Demonstration of End-to-End Interoperability between Legacy 100MbE and a 40Gb/s Optical Label Switched Network Layer

John M. Garcia, Henrik N. Poulsen, Daniel J. Blumenthal
Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106
johngarcia@ece.ucsb.edu

Abstract: Demonstration of burst mode assembly, detection, and recovery of 100MbE frames as 40Gb/s optical packets enabling scalable interoperability with legacy networks. End-to-end adaptation with $<10^{-4}$ packet loss over an 18dB range of received power is demonstrated.
OCIS codes: (060.4250) Networks; (060.4259) Networks, packet-switched; (060.6719) Switching, packet

1. Introduction

Optical packet switching (OPS) and optical label switching (OLS) are potential approaches to scale future packet networks in terms of power and footprint while meeting the growing bandwidth demands of today's networks [1]. One proposed approach to OLS is to transmit a high bit rate payload preceded by a relatively low bit rate optical header containing control and routing information [2]. This allows the optical header to be processed by low-speed electronics while the payload is processed entirely in the optical domain on a per-packet-basis. As deployment of OLS core routers becomes more feasible, OLS edge routers will be required to achieve interoperability with legacy network architectures.

Previously, an edge router for a multi-hop OLS network was demonstrated with a labeled packet format that was limited to low bit rates (155Mb/s optical header and 2.488Gb/s payload) [3]. End-to-End asynchronous adaptation of Internet protocol (IP) packets to an optically labeled format has been shown with packet forwarding and header rewrite [4], except it was limited to 3.125Gb/s headers and 12.5Gb/s payloads. In [5], buffering and routing was performed on 10GbE frames that are converted to and from an optically labeled format that consists of an 80Gb/s (8λ at 10Gb/s) payload with a 10Gb/s header. However, the NxN switch fabric requires N^2 active switches due to the wavelength division multiplexed (WDM) format of the optical packets.

In this paper we present a scalable end-to-end adaptation between 100Mbit Ethernet and 40Gb/s optically labeled packets that can be optically buffered, forwarded, and routed by compact monolithic photonic integrated circuits (PICs) [6]. Dynamic generation of labeled headers at the ingress layer is performed via an electronic IP/label table lookup. At the egress, real-time assembly of 40Gb/s payloads is performed by means of burst mode electronic payload identifier recovery in a manner similar to [4]. Successful packet adaptation is achieved with a packet loss rate less 10^{-4} measured by a commercial Ethernet packet tester. The main focus of this paper is on the interface between the optics and electronics that are relevant to an OPL edge router. As such, no traffic shaping or aggregation is performed.

2. Optically Labeled Packet Format and Adaptation Layers

Packets are arriving at the ingress layer are converted to the optical packet format shown in Fig. 1(a). It consists of a 10Gb/s non-return-to-zero (NRZ) optical header (OH) that is processed electronically by a core optical data router (ODR) and a 40Gb/s return-to-zero (RZ) optical payload (OPL) that is processed at low switching speeds in the optical domain. The OH is made up of a 64-bit OH identifier (0x89ABCDEFFEDCBA98), used by ODRs for OH detection and recovery, and a 10-bit label used to for routing purposes. Ethernet frames are encapsulated by a 64-bit payload (PL) identifier (0x98DFECA88ACEFDB9) to form a 76-byte OPL. The OH and OPL are separated by a guard band (GB) of 12.8ns. In the absence of Ethernet frames, 10Gb/s idlers are generated to maintain an adequate DC balance. There is an additional guard band of 38.4ns between optical packets and idlers.

A schematic representation of the ingress and egress adaptation layers is illustrated in Fig. 1(b). The ingress FPGA generates an OH based on the destination address located in the IP header of the asynchronously arriving Ethernet frames. The frames converted to the optical packet format are internally serialized to four 16-bit 622Mb/s output channels that are multiplexed to 40Gb/s by the external serialization stages. The OH is repeated at each of the four output channels to ensure that their resulting serialization 10Gb/s. The OPL is evenly distributed between the four output channels to guarantee that its resulting serialization is identical to the original Ethernet frame, except compressed in time: bits (0, 4, 8 …) are sent through channel 1, bits (1, 5, 9 …) through channel 2, bits (2, 6, 10 …) through channel 3, and bits (3, 7, 11 …) through channel 4.

978-1-55752-932-9/11/$26.00 ©Optical Society of America

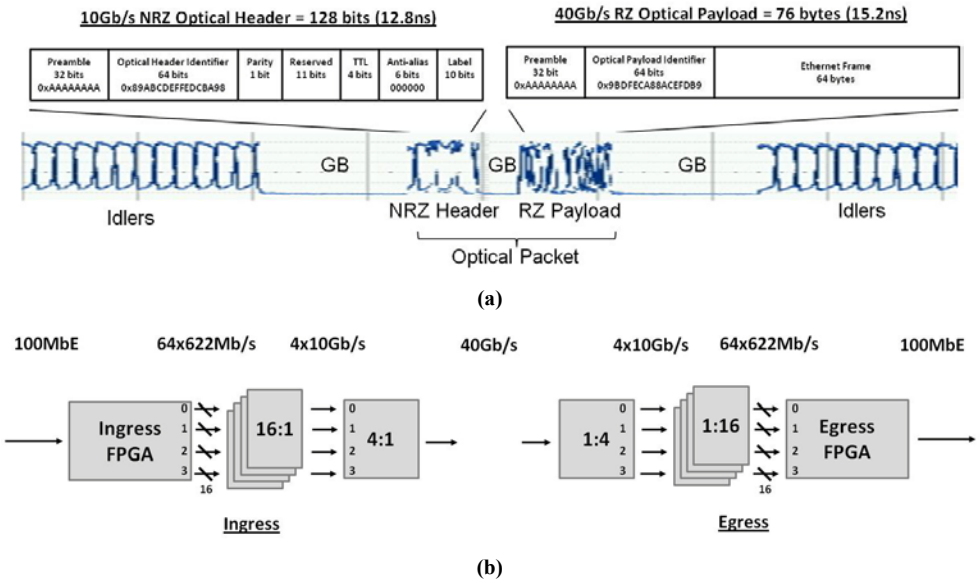

Fig. 1. (a) Optically labeled packet format and (b) (de)serialization stages used to achieve adaptation between 100MbE frames and 40Gb/s optically labeled packets.

At the egress layer of adaptation, the 40Gb/s optically labeled packets are sent through the deserialization stages depicted in the right part of Fig. 1(b) and are demultiplexed down to four 16-bit 622Mb/s channels that are forwarded to the egress FPGA. Ideally, data out of the 1:4 deserializer would be evenly distributed at the 1:16 deserializers with bit 0 coming out of port 0. In fact, data from the 1:4 deserializer will be shifted out of sequence due to the fact that optically labeled packets are arriving asynchronous to the time slot of the deserializer. Even though this effect is only shown for the 1:4 deserializer, it is also present in the each of the 1:16 deserializers.

The egress FPGA utilizes the 64-bit PL identifier to circumvent the shifted sequencing issue. Detection and recovery of the PL identifier is performed by scanning each 622Mb/s input channel for one of four 16-bit identifiers (0xF497, 0x0FF0, 0x33CC, and 0x55AA) that correspond to the demultiplexed values of the 64-bit PL identifier. Correct sequencing is determined once all four 16-bit identifiers are successfully detected within a clock cycle of the FPGA clock domain (155MHz). Occasionally, the 16-bit identifiers may have different arrival times that span more than one FPGA clock cycle. A synchronization stage based on configurable D flip-flop (DFF) delays is used to synchronize the four channels. Once correct sequencing is determined, the Ethernet payload is written into an asynchronous memory block to allow the payload to transition into the 100MbE clock domain. Prior to that, a checksum is performed on the IP header to ensure that it has not been corrupted. Payloads with verified IP header checksums are committed to memory and are transmitted over a 100BASE-TX interface.

3. Experimental Setup and Results

The setup in Fig. 2 is used to evaluate the performance of the adaptation layers. A stream of 50 million 64-byte Ethernet frames with 0.96μs of inter-packet gap (IPG) is generated by a commercial packet tester (SmartBits 6000). The 100MbE frames are converted to 40Gb/s optically labeled signals that are used to drive an optical transmitter comprised of a tunable laser (TL) followed by two Mach-Zehnder modulators (MZM) connected in tandem and an Erbium doped fiber amplifier (EDFA). The alternating NRZ/RZ signals are generated by a 20Gb/s 2:1 multiplexer with an RZ-enable signal and a 50Ω-to-ground termination as inputs. This creates a burst of 20GHz clock tones ($V_{pp} = 2V_{\pi}$) that are time-aligned with the 40Gb/s payloads in order to carve RZ pulses. The optical packets are transmitted at an average optical power of -3dBm to a pre-amplified optical receiver directly followed by the egress adaptation layer. A variable optical attenuator (VOA) is placed before the pre-amplified receiver to vary the optical signal to noise ratio (OSNR) of the transmitted signals. The egress layer converts 40Gb/s optically labeled packets into Ethernet frames and sends them to the commercial packet tester where Layer-3 packet recovery measurements are performed through 32-bit cyclic redundancy check (CRC) computations. Real-time Layer-2 statistics are provided by the egress FPGA to a graphical user interface (GUI) via a UART-to-USB bridge. Oscilloscope traces of the 40Gb/s labeled packets at the 16:1 serializers (left), the 4:1 serializer (top-right), and the optical transmitter can be seen in Fig. 2(b).

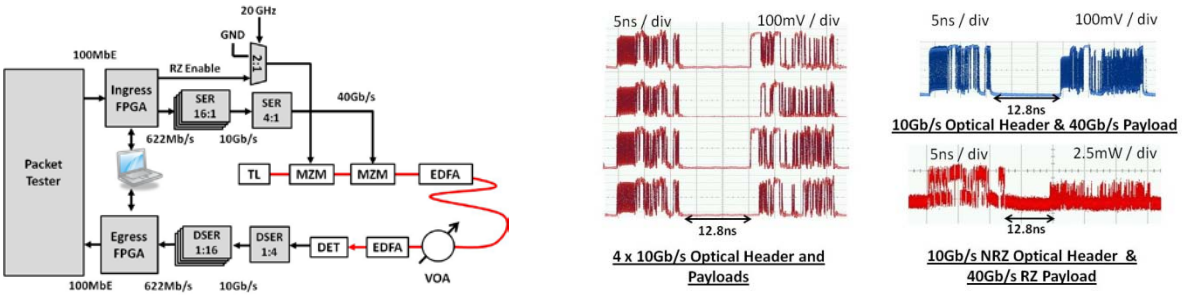

(a) **(b)**

Fig.2. (a) Experimental setup (Packet length = 64 bytes, stream length = 50M packets, IPG = 0.96µs, λ = 1550nm, P_{TX} = -3dBm) and (b) oscilloscope traces of optically labeled packets at different ingress stages of serialization.

Fig. 3(a) shows the Layer-2 and Layer-3 measurement results of recovered PL identifiers and recovered Ethernet packets respectively. Greater than 99.99% Ethernet packet recovery over a dynamic receiver input power range of 18dB is achieved. There is a negative power penalty associated with the PL identifier recovery that is attributed to the fact that one is less likely to encounter an error when attempting to detect a 64-bit identifier as opposed to computing the 32-bit CRC for an entire Ethernet frame. The egress loss rate measurements are plotted on a logarithmic scale in Fig. 3(b). A packet loss rate less than 10^{-4} and 10^{-5} is demonstrated over a range of 18dB and 4dB of received power respectively. One will also note that the egress performance is dominated by the imperfect deserialization of channel 4.

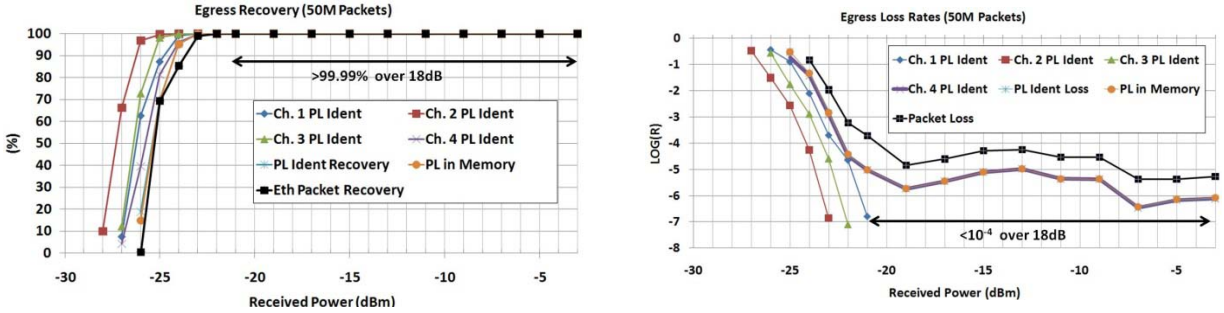

(a) **(b)**

Fig.3. (a) Experimental results showing egress recovery and (b) loss ratios.

4. Conclusion

We demonstrate real-time generation of optically labeled headers, deserialization and assembly of asynchronously arriving 40Gb/s payloads using custom burst mode adaptation layers. Successful end-to-end interoperability between legacy 100MbE and a 40Gb/s optically labeled packet format is shown with less than 10^{-4} packet loss rate over an 18dB range of receiver input power using a commercially available packet tester.

5. Acknowledgement

This work is supported by DARPA/MTO and ARL under LASOR award #W911NF-04-9-0001.

6. References

[1] D. J. Blumenthal, "Routing packets with light," *Sci. Amer.*, vol. 284, no. 1, pp. 80–83, Jan. 2001.
[2] D. J. Blumenthal, P. R. Prucnal, and J. R. Sauer, "Photonic packet switches: Architectures and experimental implementations," *Proc. IEEE*, vol. 82, no. 11, pp. 1650-1667, Nov. 1994.
[3] S. J. B. Yoo *et al,* "High performance optical-label switching packet routers and smart edge routers for the next-generation Internet," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 7, pp. 1041–1051, Sep. 2003.
[4] Reza Nejabati *et al,* "The 'OPORON' Project: Demonstration of a Fully Functional End-to-End Asynchronous Optical Packet-Switched Network," *J. Lightwave Technol.*, vol. 25, no. 11, pp. 3495-3510, Nov. 2007.
[5] H. Furukawa *et al*, "Demonstration of 10 Gbit Ethernet/Optical-Packet Converter for IP Over Optical Packet Switching Network," *J. Lightwave Technol.*, vol. 27, no. 13, pp. 2379-2390, July 2009.
[6] J. P. Mack *et al*, "Asynchronous 2x2 Optical Packet Synchronization, Buffering, and Forwarding," in *Optical Fiber Communication Conference*, OSA Technical Digest (CD) (Optical Society of America, 2010), paper OThN1.