

A Low-Energy Rate-Adaptive Bit-Interleaved Passive Optical Network

Dusan Suvakovic, Hungkei Chow, N. Prasanth Anthapadmanabhan, Dora van Veen, Adriaan J. van Wijngaarden, Tolga Ayhan, Christophe Van Praet, Guy Torfs, Xin Yin and Peter Vetter

Abstract—Energy consumption of customer premises equipment (CPE) in the new generation of time-division multiplexing (TDM) passive optical networks (PON) operating at 10 Gb/s, has become a serious problem both in terms of the global network energy consumption and the CPE battery life. The proposed low energy passive optical network (PON), based on a novel bit-interleaving downstream protocol, reduces the protocol processing energy by a factor of 30 and enables a significant reduction in the total CPE energy consumption over the standard 10 Gb/s PON CPE. The network architecture, protocol and the key enabling techniques for its implementation, including dynamic traffic interleaving, rate-adaptive descrambling of decimated traffic and downsampling clock and data recovery (CDR) circuit, are described. Detailed analysis of the CPE energy consumption and comparison with the standard PON CPE is also included in the paper.

I. INTRODUCTION

The steady increase of the global Internet traffic, network energy consumption has become a growing concern for network operators, networking equipment vendors and chip makers alike. The rising electricity costs put pressure on the profitability of the network operators, and the increased power density cause heat dissipation problems in network gear. The environmental impact of the network's energy consumption is not negligible either, as its current carbon footprint equals that of the airline industry, and it is likely to further increase. In recent years, the scientific community initiated research on network energy consumption [1], [2], [9], [10]. Recent papers analyze and model the energy efficiency of the network and its components, provide insight into the nature and lower bounds of energy consumption in communication networks, and present more energy-efficient networking gear and network architectures.

The work presented here falls into the latter category, while relying heavily on the insights of the other two. Our choice of the topic is driven by the known fact that most of the network energy consumption globally, is consumed in wireless and wireline access networks [1]. Focusing further on the wireline access network, our interest is narrowed to the energy

consuming behavior of passive optical networks (PON), featuring higher energy efficiency compared to other standard fixed access technologies, which makes them a natural starting point for work on the further improvements. Also, considering that, for all PON standards [3], almost 90 % of energy is consumed in the customer premise equipment (CPE) units (also known as optical network termination (ONT)), we further focus on minimizing the PON CPE energy consumption.

This work is influenced by our awareness of the slowdown in power scaling of CMOS technology nodes, making it unrealistic to expect any substantial improvement in energy efficiency of communication systems to result from their implementation in the next CMOS generations.

Power management and virtualization of computing resources are the dominant power saving methods in modern electronics systems. In our study of the PON CPE operation and the ways to improve its energy efficiency, both concepts are taken into account.

Sleep modes, which are a subset of power management techniques, have so far been the only ONT energy saving method adopted by the PON standards bodies [3], [13] or studied by the wider research community [11]. As discussed in Section II, such research has produced evidence of only moderate energy saving potential of the sleep modes. The actual impact of the application of aggressive sleep modes on the quality of service (QoS) and experience (QoE) as well as the complexity of traffic management and the effects of burstiness of PON traffic on the TCP flow performance has not yet been studied thoroughly.

Unlike the previous work, the solution presented in this paper takes a disruptive approach to CPE energy savings that departs from the standardized PON protocols. It redefines the functionality of the CPE and, as a result, achieves large savings that are guaranteed regardless of the traffic volume or patterns, without compromising the quality of service.

As described in Section II, the essence of the new, bit-interleaved PON (Bi-PON) protocol is that it enables the ONT to detect and extract its own downstream traffic by performing a simple, PHY layer downsampling operation, instead of a complicated XG-PON encapsulation method (XGEM) [3] or 10G-EPON media access control (MAC) [13] processing of the entire, mostly unrelated, PON downstream traffic [22].

Another important feature of the proposed Bi-PON protocol is that it allows dynamic changes of the interleaving pattern and the bandwidth allocated to different users, thus adapting to real traffic conditions. By doing so, Bi-PON avoids the bandwidth inefficiency of protocols based on static interleaving of traffic, such as Synchronous Digital Hierarchy (SDH) [6].

Parts of the material in this paper were presented at the IEEE Online Conference on Green Communications (GreenCom), Sep. 2012, and the IEEE Int'l Conf. on Communications (ICC), Budapest, Hungary, June 2013.

D. Suvakovic, H. Chow, P. Anthapadmanabhan, D. van Veen, A.J. van Wijngaarden and P. Vetter are with Bell Laboratories, Alcatel-Lucent, 600 Mountain Avenue, Murray Hill, NJ 07974, USA. Email: dusan.suvakovic@alcatel-lucent.com.

T. Ayhan is with the Photonics and Networking Research Laboratory, Stanford University, Stanford, CA.

C. van Praet, G. Torfs and X. Yin are with the Department of Information Technology INTEC/IMEC, Ghent University, Ghent, Belgium.

This paper is organized as follows. Section II briefly discusses the operation and energy inefficiency of standard XGPON1 ONT and presents the Bi-PON concept, operation and its key features. Section III presents the design of three key modules that enable Bi-PON protocol implementation. Section IV presents the extension of Bi-PON protocol across network hops. The analysis of the energy consumption of Bi-PON and a comparison of the energy consumption in XGPON is presented in Section V, which is corroborated by experimental results obtained from setups with ASIC and FPGA versions of the Bi-PON and XGPON ONT. Section VI provides the conclusions and summarizes the key features of Bi-PON.

II. PROTOCOL DESIGN FOR ENERGY EFFICIENCY

A. Energy Efficiency Limitation in Standard PON Protocols

In all standard TDM PON protocols, it is assumed that the OLT sends a sequence of arbitrarily ordered packets. Since this sequence is received by all ONTs, each one of them needs to check the destination address of every received packet to determine whether it matches the address of its own user. It then drops all the packets sent to addresses other than its own. Although the total number of packets each ONT selects to forward to its user is rather small, all packets sent by the OLT have to undergo most of the ONT processing, including: deserializing, word alignment, descrambling, forward error correction (FEC) decoding, packet delineation and PON-specific MAC parsing, as illustrated in Fig. 1. Since the standard subscriber count per PON ranges from 32 to 128, as much as 97% to 99.2% of all processed payload is dropped. When compared with processing required to receive the same amount of user payload over a point-to-point link, the standard PON ONT performs, on the average, 30-99 times more work, consuming that much more energy. Or, in the packet-centric interpretation, the energy consumption per packet equals the processing energy in the point-to-point link multiplied with the number of active ONTs.

Obviously, the total ONT energy could be reduced if each ONT is put in the sleep state during time intervals in which it is not supposed to receive traffic. To maximize the length of the ONT sleep intervals, it would be necessary for the OLT to rearrange the order of downstream transmission by grouping packets by destination, in large, back-to-back bursts.

Driven by this idea, both the IEEE and ITU-T PON standards bodies included sleep state control protocols in their respective standards for 10Gb/s PON as an energy saving instrument [3], [13].

In an ideal case, each ONT would be awake only while its own traffic is being sent downstream, and the total downstream energy consumption would be reduced to that of the point-to-point links. However, recent extensive research [8]–[10], [33] indicates that energy savings attainable by applying sleep modes in TDM PON are far lower, as a result of the following issues.

Firstly, the ability to schedule traffic for energy savings is limited by the overriding requirement for quality of service (QoS). Traffic management for QoS produces a schedule for

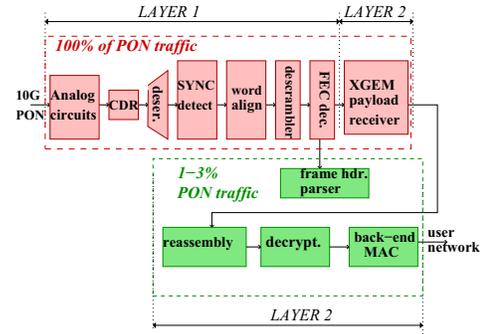


Fig. 1: Downstream architecture of XG-PON ONT.

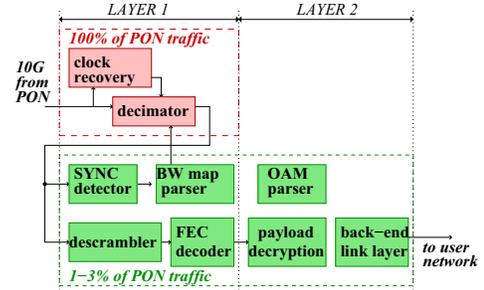


Fig. 2: Downstream architecture of Bi-PON ONT.

downstream transmission different from the one accommodating sleep modes.

Secondly, the duration of the ONT's transitions between the awake and sleep state is not negligible, while the energy consumption during that time is equal to that of the awake state. Since the ONT is not able to receive traffic during such transitions, they result in net waste of energy proportional to the frequency and duration of the transitions.

Thirdly, the arrival of upstream traffic will wake up the ONT, as needed to avoid excessive latency and packet loss, leading to a further reduction in achievable energy efficiency.

In line with the described limitations, analysis and simulation of cyclic XG-PON sleep algorithm [9], has demonstrated that substantial energy savings in the order of 30-70% can be achieved, only in the conditions of very low traffic.

In addition to the factors limiting the achievable energy savings, aggressive application of sleep modes with long sleep periods necessitates a large OLT memory, which increases the OLT cost and power consumption. Similarly, each ONT must have a relatively large packet buffer in order to store long bursts of downstream traffic, which is adding to its cost and energy consumption.

B. Bit-Interleaving Protocol

Whereas in all standard PONs, sleep modes are used to mitigate the inefficiency of the given protocol, Bi-PON protocol has been designed specifically for green operation. The ideal CPE from the point of view of home energy consumption is a completely virtualized one, with wall network outlets, but no active equipment on the premises. The rationale for this idea is that performing the virtual ONT functionality as part of the

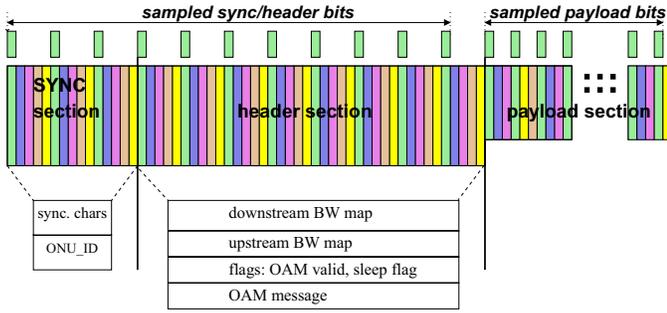


Fig. 3: Bi-PON frame structure.

OLT in the operator’s central office (CO) would not increase the OLT energy significantly given that its packet processing and traffic management functions are already sorting packets by destination. On the other hand, the complete ONT energy consumption would be eliminated, leading to approximately 90% energy savings across the access network.

However, the hidden cost of the virtual CPE implementation would be in the deployment of DWDM with thousands of different colors, which would result in a drastic increase of CO energy consumption as well as a prohibitively high capital expenses for building such access networks.

The next best approach would be to limit the ONT functionality to “tuning into its own channel”, using the analogy with the widespread low power AM/FM radio receiver. Such sampling operation can also be viewed as an equivalent to color filtering in a DWDM receiver, however requiring the use of active electronics. The baseband TDM equivalent of such “channel” would be a sequence of bits sampled from the original, full rate, downstream bit sequence by using a certain sampling rule.

Obviously, the bit ordering rule simplest for sampling, is a sequence of equally spaced bits that can be selected by using a slow, downsampling clock “tuned” properly in terms of frequency and offset [14]. It is intuitively clear that implementation of such operation requires far less energy than the ONT in a standard PON, because it completely eliminates processing of unrelated traffic. Also, since the operation of this scheme does not rely on the use of sleep modes, it does not have adverse effects on the QoS. As illustrated in Fig. 2, such operation would allow the ONT to find and sample its own traffic in the clock and data recovery stage, thus dropping the unrelated traffic much earlier than the XG-PON ONT of Fig. 1. As a result, the amount of processing and energy consumption in the subsequent ONT stages would be reduced dramatically.

The problem to be solved is to specify and implement a downstream protocol that will tell each ONT how to select its own “channel” or “lane” (the term used in the rest of this paper), while preserving the dynamic downstream bandwidth allocation capability comparable to that of the conventional TDM PON.

The protocol features and operation are explained in this section, whereas the key implementation solutions are described in Section III.

The structure of the Bi-PON downstream frame is shown in Fig. 3, where different colors denote different lanes i.e.

sequences of equally spaced bits sent to different ONTs. The frame has a fixed length and consists of the synchronization character section, the header section and the payload section. Interleaving of the entire frame, including the synchronization and header section, allows for the use of a slower receiver clock in the ONT. A lane plays the role of a virtual point-to-point connection between the OLT and the ONT. A Bi-PON lane is specified by its downstream bandwidth (BW) map parameters *offset rate* that are included in the downstream frame header. *offset* is defined as the bit-distance of the position of the first lane bit from the position of the first payload bit in the frame, whereas *rate* represents the bit-distance between adjacent lane bits.

In the synchronization and the header sections, the lane assignment is static, with each ONT being assigned the same constant rate and equal number of bits per section. Further, the offset of each lane in these sections, from the first bit of the frame, is chosen to be equal to its identification number (ONU_ID). In the payload section, the lane assignment is dynamic and the bandwidth assigned to each ONT may be changed in every frame period.

The synchronization section is a collection of bit-interleaved, independent frame synchronization characters. Each such character consists of a constant bit sequence followed by the ONU_ID number. Transition density necessary to maintain proper clock recovery in the receiver is ensured by inverting all synchronization characters associated with odd ONU_ID numbers. The ONT synchronization is achieved by initial arbitrary phase selection for the header sampling clock, followed by detecting the synchronization character and calculating the relative phase of its own synchronization lane based on the ONU_ID value found in this character. This way, the initial ONT synchronization is completed in two frame periods. Details of the synchronization algorithm are explained in [22].

The header section of the Bi-PON frame received by one ONT follows the ONT’s synchronization pattern in the lane with a constant rate ($rate_h$) and offset ($offset_h$) equal to its own ONU_ID. The header includes downstream and upstream bandwidth allocation fields, the operations, administration and management (OAM) message field as well as any number of other optional fields. A synchronized ONT reads only the contents of its own lane from the header section of the Bi-PON frame.

The downstream bandwidth map field contains the ($rate_p$, $offset_p$) pair, specifying the location of the ONT’s downstream payload lane in the current frame and enabling the ONT to extract and process only its own payload. The allowed bit rates for user payload lanes were chosen to be equal to the full downstream rate divided by a power of two, i.e., $2^{-rate_p} \times 10 \text{ Gb/s}$, where $rate_p = 0, 1 \dots$

The use of the downstream bandwidth map enables flexible and dynamically adjustable bandwidth allocation. The lane rate for each ONT can be changed independently in every frame period and in a wide range from zero to the maximum bandwidth allowed. This mechanism guarantees flexibility of bandwidth allocation comparable to that of XG-PON and other PON protocols that do not include an explicit downstream

bandwidth map, but still perform downstream bandwidth allocation in accordance with the traffic volume and traffic management policies.

The header section may include more than one downstream bandwidth map for one ONT, allowing the ONT to receive multiple payload lanes to be used for different services, unicast, multicast and broadcast traffic or different users served by the same ONT.

In our experimental implementation of the Bi-PON protocol, the allowed bit rates for a user payload lane were limited $2^{-(rate_p+3)} \times 10 \text{ Gb/s}$, where $rate_p = 0, 1 \dots 7$, ranging from 9 Mb/s to 1.25 Gb/s. The maximum lane rate of 1.25 Gb/s in the experiment does not result from any fundamental limitation of the proposed protocol but rather from the intent to create a prototype compatible with a standard XG-PON ONT with a GigE UNI [35]. Since the choice of such UNI speed makes the latency reduction benefit of the 10 Gb/s rate unavailable to the user, the rate of the Bi-PON payload lane can be limited to 1.25 Gb/s without performance degradation. However, since the Bi-PON protocol allows simultaneous allocation of multiple downstream payload lanes to one ONT, the total receive bandwidth of one ONT is not limited to the rate of a single lane and can be scheduled to be as high as the full PON rate. For example, an ONT capable of receiving eight payload lanes simultaneously may be designed with eight 1-GigE user interfaces and be capable of forwarding a peak traffic of 10 Gb/s.

Since extraction of the user payload from the Bi-PON downstream frame is essentially a layer 1 operation, which does not require any particular packet formatting, it is generally possible for the ONT to avoid any parsing, framing or line encoding of the received payload, provided that it has already been line encoded by the OLT, according to the standard supported by the user network interface (UNI). For example, for the chosen selection of lane rates, Bi-PON protocol allows transparent forwarding of Gigabit Ethernet (GigE) traffic 8b10b-encoded for 1000BASE-X optical GigE physical layer.

Alternatively, as implemented in our experimental ONT designs, a “lightweight” link layer protocol consisting of 3-byte packet delineation headers, can be used, in combination with an appropriate UNI line-encoder.

As such, it is also possible to provide multiple lanes to a CPE with multiple user ports, that may include the phone line, a TV coax line, a wireless (WiFi) port and several copper and/or plastic optical fiber (PoF) Ethernet ports. This approach is disruptive as it removes the conventional Ethernet switching technology from inside a customer premise. Common gateway functions, such as network address translation (NAT), firewall and switching/routing can be transferred to a virtual home gateway server at the access node or edge node. Such architecture not only simplifies the network inside the customer premises, but it also reduces the total power consumption, because the forwarding in the CPE is largely simplified from layer 3, down to layer 1 and because the processors in all the CPE are replaced by a single server in the network that achieves a much higher energy efficiency by multiplexing the same tasks for about one thousand homes [29].

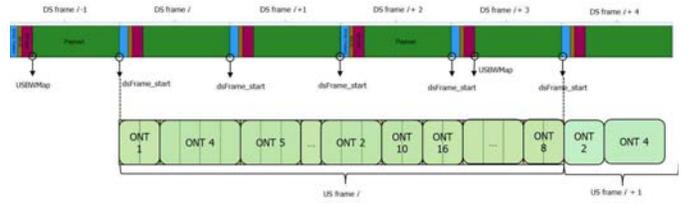


Fig. 4: Bi-PON upstream PHY frame.

Further details of the Bi-PON downstream protocol are described in [22].

With the OLT being the sole recipient of the entire Bi-PON upstream traffic, there is no inherent overhead in energy consumption associated with the processing of the upstream frame. Consequently, Bi-PON upstream transmission in Bi-PON does not require separation of traffic by bit-interleaving. Such interleaving would not even be possible since coordination of transmission by different ONTs at the bit level would be extremely difficult. The protocol for Bi-PON upstream transmission is optimized for simplicity and low-power consumption. It is based on a time-slot based burst transfer. As in XG-PON, the upstream line rate is specified to be one quarter of the downstream line rate. However, unlike in the case of XG-PON upstream frame, the duration of the Bi-PON upstream frame is specified to be four times that of a downstream frame (i.e., $4 \times 125 \mu\text{s}$), to reduce the bandwidth map processing overhead. An upstream Bi-PON frame is divided into a number of equally sized time slots. The OLT schedules the upstream transmission of all ONTs in terms of the number of allocated contiguous slots and their position in the upstream frame. An ONT determines its transmission turn and duration from its upstream bandwidth map sub-field embedded in the downstream header lane. An upstream allocation is an ordered pair of parameters (*start* and *length*), where the former specifies the position of the first time slot of the assigned burst and the latter indicates the size of the allocation, expressed as the number of time slots. Fig. 4 illustrates the relationship between the upstream and downstream transmission.

Following the US BW map is the OAM field. The length of the OAM field depends on the OAM message. Bi-PON protocol has an OAM message to enable a flexible control of sleep modes for the ONTs downstream operation. The OLT can instruct an ONT to enter a one-time or periodic sleep state, in granularities ranging from a fraction of a frame to hundreds of consecutive frames.

C. TDM PON at Higher Rates

Although the initial Bi-PON protocol design was driven only by an intent to improve the energy efficiency of 10 Gb/s ONT's, technical challenges associated with scaling of the TDM PON to a higher rate, such as 40 Gb/s, are transforming Bi-PON into an enabling technology.

Porting of the XG-PON1 protocol to an ONT running at a quadrupled speed is considered to be prohibitively expensive in terms of power consumption and hardware complexity, alike.

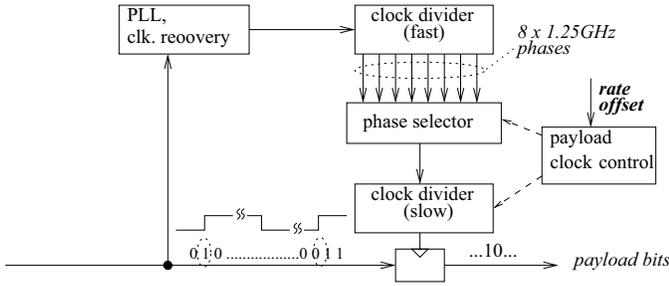


Fig. 5: Bi-PON CDR architecture.

As a result, PON standardization bodies such as FSAN, are abandoning scaling of the TDM PON and working adopting a combination of TDM and WDM technologies for further PON scaling.

Unlike in the case of XG-PON1, Bi-PON protocol implementation for the quadrupled bit rate would not represent a serious challenge given its simplicity and inherent energy efficiency. The prospect of achieving a TDM PON at the 40 Gb/s rate is further improved by the recently proposed duobinary modulation scheme [25], [26] that enables inexpensive implementation of 40 Gb/s receiver circuits for use in 40G PON ONTs. In the context of this work, the importance of such development is in that it enables its extension to higher capacity networks consisting of two or more cascaded PON stages and running the bit-interleaving protocol, as discussed in Section IV.

III. BI-PON ENABLING SOLUTIONS

A. Clock and Data Recovery Architecture

The design of the clock and data recovery (CDR) circuit is key to Bi-PON's ability to eliminate unrelated traffic early in the ONT downstream path. The downsampling CDR designed for this purpose, also achieves significant energy savings in its internal operation when compared to a conventional 10 Gb/s CDR, because it does not include deserializing and word alignment functions, used with all conventional CDRs that receive traffic at 10 Gb/s.

Bi-PON protocol provision that spaces bits destined for one ONT at a regular distance from each other, enables simple and low power CDR operation. Bi-PON CDR, shown in Fig. 5 adjusts the frequency and phase of its sampling clock according to the **(rate, offset)** parameters assigned to its payload lane. The phase of the sampling clock is controlled with the resolution of the fast, 10 GHz clock generated by the CDR. In our design [20], the use of the fast clock in the Bi-PON CDR is minimal and limited only to driving of a clock divider that produces the 8 different phases of the 1.25 GHz sampling clock. Depending on the allocated bit rate, one of the phases of the sampling clock is further divided to produce an even slower sampling clock for payload sampling. Proper phase synchronization of the generated payload clock is achieved by using the header clock as a reference.

The maximum frequency of the sampling clock is 1.25 GHz, which is 8 times lower than the Bi-PON 10 Gb/s line rate. Reception of the header section is performed using a separate

sampling clock at 39 MHz, the phase of which is adjusted according to the value of ONU_ID. These frequencies are already low enough not to require deserializing of the header and payload bit traffic for further processing by the Bi-PON ONT.

B. Descrambling with Bit-interleaved Traffic

PON downstream frames must be scrambled in order to ensure the transition density and balance of 0s and 1s in the downstream frame, which is required for reliable data recovery at the ONT. The addition and multiplication in this section are over the binary field and thus represent logical XOR and AND operations, respectively. The XG-PON standard specifies frame synchronous (additive) scrambling to be performed by the OLT by adding a pseudo-random bit sequence (PRBS) to the frame bit sequence [3]. The PRBS is generated by a linear feedback shift register (LFSR), specified by its initial state and generator polynomial. Descrambling in XG-PON, performed by each ONT, is exactly the same operation as scrambling which when applied to received scrambled frame, produces the original unscrambled frame.

In Bi-PON, scrambling is performed in the same way as in XG-PON. However, a different descrambling technique is required because a Bi-PON ONT operates at a lower clock rate and, more importantly, receives only a sampled set of bits after the decimator, which extracts only the bits corresponding to the Bi-PON lane assigned to the ONT. Therefore, the Bi-PON ONT requires a special rate- and offset- adaptive descrambler which can operate directly on the decimated data. Our solutions to this problem, described below, take advantage of some properties of shift-register sequences. These are the skipping descrambler and the descrambler using helper bits.

The main idea of the skipping descrambler is to skip forward by multiple cycles (w.r.t. the original LFSR in the scrambler) within a single clock cycle such that we produce exactly the decimated PRBS required for descrambling the decimated data received at the ONT.

Suppose an n -bit LFSR is used in the scrambler. The LFSR state $\mathbf{s}(m)$ in the m -th clock cycle is defined by the states $s_i(m)$ of the n registers as $\mathbf{s}(m) = [s_1(m), s_2(m), \dots, s_n(m)]^T$. Suppose the linear recurrence satisfied by the LFSR is given by

$$s_1(m+1) = \sum_{i=1}^n a_i s_i(m), \text{ where } a_i \in \{0, 1\} \quad \forall i.$$

Then, it is clear that $\mathbf{s}(m+1) = A \cdot \mathbf{s}(m)$ where

$$A = \left[\begin{array}{c|c} a_1 & a_2 & \dots & a_n \\ \hline & I_{n-1} & & 0_{n-1} \end{array} \right],$$

and I_{n-1} and 0_{n-1} are the $(n-1) \times (n-1)$ identity matrix and the zero vector of length $n-1$ respectively. As a result, for a constant skip value k , the LFSR state k cycles ahead from the current state can be obtained as $\mathbf{s}(m+k) = A^k \cdot \mathbf{s}(m)$. Due to the linearity, it is obvious that $\mathbf{s}(m+k_1+k_2) = A^{k_1} \cdot (A^{k_2} \cdot \mathbf{s}(m))$. Consequently, for any r -bit number $k = \sum_{i=0}^{r-1} b_i \cdot 2^i$

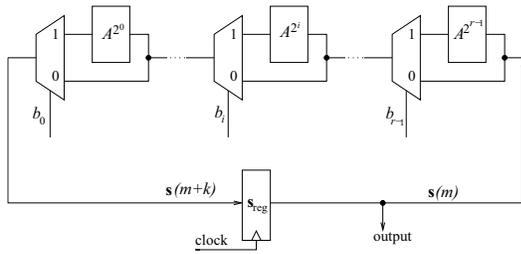


Fig. 6: Skipping Bi-PON descrambler.

with $b_i \in \{0, 1\}$, we can calculate $s(m+k)$ from $s(m)$ in the following way:

$$s(m+k) = (B_{r-1} \cdot (B_{r-2} \cdot \dots (B_0 \cdot s(m))))), \quad (1)$$

where

$$B_i = \begin{cases} A^{2^i}, & \text{if } b_i = 1, \\ I_n, & \text{otherwise.} \end{cases} \quad (2)$$

The skipping descrambler illustrated in Fig. 6 essentially implements the equations (1) and (2). Based on the **offset** value, the descrambler first adjusts its initial phase to an arbitrary number of shift-steps from the beginning of the frame in a single clock cycle by using the control inputs b_0, \dots, b_{r-1} . Subsequently, the descrambler is run by setting the control inputs to a constant value matching the specified **rate** for decimation.

In the case of the descrambler using helper bits, we suppose that scrambling is performed by a n -bit LFSR which generates a maximum-length sequence (MLS), which is typical in practice. We also assume the decimation **rate** k is a power of 2, which is true for our Bi-PON implementation. It is known [23] that if an MLS is decimated with rate k starting with any offset, then the resulting sequence of bits is the same as the original MLS except for a phase shift when k is a power of 2. Therefore, in this case, a simpler descrambling solution is potentially possible by using the same LFSR that is used in the scrambler to produce the decimated MLS for descrambling the user's decimated bit stream. However, the main problem is to initialize the LFSR in the descrambler to the appropriate state so that the MLS with the desired phase shift (i.e., the decimated MLS) is produced. A quick observation reveals that the desired initial state, say for an n -bit LFSR, is exactly the same as the first n bits of the decimated MLS.

Based on this observation, it is assumed in this section that the sender has the capability to insert a few additional bits, termed *helper bits*, in the bit stream for each ONT which can be used to assist the ONT in determining the appropriate LFSR initialization. In the payload sections of the frame, it is indeed feasible to insert a few additional bits before the actual payload begins. The effective throughput to the ONT is not affected much because only a small number of additional bits are needed. The descrambling solution works as follows.

For every transmitted frame, before sending it to the scrambler, the OLT first inserts n helper bits with value 0 at the beginning of the payload data for every ONT. These helper bits are inserted exactly in the Bi-PON lane for the ONT, i.e.,

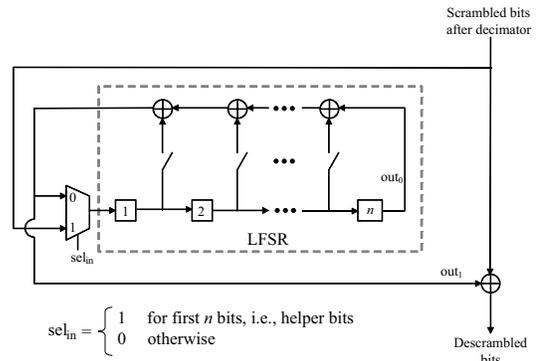


Fig. 7: Bi-PON descrambler using helper bits.

at bit intervals specified by the **rate** starting from the **offset**. This bit-interleaved frame is then sent to the scrambler¹.

The descrambler implementation is shown in Fig. 7. At the start of every frame, the ONT uses the first n received bits from its Bi-PON lane to initialize the LFSR in the descrambler. This is accomplished by shifting-in the first n bits one-by-one using the multiplexer with sel_{in} as defined. When the OLT inserts helper bits with value 0, the first n bits arriving at the ONT are exactly the first n bits of the decimated MLS corresponding to ONT, and thus provide the appropriate initialization. Finally, the output sequence out_1 is used to descramble the received data. Note that the first n bits of the descrambled data are ignored in every frame as they correspond to the helper bits.

For further details on these descrambling solutions, and other alternatives, we refer the reader to [24].

C. Real-time Interleaver

Dynamic scheduling and real-time interleaving of downstream traffic are key to Bi-PON performance. The required interleaving at 10 Gb/s according to bandwidth maps that randomly change between successive frames, makes Bi-PON interleaver design and implementation particularly challenging. Failure of the interleaver to keep pace with the line rate would reduce the effective bandwidth of the PON downstream link, erasing the bandwidth advantage of the 10 Gigabit optics and potentially degrading the quality of service. The architecture of a real-time Bi-PON interleaver, described here, has been verified to meet all mentioned functional and performance requirements, while using conventional hardware building blocks.

Downstream architecture of the Bi-PON OLT is shown in Fig. 8. As for a conventional PON, packets arriving from the core network are stored in the main memory before being forwarded, over the PON, to the end user. While packets are stored, common access node functions, such as packet processing, traffic management and downstream scheduling of packets is performed. In Bi-PON, an additional step is taken, whereby time windows are allocated for traffic associated with each user. This function is performed by the Bi-PON lane scheduler, which converts the bandwidth distribution information from

¹The method can be further optimized such that helper bits need to be inserted only in those frames where the ONT's allocation changes.

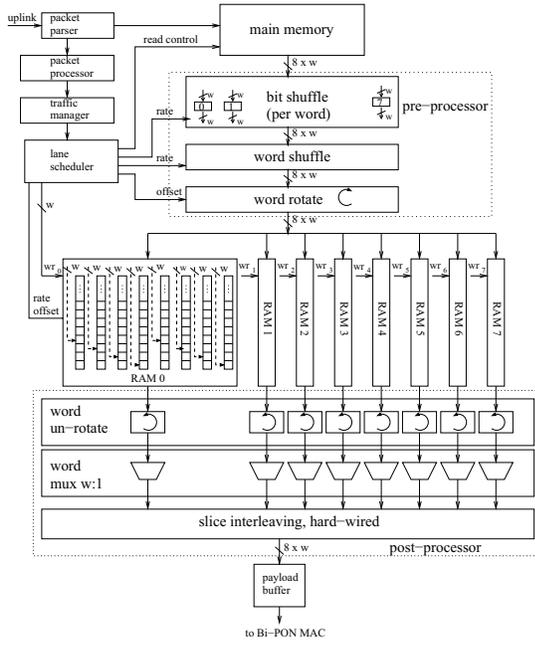


Fig. 8: Bi-PON OLT downstream architecture with interleaver.

the traffic manager into a downstream bandwidth map, with each user service being assigned a specific bit rate and offset in the Bi-PON frame. As shown in Fig. 8 the **rate** and **offset** parameters used as the only control input for the Bi-PON interleaver. Bi-PON interleaver architecture consists of eight identical slices, each one including one interleaving RAM module as well as pre-processing and post-processing logic. The 8-way slicing simplifies interleaving by taking advantage of the protocol specification that the maximum rate lane carries 1/8 of the total traffic. Therefore, one RAM module stores bits belonging either to one maximum rate (1.25 Gb/s) lane or to a collection of lower bandwidth lanes that are non-overlapping subsets of the same 1.25 Gb/s lane. Each RAM module in each slice is further vertically divided into 8 separately addressable w -bit wide sub-modules, where different write addresses are used for each sub-module in order to achieve interleaving.

The goal of pre-processing is to create an arrangement of the bits in the interleaver memory that is as close to the final interleaving as possible, while avoiding having multiple bits contending to be written over the same RAM column, since such contention would degrade the interleaver throughput by extending the RAM write time to multiple clock cycles.

Pre-processing at each slice consists of three stages: re-ordering of the input bits within each w -bit word, re-ordering of w -bit words and finally, rotational shifting of w -bit words.

In the post-processing stage, the pre-interleaved contents of the eight slices are read one $8 \cdot w$ -bit word at a time, followed by a rotational shift, inverse to the one performed in the pre-processing stage. Subsequently, 1 bit out of each w -bit word is selected using a uniform bit selection rule across all 8 w -bit words. Finally, the selected bits from each RAM module are statically interleaved by hard-wiring.

The required re-ordering of both bits and words, performed in pre-processing stage, is found to match the well known

“perfect shuffle” [19] operation, repeated the number of times determined by the value of the **rate** parameter of the particular lane.

The correctness of the interleaving procedure described above, has been verified in simulation and implementation.

Given the repetition period of the Bi-PON interleaving pattern and the chosen architecture, the minimum size of a RAM module is $128 \cdot BW$, where $BW = 8 \cdot w$ is the RAM word width. For example, if $BW = 128$ bits, the size of each RAM module will be 2kByte and the total on-chip memory size will be 16kByte, which is considered to be small and inexpensive both in the FPGA and ASIC implementation. The complexity of the control, pre-processing and post-processing logic is also rather low and its size is estimated to be about 4,000 logic gates, making the total complexity of the interleaver low.

The operation of the Bi-PON interleaver is controlled solely by using the **(rate_p, offset_p)** parameters. The control logic fills up each one of the 8 RAM modules with payload lanes assigned to a particular RAM, typically with segments of different packets, which can be as small as w bytes or as large as 2 kB.

IV. CASCADED BIT-INTERLEAVING PON

In this section, the extension of the Bi-PON protocol to a network with multiple PON stages, separated by active repeaters, is described. Possible application of such cascaded PON are discussed.

A. Principle of Operation

The only part of the Bi-PON ONT that does not save energy as a result of running the bit-interleaving protocol is its PON interface, including the optic and electronic components equivalent to those used in XG-PON ONT. These components are exposed to the full 10 Gb/s line rate, consuming as much energy as the equivalent components used in a standard XG-PON ONT.

Cascaded Bit-Interleaving PON (CBi-PON), shown in Fig. 9, is an extension of the Bi-PON. It consists of two PON or more stages, separated by several active repeater devices. The role of the repeater is to enable reduction of the downstream line rate of the Bi-PON ONT, to a rate compatible with the rate of the ONT’s UNI interface. This way, lower speed grade and consequently, lower power transceiver components can be used in the ONT PON interface.

The repeater has two ports, one for each PON stage. It receives traffic from the OLT, at the full PON rate 10 Gb/s, over its uplink interface. Its downstream receiver operation is identical to that of a Bi-PON ONT, described in Section II-B. Each repeater receives the interleaved downstream frame and fetches the bits belonging to payload lanes assigned to it in the downstream bandwidth map, using the rate and offset information from the frame header. The repeater then descrambles the decimated bits. If forward error correction (FEC) is used for the first stage, the descrambled bits are processed by the Forward Error Correction (FEC) decoder and subsequently forwarded to the second stage network from the repeater’s downlink interface.

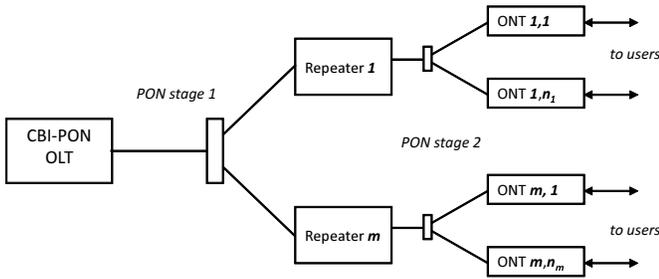


Fig. 9: Cascaded bit-interleaving PON.

The cascaded bit-interleaving scheme requires the OLT to create a nested downstream frame structure where payload of the first-stage frame consists of a collection of interleaved second-stage Bi-PON frames. The OLT first forms the second-stage frames, using the decimation rate and offset information of the end-ONTs. The structure of these frames is equivalent to that of the original Bi-PON frame, but with a shorter payload section. Once the second-stage frames are formed, their payload and bandwidth map sections are scrambled. Then, the OLT interleaves all second-stage frames and places them into the payload section of the first-stage frame, which is subsequently FEC encoded, scrambled and transmitted over the PON.

The OLT schedules and assigns upstream allocation to all repeaters as well as the individual end-ONTs, by embedding the upstream bandwidth maps in the header sections of their respective downstream frames. Each end-ONT sends its upstream data to the repeater in the specified time interval, in a burst. The repeater buffers the upstream traffic received from its end-ONTs until the beginning of its own upstream transmission window and then forwards its buffer content to the OLT, again in the burst-mode. It should be noted that the repeater simply forwards its upstream data to the OLT, without performing word alignment or decoding, which is possible because no data is processed at the repeater. Since in the CBI-PON architecture, each end-ONT processes its downstream at a rate lower than the OLT transmission rate and the distance between the ONT and the repeater can be short (e.g. less than 100 meters), the ONT PON interfaces (optical front-end transceiver, oscillators, PLLs, transimpedance amplifier and limiting amplifier) consume less power. The repeater power overhead is low because it is shared by a number of end-ONTs, whereas its complexity is low given that the routing functionality is centralized in the OLT.

B. Metro-Access Convergence

CBI-PON can be used to reduce the complexity of access nodes between the first mile and the metro-aggregation network, by eliminating the need for L2 switching, packet processing, buffering and traffic management in an access node. This way, the power consumption and deployment costs would also be significantly reduced.

Various long reach PON systems have been reported that extend the reach and the splitting factor by using power hungry optical amplifiers or optical-electrical-optical (OEO)

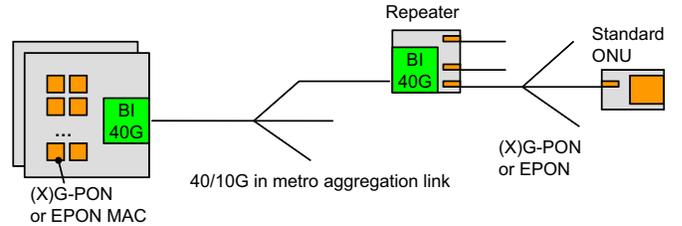


Fig. 10: A 40 Gb/s Bi-PON in metro aggregation network carrying standard GPON format for the first mile segment.

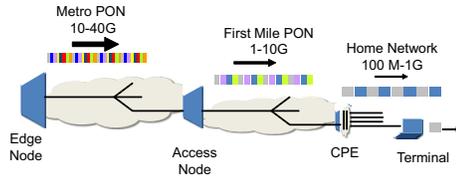


Fig. 11: Cascaded bit-interleaving applied to metro, access, and shared internal bus of customer premises equipment (CPE).

repeaters (e.g. [27]). The increased line rate and greater optical budget, needed to achieve the higher splitting factor, significantly add to the cost of the optical transceivers at the ONU. The advantage of bit-interleaving in such long reach access architecture is that the OEO repeater down-samples the higher rates (e.g., 10 Gb/s up to 40 Gb/s) in the metro aggregation section to lower rates (e.g., 1 Gb/s up to 10 Gb/s) in the first mile, hence relaxing the requirements and cost of the ONU transceivers as well as the repeater transceivers facing the drop side. A further advantage of this approach, illustrated in Fig. 10, is in that it supports legacy ONU and reuses standard PON MAC implementations in the OLT. This is achieved by transparently carrying the standard protocols for the first mile segment, over the bit-interleaving network across the metro aggregation section.

A more disruptive metro-access architecture, shown in Figure 11, uses the CBI-PON protocol for switching of traffic across the entire path from the edge node to the customer premises and possibly, even into the home network. A hierarchical bit-interleaver at the edge node arranges the bits such that one or more levels of repeaters can make selection of bits to be forward at a lower rate, again in an interleaved format.

V. DISCUSSION OF ENERGY CONSUMPTION AND EXPERIMENTAL RESULTS

A. Digital Power Consumption Models

To obtain a better insight into the nature of dynamic power consumption P_{daB} of the Bi-PON ONT in the active state, we have modeled its power consumption with equation (4) relative to the dynamic power consumption P_{daX} of the XG-PON ONT, using the well known formula [34] for dynamic power consumption is expressed in equation (3). In equation (3), f_{clk} represents the digital clock frequency, C_{sw} is the average switched capacitance and V_{dd} is the supply voltage.

$$P_d = f_{clk} \cdot C_{sw} \cdot V_{dd}^2 \quad (3)$$

$$P_{daB} = P_{daX} \cdot \frac{f_{clkB}}{f_{clkX}} \cdot \frac{C_B}{C_X} \cdot \left(\frac{V_{DDB}}{V_{DDX}} \right)^2 \quad (4)$$

In equation (4), the dynamic power P_{daB} of the Bi-PON ONT is expressed as the dynamic power of the XG-PON ONT P_{daX} multiplied by three scaling factors, each one of which is ≤ 1 . The first factor is the ratio of their respective operating clock frequencies f_{clkB} and f_{clkX} . Whereas f_{clkX} is constant, the Bi-PON protocol and ONT design enable scaling of f_{clkB} proportionally with the user traffic. In the ASIC implementation described above, the ONT clock frequency is adjusted to 10 MHz to process user downstream traffic rate of 10 Mb/s. Assuming the value of f_{clkX} of at least 155 MHz, the clock scaling alone will ensure 15-fold energy reduction at 10 Mb/s. Further energy saving enabled by the Bi-PON protocol is the result of the reduction of the ONT logic and total switched capacitance. In the FPGA implementation experiment described above, the total Bi-PON ONT was only 5% of that of the XG-PON ONT, suggesting a similar value of $\frac{C_B}{C_X}$ and explaining the dynamic power savings achieved in both ASIC and FPGA implementations.

The last factor in (4) is the ratio of squares of Bi-PON and XG-PON ONT supply voltages V_{DDB} and V_{DDX} , which indicates a further opportunity for energy saving available to the Bi-PON architecture, achievable by using a reduced digital supply voltage whenever $f_{clkB} < f_{clkX}$, allowing longer circuit delays. The use of dynamic voltage-frequency scaling (DVFS) has not been included in our experiments however, by a conservative estimate, additional factor of 2 in energy savings should be achievable.

$$P_{sram} = f_a \cdot N_{bl} \cdot C_{bl} \cdot V_{swing} \cdot V_{DD} + N_{cell} \cdot I_{leak} \cdot V_{DD} \quad (5)$$

ONT memory energy consumption generally depends on the memory size and locality. Energy consumption of off-chip RAM is much higher than that of the on-chip RAM, due to a drastically increased memory access energy, large memory size and typical use of DRAM. Here it is assumed that current technology allows on-chip implementation of SRAM of sufficient size for packet buffering in PON ONT and that Bi-PON and XG-PON ONTs both use only on-chip memory. SRAM power consumption, given in (5), is the sum of the dynamic and static power components, where the dynamic component depends on the frequency of memory access f_a , the number of bit lines N_{bl} , total capacitance of the bitline C_{bl} , bitline precharge voltage V_{swing} and the supply voltage V_{DD} . The static component, caused by leakage current in the SRAM cells, depends on the cell leakage current I_{leak} , number of SRAM cells N_{cell} and the supply voltage.

Each PON ONT is assumed to include two RAM units in the downstream path: FEC-RAM for temporary storage of downstream traffic during the forward error correction decoding and PKT-RAM-DS for storage of selected downstream user traffic enqueued for transmission at the user interface. An additional SRAM unit (PKT-RAM-US) is needed for storage of upstream packets. Assuming that PKT-RAM-DS is sized to store the complete payload of one XG-PON downstream frame, the total memory requirement for the XG-PON ONT is assumed to be about 2 Mbits and its total dynamic and static power is

estimated to be between 1 and 2 mW [31], which is negligible compared to P_{daX} . SRAM power in the Bi-PON ONT are even lower, because the maximum Bi-PON downstream user size is 8 times lower than that in XG-PON, which would allow for proportional scaling of the FEC-RAM and PKT-RAM-DS sizes.

B. Optical Interface Power

The downstream PON optical interface (OI) consists of a photodiode, transimpedance amplifier, a limiting amplifier and a clock recovery (CR) circuit. Since the bit rates and optical power budget are the same for Bi-PON and XG-PON, the same components can be used for the implementation of both ONTs, and their total power in the active state P_{o10G} is the same. It should be noted that, for the purpose of analysis, we have separated the clock and data recovery functions because the data recovery function is implemented as a digital circuit and is included in the digital power consumption model. Also, as explained in Section III, the complexity and power consumption of data recovery differs significantly for the two protocols. Whereas Bi-PON data recovery is a simple downsampler using a slow payload clock, the same function in XG-PON includes large and power-hungry demultiplexing and word alignment logic.

The power of the Bi-PON and XG-PON ONT OI in the active state is considered to be 750 mW, where PD, TIA and LA is estimated to be about 625 mW [33] and the CR power is assumed to be 125 mW based on our ASIC measurements.

The upstream OI, which consists of a laser and a laser driver circuit, contributes another significant component to the total power consumption of the Bi-PON ONT. Current commercial burst mode laser drivers fail to save power between the the ONT upstream transmission bursts. The reason for this is that it typically takes these circuits a few milliseconds to reach stable operation when turned on, which is too slow for a power saving burst-mode operation. To be able to switch the laser on and off fast enough, the drivers typically steer their output current between the laser and a shunt resistor. Therefore, although fast laser switching is achieved, no power saving takes place between bursts. This results in the constant power consumption, which is typically about 700 mW for the 2.5 Gb/s XG-GPON1 optical interface.

Recently, the design of a 10 Gb/s burst-mode laser driver that can switch fast enough for XG-PON burst-mode operation, has been reported [36]. This new circuit consumes 66 mW in the stand-by mode and 1116 mW in the active mode. Since the average upstream transmission time of one ONT in a 64-split PON cannot exceed 16% of the total time, the use of this driver would result in the average power consumption not exceeding 234 mW. This value would likely be lower for an equivalent laser driver designed for the Bi-PON or XG-PON1 upstream rate of 2.5 Gb/s. Such improvement in the upstream OI energy efficiency further exposes protocol processing as the main source of power consumption in the 10 Gb/s PON and increases the importance of protocol optimizations.

C. Power in the Periodic Sleep Regime

The OI is the only ONT part which can potentially be more energy efficient for XG-PON than for Bi-PON. For Bi-PON, the average power consumption of this interface when sleep modes are used, is specified by

$$P_{oB} = (P_{o10G}) \cdot (1 - sleep_B), \quad (6)$$

and for XG-PON, it is specified by

$$P_{oX} = (P_{o10G}) \cdot (1 - sleep_X), \quad (7)$$

where the active power consumption P_{o10G} of both PON OIs is the same. The actual average power is proportional to the fraction of time the interface is turned on. The proportionality factor is expressed as $1 - sleep_B$ and $1 - sleep_X$, where $sleep_B$ and $sleep_X$ are the fractions of time spent in the sleep mode for Bi-PON and XG-PON ONT, respectively.

The use of cyclic sleep states is the only available energy saving technique in the XG-PON ONT and it applies to the complete ONT functionality, including the downstream PON PD-TIA-LA interface. In the following analysis, the potential energy efficiency of the downstream OI of a XG-PON ONT is compared with that of a Bi-PON ONT in which the downstream user rate is limited to only one 1.25 Gb/s lane. Such comparison is rather unfair given that a Bi-PON ONT can support multiple 1.25 Gb/s lanes, as explained in Section II-B, however it is intentionally selected to illustrate how unlikely the traffic conditions need to be in order for XG-PON sleep modes to outperform the Bi-PON power saving approach.

Assuming that the total Bi-PON user rate is limited to 1.25 Gb/s, which is $\frac{1}{8}$ of that of XG-PON, the maximum value of $sleep_X$ can theoretically be 8 times higher than the maximum value of $sleep_B$, resulting in the XG-PON ONT OI being 8 times more energy efficient than the same interface in the Bi-PON ONT. However, as explained below, a scenario in which the XG-PON OI consumes less energy than the Bi-PON OI is not realistic for any typical ONT traffic and QoS-aware traffic management since it is practically impossible to make XG-PON ONT sleep cycles longer than those of the Bi-PON ONT.

For both protocols, the awake time for the ONT running cyclic sleep states is a multiple of a whole frame period. Since the frame length for both PONs equals $125 \mu\text{s}$, the maximum payload per ONT per frame is equal $pload_X = 1,250,000$ bits for XG-PON and $pload_B = 156,000$ bits, for Bi-PON. Then, if for a given maximum allowed cyclic sleep period $(t_s)_{\text{max}}$, the sustained average bit rate of the downstream traffic $ds_{\text{avg}} \leq pload_B / (t_s)_{\text{max}}$, the awake times and OI power will be equal to one frame period for both ONTs. The necessary condition for the XG-PON OI operation to be more energy efficient than that of the Bi-PON ONT is $ds_{\text{avg}} > pload_B / (t_s)_{\text{max}}$. The values of the sustained downstream user rates for which $P_{oX} / P_{oB} = 1, 2, 4$ and 8 as function of the sleep cycle length, are shown in Fig. 12. For sleep periods of up to 10 ms, XG-PON ONT can achieve 8-fold savings in OI interface power only if its sustained user downstream traffic is greater than 100 Mb/s, whereas XG-PON advantage is completely eliminated for user traffic lower than

15 Mb/s. For a more realistic bit rate of 3 Mb/s, typical of a compressed high definition video stream, it would be necessary to extend the sleep period to 40 ms in order for the XG-PON OI to be more energy efficient, whereas to achieve 8-fold higher efficiency than the Bi-PON OI, its sleep periods would have to exceed 100 ms, which would significantly degrade QoS for most services. Additionally, the probability of sleep periods lasting 40-100 ms is very low in the conditions of the sustained traffic due to the ont wake-ups caused by arrival of upstream traffic [8].

To complete the comparison of XG-PON and Bi-PON energy consumption in cyclic sleep modes, the values of the total active power consumption for the two ONTs must be considered, including the power associated with digital processing and the user interface. As a reference for the total active power of a XG-PON1 ONT with 1 Gigabit Ethernet LAN port, we use the value of 6 W, specified in the European ‘‘Code of Conduct on Energy Consumption of Broadband Equipment’’ [35], whereas the estimated total power consumption for a Bi-PON ONU with the equivalent functionality is 1.3 W. Given that Bi-PON’s active power consumption is significantly lower, it is clear that XG-PON cannot compensate for it through a better use of the periodic sleep regime, because this would require the traffic volume to be very high, which in turn would make the aggressive use of sleep modes impossible.

In contrast with the XG-PON ONT in which most power is consumed in digital processing, the total power of the Bi-PON ONT is dominated by the power of its PON downstream optical interface, which does not scale with the user traffic but rather with the PON line rate. However, as explained in, Section IV Bi-PON protocol creates an opportunity for scaling down of the ONT line rate through the use of the cascaded bit-interleaved PON architecture. This architecture enables reduction of the ONT line rate to 1.25 Gb/s or 2.5 Gb/s and use of EPON or GPON OI components, which would reduce the total power of the PD, TIA and LA components from 650 mW to about 200 mW. Together with the clock recovery, the total power of the Bi-PON OI functionality would be reduced from 750 mW to 300 mW, resulting in the total ONT power consumption of only 750 mW.

In the above analysis, the power consumption of the electrical and optical components in the upstream path in both ONTs was considered to be responsible for about 15% of the ONT power and the power supply losses were assumed to be 20% of the total power consumption.

D. Experimental Results

In order to verify feasibility of the Bi-PON protocol and its energy efficiency, an ASIC shown in Fig. 13, including Bi-PON CDR and ONT downstream protocol processing logic, has been implemented in a 130 nm BiCMOS process, in collaboration with INTEC, Univ. of Ghent [30]. The ASIC includes the functionalities of a PLL-based 10 Gb/s CDR and the complete downstream Bi-PON protocol processing, while occupying a silicon area of only 2.5 mm^2 . The area of the die is dominated by analog circuitry, whereas the digital logic occupies only about 20% of the total area.

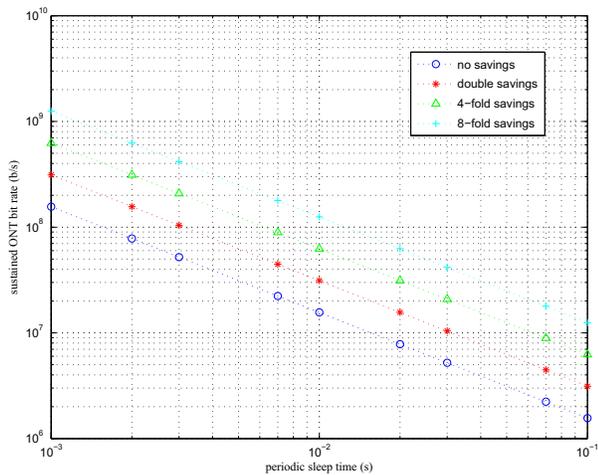


Fig. 12: Limits of XGPON potential optical power savings over Bi-PON.

The ASIC functionality has been verified in a complete 10 Gb/s PON setup, using bit-interleaved packet traffic at various user rates and measuring the power consumption of the ASIC.

Additionally, the Bi-PON ONT downstream protocol functionality has been implemented in Altera Stratix IV FPGA in order to compare its dynamic power consumption with that of the XG-PON ONT core implemented in the same FPGA and running on an identical ONT board. In this experiment, the static power consumption was subtracted from the total measured power because it was dominated by the FPGA internal power independent of our design.

Comparing the two FPGA-based implementations, the Bi-PON ONT design utilizes far less logic and memory resources than the XG-PON ONT design. This is mainly due to its simplified protocol and ability to greatly reduce the data rate of the incoming data stream to its useful content, very early in the architecture. This reduction in data rate eliminates the requirement of massive parallel data paths and processing. As a result, both dynamic and static power consumption of the Bi-PON ONT are lower than the XG-PON ONT.

The ASIC power consumption was measured separately for the analog CDR and the digital logic, for various Bi-PON bit rates. The total power of the ASIC is dominated by the analog part consuming about 130 mW under all operating conditions. The power of the protocol processing parts scales with both the traffic and the assigned payload rate, which is expected since Bi-PON digital clock frequency changes proportionally with the payload bit rate.

While passing traffic, it consumes between 50 mW at 9 Mb/s and 100 mW at 1.25 Gb/s. A similar power scaling trend is measured for the FPGA implementation of the Bi-PON ONT, with the power consumption in this implementation being about twice as high for the corresponding measurement points.

Unlike for the Bi-PON implementation, the power consumption of the implemented XG-PON downstream protocol did not show any noticeable change with the amount of traffic it was receiving, measuring a steady power of 3.7 W. Such

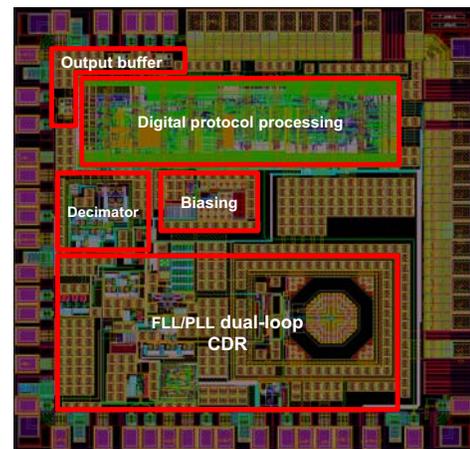


Fig. 13: Bi-PON ONT ASIC.

behavior was expected, since XG-PON protocol requires the ONT to process both the local and unrelated traffic. Plots of the measurement results obtained from the ASIC and both FPGA implementations are shown in Fig. 14.

Further, the Bi-PON upstream protocol, described in Section II has been implemented in the same FPGA and the complete functionality of the Bi-PON ONT prototype was verified with real, bidirectional traffic. The dynamic power of the upstream protocol function was measured for traffic ranging between 200 Mb/s and 1.25 Gb/s. As shown in Fig. 14, certain power scaling with the traffic exists, but it is far less pronounced than in the case of the downstream protocol processing. The reason for such behavior lies in the fact that the measurements include the power consumption of the UNI GigE receiver. Since all wired Ethernet links transmit their line code continuously, regardless of the presence of traffic, the GigE receiver in our Bi-PON upstream test setup was continuously consuming energy associated with clock and data recovery, word alignment and decoding of the line code. Therefore, the power plateau of approximately 220 mW, observed around the data points associated with lower traffic rates, represents the power of the GigE receiver. This implies that the actual power consumption of the Bi-PON upstream protocol implementation is approximately 35 mW, which is the difference between the maximum measured power consumption and the plateau value.

The result of the power measurements of the Bi-PON upstream protocol implementation further justifies the focus of this work on the optimization of the downstream protocol.

As argued in Section II, digital downstream processing is responsible for the dominant part of energy consumption in the standard 10G-EPON and XGPON ONTs, whereas Bi-PON protocol enables significant reduction of this component. The energy efficiency of the Bi-PON protocol processing has been confirmed by our experimental results and its impact on the energy efficiency of the whole ONT has been confirmed using available information about power consumption of other ONT parts, such as the physical interfaces, memory and the power supply.

The experimental results and analysis presented in this section unequivocally reveal the significant Bi-PON superi-

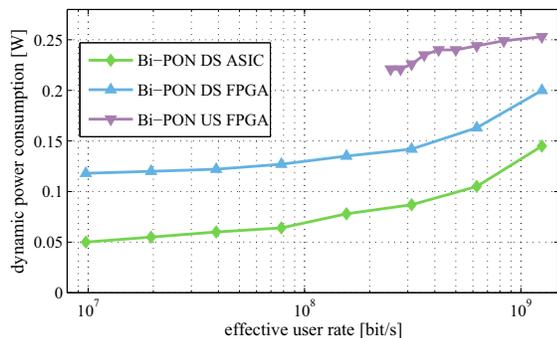


Fig. 14: User data rate vs dynamic power consumption for Bi-PON

ority in energy efficiency over XG-PON. Whereas Bi-PON outperforms XG-PON in the digital processing segment by more than an order of magnitude, this advantage is somewhat reduced due to the lack of power scalability in the PON optical interface. It has also been shown that aggressive use of sleep modes cannot change the relative efficiency in favor of XG-PON. The future progress in the design of low power optical interface components is expected to bring further advantage to Bi-PON over the conventional TDM PON protocols.

VI. CONCLUSION

The work described in this paper has demonstrated the impact of the protocol design on the network energy efficiency and pointed at the inefficiency of the standard TDM PON protocols, offering an alternative, low-energy protocol. Unlike any of the standard TDM PON protocol that require the CPE to perform switching of packets in the GEM (i.e. MAC) layer, the proposed bit-interleaving PON protocol, achieves such efficiency by performing switching in the PHY layer. This protocol not only enables major energy saving but also potentially simplifies and reduces the cost of the customer premises equipment. Moreover, due to its very low CPE power consumption, Bi-PON significantly extends the battery life and represents a much more dependable access technology than standard PON.

Bi-PON's dramatic reduction of energy consumption in the digital, protocol processing hardware, creates an opportunity for simplifying and optimizing some other optical networks, such as the converged metro-access, or access-home network. Being agnostic to higher-layer protocols, it also enables energy efficient sharing of the common PON infrastructure between multiple different networks that may use different data link layer protocols such as: Ethernet, CPRI, GPON, etc. Finally, the Bi-PON protocol is a promising solution for implementation of 40 Gb/s TDM PON, capable of keeping the CPE power consumption at an acceptable level.

Extensive implementation and prototyping work performed as part of this research, has not only provided the proof of concept for the proposed protocol, but also a very reliable insight into the Bi-PON energy consumption, based on physical power measurements.

In addition to the design of an energy saving PON protocol, this work has resulted in a few other novel solutions. One of these solutions is the technique for dynamic, bandwidth-adjustable, interleaving of different classes of traffic forwarded over a communication link. Feasibility of dynamic interleaving is key to enabling the Bi-PON as an alternative to standard TDM PON protocols because it provides flexibility of downstream bandwidth allocation equivalent to that of the standard protocols. However, the applicability of the interleaving technique is not limited to passive optical networks but can be extended to any point-to-multipoint link and possibly find use in point-to-point links as well.

Another key solution that emerged from this work is the method for schedule-specific, dynamically configurable descrambling of the decimated traffic, which enables scrambling of the interleaved Bi-PON traffic.

The third important solution is the ultra low energy decimating clock and data recovery circuit that completely eliminates the hardware and power overhead associated with deserializing and aligning of the received traffic.

Our future work will focus on higher bit rate, multi-stage, bit-interleaving architectures for converged networks.

REFERENCES

- [1] D. C. Kilper, G. Atkinson, S. K. Korotky, S. Goyal, P. Vetter, D. Suvakovic and O. Blume, "Power trends in communication networks," *IEEE J. of Selected Topics in Quantum Electronics*, vol. 17, no. 2, pp. 275–284, Mar./Apr. 2011.
- [2] A.R. Dhaini, P.H. Ho, and G. Shen, "Toward green next-generation passive optical networks," *IEEE Commun. Mag.*, vol. 49, no. 11, pp. 94–101, Nov. 2011.
- [3] *ITU-T Recommendation G.987.3*, 10 Gbps capable passive optical networks, TC layer specifications, Oct. 2010.
- [4] ABI Research, "Fiber broadband subscriptions to double by the end of 2016", www.abiresearch.com, July 2011.
- [5] C. Lange, et al., "On the energy consumption of FTTH access networks", in *Proc. OFC*, paper JThA79, 2009.
- [6] ITU-T G.832, *ITU-T Recommendation G.832*, Transport of SDH elements on PDH networks - Frame and multiplexing structures, Oct. 1998
- [7] L. Valcarenghi et al., "How to save energy in passive optical networks", in *Proc. ICTON*, 2011.
- [8] H. Bang, J. Kim, S. Lee and C. Park, "Determination of sleep period for cyclic sleep mode in XG-PON power management", *IEEE Commun. Letts*, vol.6, no. 1, Jan. 2012.
- [9] B. Skubic and D. Hood, "Evaluation of ONU power saving modes for gigabit-capable passive optical networks", *IEEE Network*, Mar./Apr. 2011.
- [10] B. Skubic, A. Lindstrom, E. I. de Betou and I. Pappa, "Energy saving strategies in fixed access networks" in *Proc. Green Communications, Theoretical Fundamentals, Algorithms and Architectures*, edited by J. Wu, S. Rangan and H. Zhang, CRC Press, Boca Raton, 2013.
- [11] A.R. Dhaini, P.-H. Ho, G. Shen, and B. Shihada, "Energy efficiency in TDMA-based next-generation passive optical access networks," *IEEE/ACM Trans. Netw.*, vol. PP, no. 99, pp. 1–14, May 2013.
- [12] N. P. Anthapadmanabhan, N. Dinh, A. Walid, and A.J. van Wijngaarden, "Analysis of a probing-based cyclic sleep mechanism for passive optical networks," in *Proc. IEEE Global Telecommun. Conf.*, Atlanta, GA, Dec. 2013.
- [13] IEEE Std 802.3av-2009.
- [14] R. Heidemann and H. Krimmel, "Cable television distribution network with video-on-demand transmission", U.S. patent 5,517,232, May 14, 1996.
- [15] "IEEE Standard for Information Technology–Telecommunications and Information Exchange Between Systems–Local and Metropolitan Area Networks–Specific Requirements Part 3: Carrier Sense Multiple Access With Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications - Section One," IEEE Std 802.3-2008 (Revision of IEEE Std 802.3-2005), Dec. 26 2008

- [16] N. P. Anthapadmanabhan, D. Suvakovic, H. K. Chow, and D. van Veen, "Power scalable descrambling methods for bit-interleaved TDM networks," in *Proc. ICC*, Budapest, Hungary, June 2013.
- [17] Common Public Radio Interface (CPRI) specifications, www.cpri.info
- [18] "IEEE Standard for Information technology–Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007)
- [19] H. Stone, "Parallel processing with the perfect shuffle," *IEEE Trans. Computers*, vol. 20, no. 2, 1971.
- [20] C. van Praet et al. "10 Gbit/s bit interleaving CDR for low-power PON," *Elec. Letts*, vol. 48, no. 21, pp. 1361–1363, Oct. 2012.
- [21] *ITU-T Recommendation G.989.1*, 40 Gigabit capable Passive Optical Network (NG-PON2) requirements, Mar. 2013.
- [22] D. Suvakovic, H.K. Chow et al. "Low energy bit-interleaving downstream protocol for passive optical networks", in *Proc. IEEE Online Conf. on Green Communications (GreenCom)*, Sep. 2012.
- [23] S. W. Golomb, *Shift Register Sequences*, Holden Day, 1967.
- [24] N. P. Anthapadmanabhan, D. Suvakovic, H. Chow and D. van Veen, "Power scalable descrambling methods for bit-interleaved TDM networks", in *Proc. IEEE Int. Conf. Commun.*, Budapest, Hungary, June 2013.
- [25] D. van Veen, V. Houtsma, P. Winzer and P. Vetter, "26-Gbps PON transmission over 40-km using duobinary detection with a low Cost 7-GHz APD-based receiver", in *Proc. ECOC*, Amsterdam, The Netherlands, p. Tu.3.B.1, Sep. 2012.
- [26] D. van Veen, D. Suvakovic, H. K. Chow, V. Houtsma, E. Harstead, P. Winzer and P. Vetter, "Options for TDM PON beyond 10G", access networks and in-house communications," in *Proc. ANIC*, Denver, p. AW2A.1, June 2012.
- [27] R.P. Davey, P. Healey, I. Hope, P. Watkinson, D.B. Payne, O.Marmur, J. Ruhmann and Y. Zuiderveld, "DWDM reach extension of a GPON to 135 km", *J. Lightwave Techn.*, vol. 24 , no. 1, pp. 29–31, 2006.
- [28] T. Ayhan, A. R. Dhaini, L. G. Kazovsky, D. Suvakovic and H. K. Chow, "Energy-efficient cascaded bit-interleaving protocol for integrated optical access/in-building networks", *IEEE GreenCom*, 2013.
- [29] J.-P. Gelas, L. Lefèvre, T. Assefa and M. Libsie, "Virtualizing home gateways for large scale energy reduction in wire-line networks", in *Proc. Electronic Goes Green Conf.*, Berlin, Germany, Sep. 2012.
- [30] C. van Praet, et al., "Demonstration of low-power bit-interleaving TDM PON," *Optics Express*, vol. 20, no 26, Dec. 2012.
- [31] F. Tachibana et al., "A 27% active and 85% standby power reduction in dual-power-supply SRAM BL power calculator and digitally controllable retention circuit," in *Proc. Int'l Solid-State Circuits Conference*, 2013.
- [32] S.-W. wong et al. "Sleep mode for energy saving PONs: advantages and drawbacks", in *Proc. IEEE GLOBECOM*, 2009.
- [33] L. Valcarengi et al. "Energy efficiency in passive optical networks: where, when and how?", *IEEE Network*, Nov./Dec. 2012.
- [34] A.P. Chandrakasan, S. Sheng and R.W. Brodersen, "Low-power CMOS digital design", *IEEE J. Solid-State Circuits*, vol. 27 , no. 4, pp. 473-484, 1992.
- [35] *European Commission Joint Research Centre*, Code of Conduct on Energy Consumption of Broadband Equipment, version 4.1, Jan. 2013.
- [36] H. Kozumi, M. Togashi, M. Nogawa and Y. Ohtomo, "A 10Gb/s burst-mode laser diode driver for burst-by-burst power saving," in *Proc. of IEEE ISSCC*, 2012