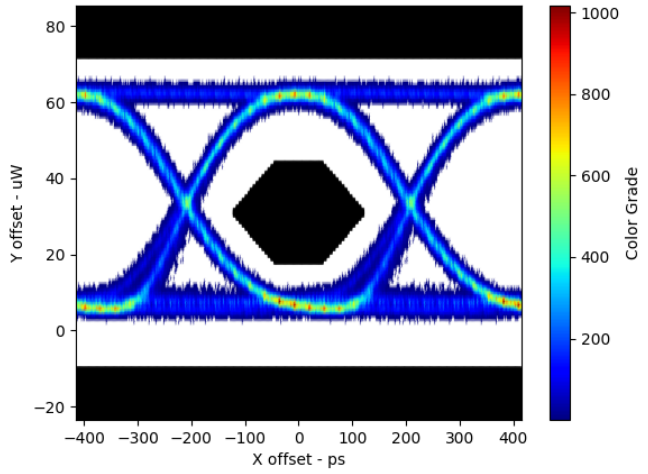

Fig. 16. Example of an eye diagram of an ONU transmitter.

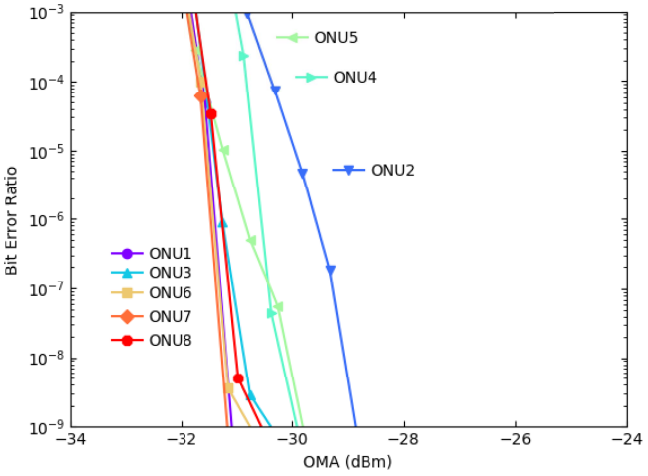

Fig. 17. Example of a BER in the upstream direction (OLT Rx sensitivity).

effective data rate of the upstream for a given ONU is much lower than in the downstream.

The sensitivity of the OLT modules is also measured by performing a BER versus OMA measurement as shown in Fig. 17. The BER is checked simultaneously for all eight ONUs present in the system and the OLT sensitivity is selected as the worst observed case.

The measurement is repeated for fixed attenuations of 0 and 9 dB and applied to the VOA2 to ensure the OLT can cope with 9 dB power difference between ONU modules. No significant difference was observed.

A histogram summarizing the upstream performance measurements for all modules during the quality assurance process is shown in Fig. 18. It is observed that the ONU modules meet the specification for minimum transmitter power. Similarly, the OLT modules have met the specification for maximum receiver sensitivity with a margin of around 2 dB.

3) Discussion on Power Budget: In Table III, we can observe the typical attenuation of a PON network for a 1:64 and 1:128 split ratio. We assume in both scenarios that 1:2 and 1:32 (1:64) splitter are present in the system.

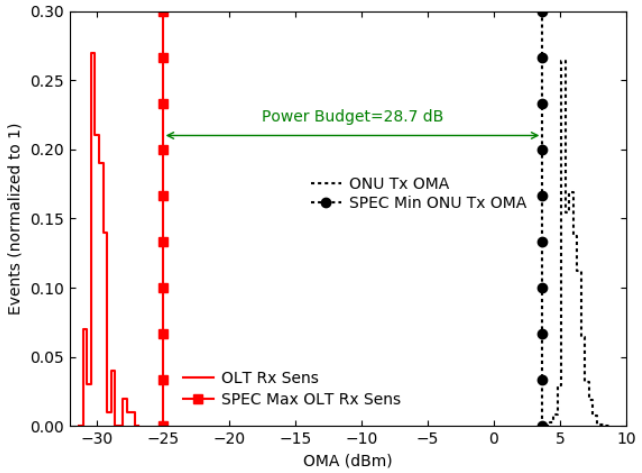

Fig. 18. Histogram power budget upstream. All sensitivity values are displayed for BER $= 10^{-9}$.

TABLE III
PON ATTENUATION OF PASSIVE COMPONENTS

Device	Attenuation (dB)
Splitter 1:32 (1:64)	17.0 (20.5)
Splitter 1:2	3.7
Fiber - 200 m, 0.5 dB/km	0.1
Connectors - 4 units, 0.5 dB/unit	0.1
Total	**22.8 (26.3)**

The total expected attenuation is 22.8 and 26.3 dB for the two considered scenarios.

The power budget detailed in the quality assurance specifications for the downstream and upstream directions are 27.7 and 28.7 dB, respectively, with a BER $= 10^{-11}$ and BER $= 10^{-9}$. Considering a BER $= 10^{-12}$ target in both directions, we estimate an additional penalty of 1 and 2 dB for the downstream and upstream directions with respect to the sensitivity measurement in the quality assurance. Finally, since the downstream sensitivity is specified without FEC (i.e., pre-FEC), we consider a gain of 3 dB to the power budget from the FEC coding gain as was observed in Fig. 14.

The system power budgets are thus 29.7 and 26.7 dB in the downstream and upstream directions, respectively. The optical margins for the 1:64 scenario are therefore, 6.8 and 3.9 dB; both margins are very comfortable for a 1:64 split ratio.

For the 1:128 scenario, we obtain a margin of 3.4 dB for the downstream direction, which is considered an acceptable margin. For the upstream direction, a small but positive margin of 0.4 dB is obtained. As observed in Fig. 18, the receiver performance of all OLT modules measured was always at least 2 dB above the maximum sensitivity specification. This allows us to conclude that in both directions, the margin for the 1:128 scenarios would be above 2.4 dB for the modules analyzed.

B. Timing Performance

A careful evaluation of TTC-PON timing performance was carried out. The setup used for the timing measurements presented in this article is shown in Fig. 19. The OLT receives an

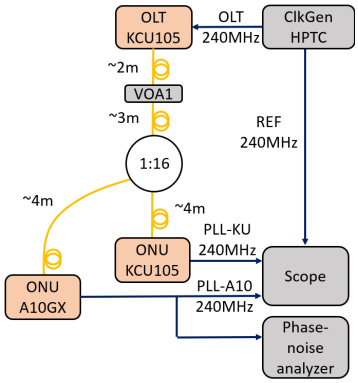

Fig. 19. Setup for TTC-PON timing measurements.

240-MHz signal from the high-precision timing clock (HPTC) generator. The HPTC is a low-cost custom-designed clock generator for timing measurements developed at CERN. The design is available on GitLab [23]. In this setup, two ONUs were implemented using the KCU105 and A10GX evaluation boards as these contain the FPGAs in which TTC-PON will be implemented.

The ONU implemented in the KCU105 is equipped with the TTC-PON FMC. The ONU in the A10GX uses the on-board SFP cage and an external PLL evaluation board. In both cases, the external PLL for the ONU implementation is an Si5344 integrated circuit with a bandwidth of around 2.7 kHz and configured in zero-delay mode. The clock being analyzed in both cases is a 240-MHz recovered cleaned clock delivered by the ONU external PLL as shown in Fig. 8. For the ONU-A10GX, the clock comes directly from the output of a PLL while for the ONU-KCU105 there are two additional clock multiplexing components related to the TTC-PON FMC design.

The timing performance of this network is evaluated and the results are further detailed in Sections VIII-B1–VIII-B3.

1) Phase Measurement: Initially, the OLT fine roundtrip phase variation measurement accuracy was evaluated. To perform the measurement, the phase of the ONU-KCU105 transmitter data is shifted over a span of 800 ps in 1.63 ps steps, thanks to the Xilinx Tx phase interpolator (PI). The phase measured by the OLT is acquired 100 times for each ONU phase.

We observe in Fig. 20 the integral nonlinearity (INL) for the phase variation measured in the OLT. The INL includes both the OLT phase-measurement circuitry and the Tx PI nonlinearitites. The peak-to-peak of the mean INL is around 10 ps over the measured range.

2) Phase Determinism After Resets: If the system is restarted, it requires a deterministic phase to avoid timing calibrations, which have to be performed very often. To evaluate the phase determinism after restarts, a phase measurement is performed with an oscilloscope DSA91204A. An edge–edge measurement is performed between the ONUs recovered cleaned clocks and a reference clock sourced from the HPTC. The sampling rate is set to 40 GS/s with an acquisition window

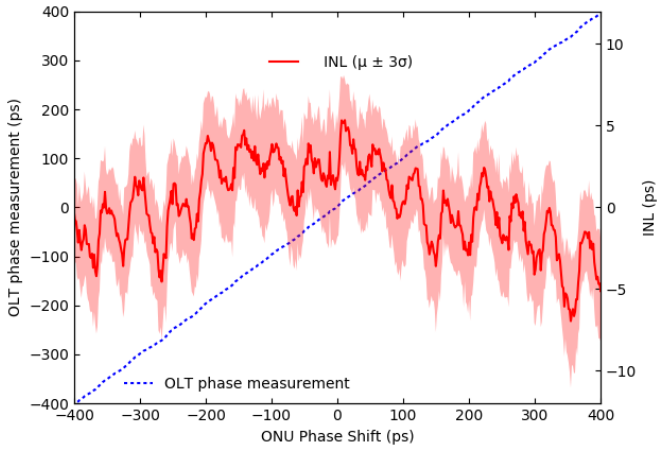

Fig. 20. OLT roundtrip phase variation measurement and INL for different ONU Tx phases.

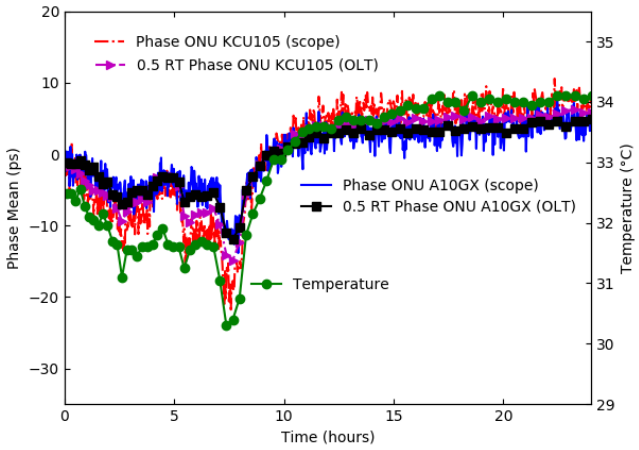

Fig. 21. Phase determinism over time. A full system reset is applied before each sample acquisition.

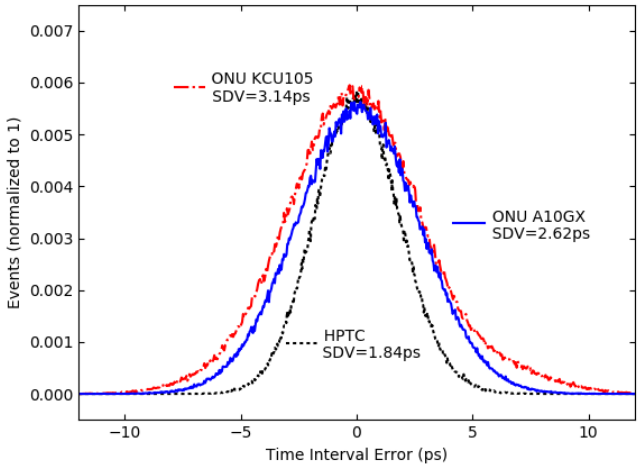

Fig. 22. TIE histogram. A total of 600k phase measurements is present per curve.

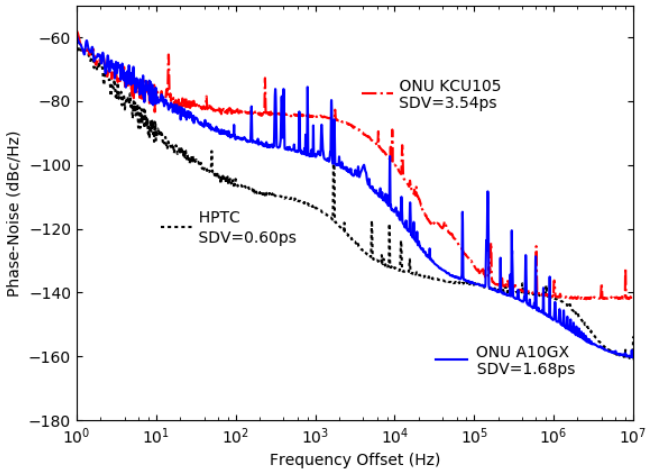

Fig. 23. Phase noise. The SDV values are integrated from 1 Hz to 10 MHz. All clocks have a frequency of 240 MHz.

of approximately 2.5 ms. For each phase measurement, around 600 kS is accumulated.

The setup is fully reset sequentially before each measurement in the following order: reset the OLT core and the MGT, reset the ONU cores, and the MGTs and reset the ONU external PLLs. During the same test, the roundtrip phase variation monitoring is also recorded. The fine phase monitoring feature is used to estimate the downstream clock phase variation by assuming a model where the downstream phase variation is half of the roundtrip phase variation.

Fig. 21 shows the mean phase measured for 800 reset sequences over 24 h. No phase jumps larger than 10 ps after a reset, and the mean phase variation has a direct correlation with the measured ambient temperature.

The average roundtrip phase variation monitoring yields a very good estimation of the downstream phase variation. If required, this feature can be exploited for the implementation of a timing compensation scheme on the TTC-PON system.

3) Jitter: To evaluate the clock quality in the time domain, a TIE jitter measurement is performed with an oscilloscope. The oscilloscope is set-up in a similar configuration described in the phase-determinism measurement. Fig. 22 shows the histogram of the Clock TIE jitter measurements (for a single-shot acquisition of the scope) using a constant-frequency clock-recovery scheme. The TIE standard deviation (SDV) is around 2.6 ps for the ONU-A10GX and 3.1 ps for the ONU-KCU105.

The jitter in the frequency domain was evaluated using a phase-noise measurement, performed with an FSWP8 phase-noise analyzer. Fig. 23 shows the quality of the HPTC clock generator, with an integrated phase noise from 1 Hz to 10 MHz of approximately 0.6 ps. The integrated phase noise, for a similar integration range, is 1.7 ps for the ONU-A10GX and 3.5 ps for the ONU-KCU105.

IX. Conclusion

A TTC system based on the 10 Gbit/s XG-PON technology from ITU is presented. TTC-PON is a P2M bidirectional system which enables the reuse of the current optical fiber infrastructure already installed in the LHC experiments for legacy TTC implementations. The system offers the potential

for future upgrades accommodating up to 128 slave nodes in a purely PON.

The tailored communication protocols adopted in the system were discussed. They target the optimization of the most important figures of merit for HEP timing applications, namely line synchronization with the collision rate, fixed and deterministic latency, and low upstream waiting time. The TTC-PON system was optimized for the LHC experiments; however, other PON-based networks may benefit from the underlying concept of optimizing the dead time allocated to ONU nodes.

The hardware, firmware, and software for a TTC-PON network are presented and the sources are available online. The monitoring features allow continuous evaluation of the link and timing performance in real time. The technique described for picosecond-level phase measurement of the ONU recovered clock can be utilized for long-distances timing distribution systems based on PON technology.

The results of the quality assurance tests for more than 1000 ONU modules and 100 OLT modules, which will be installed in the ALICE and LHCb experiments, are presented. The characterization of the system timing shows the high level of performance achieved, making the TTC-PON a back-end solution capable of coping with the most stringent timing requirements of the LHC experiments.

ACKNOWLEDGMENT

The authors would like to thank our colleague Anthony Frederick Bulling from CERN for reviewing the manuscript and Paolo Novellini from Xilinx Italy for all the support provided during the implementation of the system and for the useful discussions on the implementation of burst-mode clock and data recovery in field-programmable gate arrays (FPGAs).

REFERENCES

[1] B. G. Taylor, "TTC distribution for LHC detectors," in *Proc. IEEE Nucl. Sci. Symp. Conf. Rec.*, vol. 1, Nov. 1997, pp. 87–91.

[2] S. Baron, T. Mastoridis, J. Troska, and P. Baudrenghien, "Jitter impact on clock distribution in LHC experiments," *J. Instrum.*, vol. 7, no. 12, Dec. 2012, Art. no. C12023.

[3] E. Mendes, S. Baron, C. Soos, J. Troska, and P. Novellini, "Achieving picosecond-level phase stability in timing distribution systems with Xilinx ultrascale transceivers," *IEEE Trans. Nucl. Sci.*, vol. 67, no. 3, pp. 473–481, Mar. 2020.

[4] CERN. (Aug. 2020). *High Luminosity LHC Project Plan*. [Online]. Available: https://hilumilhc.web.cern.ch/content/hl-lhc-project

[5] C.-H. Lee, W. V. Sorin, and B. Y. Kim, "Fiber to the home using a PON infrastructure," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4568–4583, Dec. 2006.

[6] I. Papakonstantinou *et al.*, "A fully bidirectional optical network with latency monitoring capability for the distribution of timing-trigger and control signals in high-energy physics experiments," *IEEE Trans. Nucl. Sci.*, vol. 58, no. 4, pp. 1628–1640, Aug. 2011.

[7] M. Krivda *et al.*, "The ALICE central trigger processor (CTP) upgrade," *J. Instrum.*, vol. 11, no. 3, 2016. Art. no. C03051.

[8] B.-G. Ahn and Y. Park, "A symmetric-structure CDMA-PON system and its implementation," *IEEE Photon. Technol. Lett.*, vol. 14, no. 9, pp. 1381–1383, Sep. 2002.

[9] *10-Gigabit-Capable Passive Optical Networks (XG-PON): General Requirements*, document ITU G.987.1, 2016. [Online]. Available: https://www.itu.int/rec/T-REC-G.987.1-201603-I/en

[10] *40-Gigabit-Capable Passive Optical Networks (NG-PON2): General Requirements*, document ITU G.989.1, 2016. [Online]. Available: https://www.itu.int/rec/T-REC-G.989.1-201303-I/en

[11] Y. Akaiwa, *The Elements of Digital Communication System*. 2015, ch. 3, pp. 127–187. [Online]. Available: https://onlinelibrary.wiley.com/doi/book/10.1002/9781119041139

[12] R. C. Bose and D. K. Ray-Chaudhuri, "On a class of error correcting binary group codes," *Inf. Control*, vol. 3, no. 1, pp. 68–79, Mar. 1960.

[13] A. X. Widmer and P. A. Franaszek, "A DC-balanced, partitioned-block, 8B/10B transmission code," *IBM J. Res. Develop.*, vol. 27, no. 5, pp. 440–451, Sep. 1983.

[14] *Frame Alignment and Cyclic Redundancy Check (CRC) Procedures Relating to Basic Frame Structures Defined in Recommendation G.704*, document ITU G.706, 1991. [Online]. Available: https://www.itu.int/rec/T-REC-G.706-199104-I/en

[15] *10-Gigabit-Capable Passive Optical Networks (XG-PON): Physical Media Dependent (PMD) Layer Specification*, document ITU G.987.2, 2016. [Online]. Available: https://www.itu.int/rec/T-REC-G.987.2-201602-I/en

[16] E. Mendes, S. Baron, and C. Soos. (Mar. 2020). *GitLab for TTC-PON Project. CERN. TTC-PON (Rev_1_0_2)*. [Online]. Available: https://gitlab.cern.ch/TTC-PON/official_release

[17] E. B. S. Mendes, S. Baron, D. M. Kolotouros, C. Soos, and F. Vasey, "The 10 G TTC-PON: Challenges, solutions and performance," *J. Instrum.*, vol. 12, no. 02, Feb. 2017, Art. no. C02041.

[18] R. Giordano and A. Aloisio, "Fixed-latency, multi-gigabit serial links with Xilinx FPGAs," *IEEE Trans. Nucl. Sci.*, vol. 58, no. 1, pp. 194–201, Feb. 2011.

[19] D. M. Kolotouros, S. Baron, C. Soos, and F. Vasey, "A TTC upgrade proposal using bidirectional 10G-PON FTTH technology," *J. Instrum.*, vol. 10, no. 04, Apr. 2015, Art. no. C04001.

[20] M. Lipinski, T. Wlostowski, J. Serrano, and P. Alvarez, "White rabbit: A PTP application for robust sub-nanosecond synchronization," in *Proc. IEEE Int. Symp. Precis. Clock Synchronization Meas., Control Commun.*, Sep. 2011, pp. 25–30.

[21] E. B. De Souza Mendes, S. Baron, and M. Taylor, "TCLink: A timing compensated high-speed optical link for the HL-LHC experiments," *PoS*, vol. TWEPP2019, p. 57, 2020, doi: 10.22323/1.370.0057.

[22] D. M. Kolotouros *et al.*, "Metrics and methods for TTC-PON system characterization," *J. Instrum.*, vol. 9, no. 01, Jan. 2014, Art. no. C01015.

[23] E. Mendes and S. Baron. (Mar. 2020). *GitLab for HPTC Project. CERN*. [Online]. Available: http://gitlab.cern.ch/HPTD/hptc