Demonstration of Industrial Network Applications by PHY Softwarization for Fully Virtualized Access Networks

Takahiro Suzuki, Yushi Koyasako, Sang-Yuep Kim, Jun-ichi Kani and Tomoaki Yoshida

NTT Access Network Service Systems Laboratories, NTT Corporation, 1-1 Hikari-no-oka, Yokosuka-Shi, Kanagawa, 239-0847 Japan

takahiro.suzuki.pc@hco.ntt.co.jp

Abstract: This paper demonstrates that PHY processing for industrial network applications and PON can be realized on a common general-purpose server for flexible service creations; it realizes the level of latency performance required by IRT communication. © 2022 The Author(s)

1. Introduction

Traditional optical access systems are changing to the architecture based on general-purpose components and software in order to break away from the development inefficiency of dedicated hardware with its long development cycles and high initial costs. The SDN-enabled broadband access (SEBA) promoted by telecom operators and vendors architects central offices (COs) with whitebox optical line terminals (OLTs) and general-purpose servers. It enables flexible service creation with softwarization of upper-layer functions such as authentication workflow and edge applications [1]. For flexible lower-layer service creation, further softwarization of OLT hardware has been studied. To realize low-latency networks for remote control of machinery, the softwarized low latency dynamic bandwidth allocation (DBA) mechanism has been proposed [2]. In the physical layer (PHY), software implementation of 10G-EPON physical coding sublayer (PCS) has been already demonstrated using a general-purpose server with graphics processing units (GPUs) [3, 4]. However, the demonstrations to date have been limited to PON applications. Traditionally, application-specific networks have been provided in only local networks using dedicated systems as shown in Fig. 1(a). By implementing specific networks in the access segments, edge computing services that extend the specific network to the access segments can be provided as shown in Fig. 1(b). They provide service providers with lower initial costs for new service creations and simplification of system management. An example of application-specific communication is the industrial protocol having unique PHY requirements. While PHY processing of communications based on Ethernet and PON is mainly composed of line coding and forward error correction (FEC), isochronous real-time (IRT) communication for industrial networks, such as Time-Sensitive Networking (TSN), PROFINET, and EtherCAT, include a cyclic transfer function of Ethernet frames for deterministic networks as shown in Fig. 1(c), which have requirements of 1 ms or less latency [6, 7].

In this paper, we demonstrate that industrial PON that combines PON PHY with IRT communication as a new transmission service can be realized on a general-purpose server for flexible service creations by softwarization of time slot allocation and hardware-assisted cyclic transfer. Our software implementation successfully allocates bandwidth for frames with IRT requirements and satisfies the level of latency performance required by IRT communication.

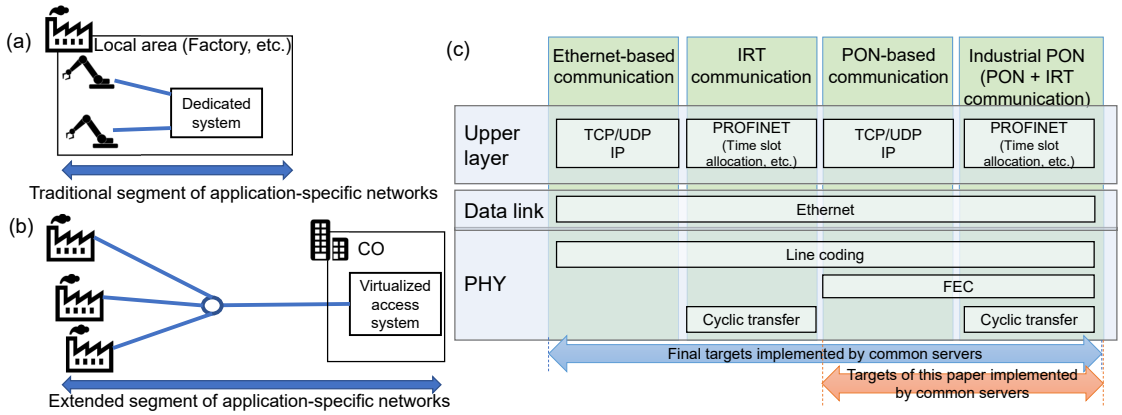

Fig. 1. (a) Limited segment where application-specific networks are provided, (b) extended segment where application-specific networks are provided, and (c) differences in PHY specification for each communication type.

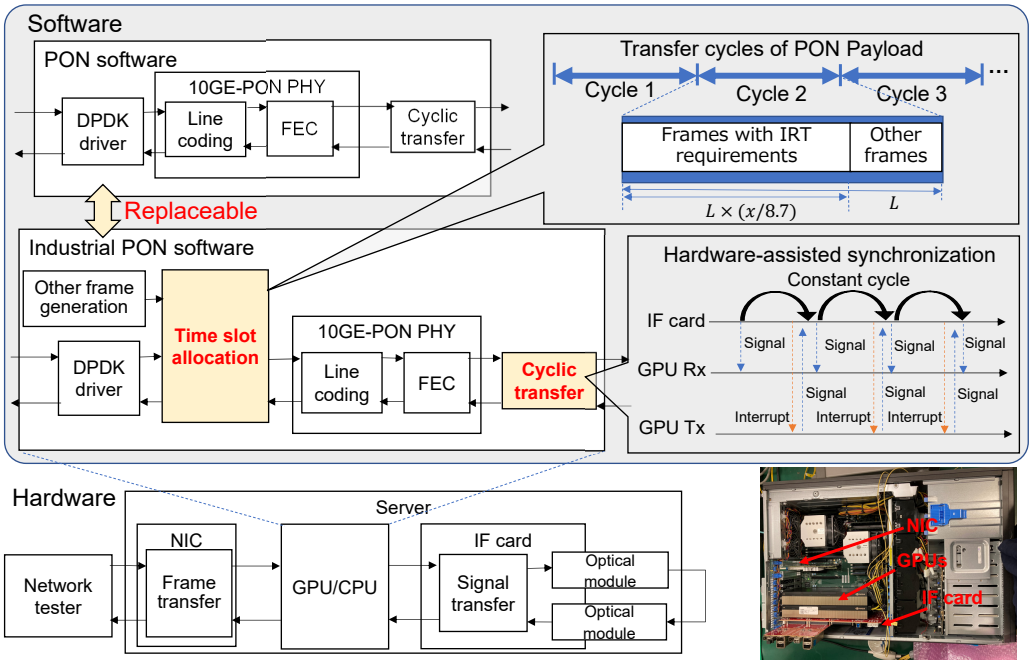

Fig. 2. Proposed PHY implementation of PON and industrial PON on a common server.

2. Industrial PON softwarization on a common general-purpose server

Figure 2(a) depicts our demonstration platform that softwarizes the functions of IRT communications combined with 10G-EPON PHY on a general-purpose server. Our platform is based on the implementation of 10G-EPON PHY softwarization; latency is measured by a network tester [5]. The novelties of this paper are that we implement time slot allocation and utilization of cyclic transfer based on hardware-assisted synchronization to realize primary functions of industrial protocols. In the configuration, the network tester generates and receives Ethernet frames to measure the throughput, latency, and jitter of the softwarized PHY. The network interface card (NIC) passes the Ethernet frames between the network tester and GPU/CPU. CPU sends and receives the Ethernet frames using the Data Plane Development Kit (DPDK) driver, which is open-source software to realize high-speed packet processing. The time slot allocation allocates dedicated time slot to the frames with high real-time requirements. 10G-EPON PHY plays the role in encoding the input frames and decoding 10G-EPON frames with line coding and forward error correction (FEC). The cyclic transfer includes polling for 10G-EPON frame reception from interface (IF) card and an interrupt-based signal output function to IF cards. IF cards are responsible for transferring the signal of 10G-EPON frames without any processing and the frames are looped back from optical module. Since functions of PON and industrial PON are implemented entirely in software, they are replaceable by changing the startup program.

For time slot allocation, we refer to the PROFINET protocol [7] and configure the time slots using PON frames. PROFINET features IRT communications by defining bus cycles having exact start times and allocating constant time to the IRT channel in the bus cycle. In our frame composition, given that the PON frame is periodically transferred under transfer control, we simulate it by allocating a constant length to the frames with IRT requirements. The other frames without IRT requirements are allocated to the remaining time slots, which are generally utilized for management data communication of industrial devices. In this experiment, frames generated by a network tester are treated as frames with IRT requirements, and the other frames are generated in the server. When allocating an allocated throughput of x Gbps for frames with IRT requirements, the length of $L \times (x/8.7)$ is assigned for PON payload length L, where 8.7 Gb/s is the theoretical throughput of 10G-EPON. A ratio of (x/8.7) of the PON payload is assigned to frames with IRT requirements, and a length of ratio 1-(x/8.7) is assigned to other frames.

For cyclic transfer, given that IRT communication sets precise hardware-level jitter requirements, we utilize hardware-assisted synchronization for industrial network applications, which is also implemented in 10G-EPON PHY softwarization in order to mitigate fluctuations in software processing. IF card is driven by a hardware clock to periodically send and receive PHY-encoded signals and interrupt signals. It also adds a data update flag to the PHY-encoded signals, and the receiver side on the GPU polls the flag to receive the signal and start PHY decoding. The transmitter side on the GPU receives interrupt signals to notice the start timing of data transfer, and PHY-encoded signals are sent to the IF card. 10G-EPON PHY processing time must be much shorter than the constant cycle of the cyclic transfer.

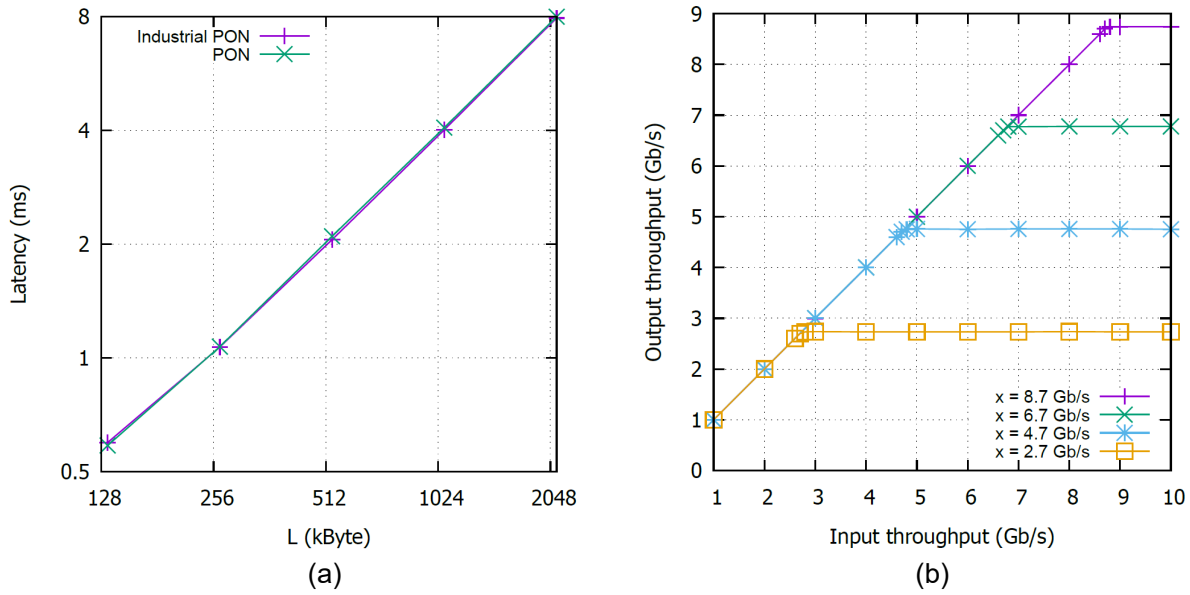

Fig. 3. (a) latency throughput and (b) throughput performance between the our industrial PON and PON on the same platform.

3. Performance evaluation

We evaluated the softwarized industrial PON by the setup shown in Fig. 2. The server used Intel Xeon E5-2699v4 CPUs. The two GPUs, NVIDIA Tesla A100, were connected to the server via PCIe gen3 x16. As the IF card, we utilized Tokyo Electron Device TB-7VX-690T-PCIEXP; it was connected to the server via PCIe gen3 x8. The small form-factor pluggable (SFP)+ optical module connected to the IF board generated and received 10.3125-Gb/s signals. A standard SFP+ module (Sumitomo Electric Industries SPP5100ZX-GL) with a built-in avalanche photodiode (APD) was used. Idle signals were inserted between PON frames. The latency and throughput were measured by a network tester (VIAVI MTS 5800).

Figure 3(a) shows the measured latency for each payload length L. The latency decreases as L becomes shorter, but it was measured within the range where frame transfer is possible. The result shows that our implementation of industrial PON achieves latency of 0.596 ms comparable to PON PHY when L is 133 kByte. This value is close to the 1 ms or less latency target of IRT communication. Figure 3(b) shows the measured output throughput for each input throughput of simulated frames with IRT requirements in the network tester. Input throughput is the throughput of frames sent by the network tester to the server's NIC. The throughput was measured for allocated throughput $x = 8.7, 6.7, 4.7, 2.7$ Gb/s in our industrial PON. For example, when x is set to 6.7 Gb/s, 8.7-6.7=2.0 Gb/s is assigned to other frames and the output throughput saturates to 6.7 Gb/s. The results show that the bandwidth of the frames with IRT requirements are correctly allocated for a throughput of 8.7 Gb/s, and the time slot allocation is successfully implemented. These results show that our implementation successfully demonstrated the performance attributes desired for industrial networks and the feasibility of applying our PHY softwarization approach to various transmission applications.

4. Conclusion

This paper demonstrated an industrial network application by softwarizing industrial PON as a new service, which enables time slot allocation and hardware-assisted synchronization for cyclic frame transfer. The real-time experiment showed that our implementation of industrial PON achieved latency of 0.596 ms, and the feasibility of prototyping new transmission services by softwarizing PHY was demonstrated.

References

1. S. Das, "From CORD to SDN Enabled Broadband Access (SEBA) [Invited Tutorial]," in Journal of Optical Communications and Networking, vol. 13, no. 1, pp. A88-A99, Jan. 2021.
2. D. R. Mafioletti et al., "Demonstration of a low latency bandwidth allocation mechanism for mission critical applications in virtual PONs with P4 programmable hardware," 2022 Optical Fiber Communications Conference and Exhibition (OFC), pp. 1-3, 2022.
3. T. Suzuki, et al., "Demonstration of Fully Softwarized 10G-EPON PHY Processing on A General-Purpose Server for Flexible Access Systems", Journal of Lightwave Technology, vol. 38, Issue 4, pp. 777-783, Feb. 2020.
4. T. Suzuki, et al., "Software Implementation of 10G-EPON Upstream Physical-Layer Processing for Flexible Access Systems," in Journal of Lightwave Technology, vol. 37, no. 6, pp. 1631-1637, Mar. 2019.
5. T. Suzuki, et al., "Low-Latency PON PHY Implementation on GPUs for Fully Software-Defined Access Networks," IEEE Network, vol. 36, no. 2, pp. 108-114, Mar./Apr. 2022
6. IEEE Std 1588-2019: "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," 2019.
7. "PROFINET System Description", PROFIBUS Nutzerorganisation, Oct. 2014.