All-optical asynchronous variable-length optically labelled 40Gbps packet switch

D. Wolfson (1), V. Lal (1), M. Masanovic, H. N. Poulsen (1), C. Coldren (2), G. Epps (3), D. Civello (3), P. Donner (3), and D. J. Blumenthal (1)
1: University of California, Santa Barbara, ECE Department, USA, email: dwolfson@ece.ucsb.edu
2: Agility Communications, 3: Cisco Systems

Abstract *We describe the operation of a packet switch that dynamically forwards 40Gbps optical packets based on asynchronously recovered 10Gbps optical headers using highly-integrated InP photonic circuits. Results of both Layer 1 and Layer 2 performance measurements are presented.*

Introduction

Key challenges in scaling high-capacity commercial electrical routers are the increased power and space requirements. All-optical packet switching is a very promising technique to address both issues, i.e., power consumption can be reduced because only low speed optical headers need to be processed by electronics and expected advances in monolithic integration can substantially decrease the footprint of highly advanced optical processing chips and thereby the footprint of optical packet switches.

In this paper, we report for the first time on an all-optical packet switch developed under the LASOR project with the following key features: i) Fully asynchronous operation with variable length packets representative of Internet traffic mixture (IMIX), ii) first time Layer 1 *and* Layer 2 performance measurements of a functioning 40Gbps optical packet switch, iii) 40Gbps packet forwarding and 10Gbps header rewrite using a monolithic InP packet forwarding chip (PFC).

Optical Packet Switch – principle of operation

The optical packet switch is shown in Figure 1. Input optical packets consist of 40Gbps RZ payloads and 10Gbps NRZ headers. The optical headers contain various information fields in addition to an Optical Header Identifier (OHI) that uniquely identifies the header and a label field indicating the egress port (see right side of Figure 1 for details).

The packet stream is optically tapped and the optical headers are asynchronously recovered in the 10Gbps burst mode clock and data recovery (CDR). After optical header recovery and processing in the Electronic Route Processor (ERP), the ERP sends an erase signal to a blanking SOA to erase the optical header. The ERP also computes a new packet λ and a new optical header. The fixed optical delay (~60m fiber = ~300ns) takes the electrical processing time into account. These signals are sent to the Packet Forwarding Chip (PFC), which is an InP-based Mach Zehnder Interferometric wavelength converter with an on-chip fast tuneable CW laser and header re-write modulator and is described in more detail in [1]. An Arrayed Waveguide Grating Router (AWGR) is used in combination with the PFC to forward the optical packets to the correct egress port on a per-packet basis.

Details of the setup

The optical packet transmitter generates variable length 40Gbps RZ payloads (40→1500) bytes with 10Gbps NRZ headers (128 bits). We use a repeating 60 packet long packet stream approximating Internet traffic (IMIX) with a 7:4:1 packet payload size ratio, i.e., 7x40 byte payloads (8ns), 4x570 byte payloads (114ns), and 1x1500 byte long payloads (300ns). Figure 1 shows a screen shot with the 3 different payload sizes. The label fields in the optical headers are coded in such a way that every other packet is forwarded to egress port 1 and 2, respectively. It should be stressed that using a packet stream with variable packet lengths demonstrates the system's ability to control and operate the packet switching components under a realistic scenario. A total guard band of 100ns is used per packet with 30ns before the header, 30ns between the header and the payload and 40ns after the payload. This is to take into account the optical and electrical signal processing times as well as timing uncertainty during header rewrite. The focus of this experiment was not to minimize the guard bands and the current values are not representative of a lower limit

The 10Gbps burst mode asynchronous CDR recovers the clock and optical headers within 20 bits. The CDR is based on a narrow band filter and a feedback loop in the clock path and a D flip-flop in the data path. The optical packet stream is also processed by a Payload Envelope Detect (PED) circuit that generates a logic envelope signal corresponding to the payload location and length [2]. The PED enables the ERP to maintain correct timing of all controls signals to the optical forwarding plane. The PED circuit is based on fast payload clock recovery followed by a Schottky diode. The erase stage is an SOA with turn on/off times of ~2ns and ratios in excess of 40dB. Figure 2 shows input/output signals from the erase stage in path B.

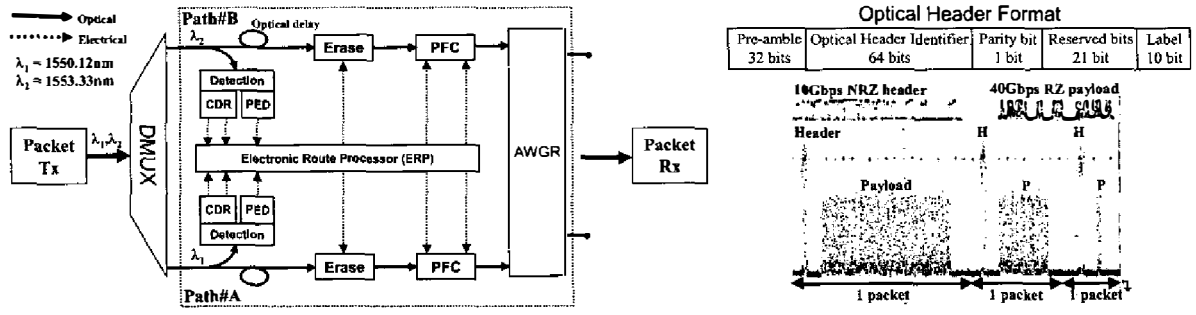

Figure 1: System diagram of the optical packet switch (left), the optical header format (right, top) and the generated Internet based packet stream (right bottom) – H: 10Gbps NRZ header, P: 40Gbps RZ payload.

Figure 2: Optical input and control signal to the erase stage in path B (top) and output signal (bottom).

Layer 1 and Layer 2 measurements

We define Layer 1 measurements as the packet BER of optical headers and payloads. Layer 1 measurements are used in this experiment to verify that the PFC is working error free for payload wavelength conversion and for header rewrite. We define Layer 2 measurements as capturing the optical packet stream and calculating how many headers or payloads are captured. The optical headers and payloads are uniquely identified by a 64 bit field as shown in Figure 1 for optical headers. We use a SHF 50Gbps BERT (10000 series) with 128 Mbit of internal RAM to store captured data and proprietary software to transfer the data to a PC and to analyze it.

Figure 3 shows Layer 2 performance of the CDR and PED in both paths of the optical packet switch. The right part of the figure shows the CDR output (top) and PED signal (bottom), respectively. The PED circuits have a large input power dynamic range (IPDR) where 100% header recovery is possible. The spike at lower input powers is due to noise that is interpreted as a valid PED. The IPDR of the CDRs is limited to <3dB due to the D flip-flop used in our design. CDR#B is able to recover 100% of optical headers, whereas CDR#A does not exceed ~90%. This is due to reflections in the clock path on the PCB board of CDR#A.

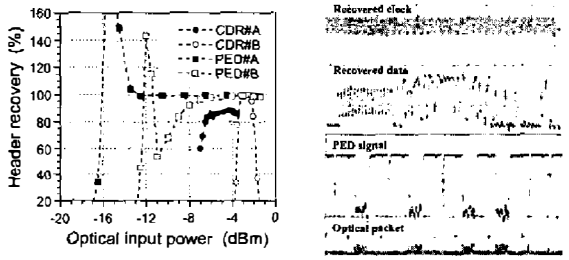

Figure 3: Header recovery in the PED and CDR in both paths of the optical packet switch (left), recovered clock and data from the CDR (right, top) and PED signal (right, bottom).

The measured packet BER of the PFC is shown in Figure 4 for both 40Gbps RZ payload conversion (left) 10Gbps NRZ header rewrite (right). For the payload BER measurements, header re-write was not performed and vice versa for the header BER. 40Gbps payload conversion resulted in power penalties ranging from 4-7dB. This penalty is mostly caused by extinction ratio degradation as can be seen in the inset. Error free performance with no additional penalty between PRBS data and packets was observed. The results for 10Gbps header re-write also shows error free performance. The large difference in sensitivity is due to the optically pre-amplified receiver.

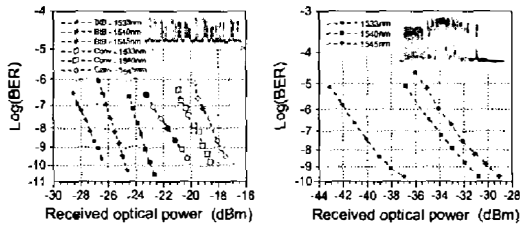

Figure 4: BER curves for 40Gbps payload conversion (left) and 10Gbps header rewrite (right) using the PFC in path B.

Figure 5 shows the variable length optical packet stream after payload conversion and header rewrite but without λ-switching (top) and with per-packet switching at egress port 1 (middle) and egress port 2 (bottom). Every other packet in the incoming packet stream is correctly forwarded to egress port 1 and 2, respectively. The pedestals after λ–switching is caused by the fact that when packets are forwarded to, e.g., egress port 2, no light is coming out of egress port 1. To obtain these results, the PFC has successfully performed the following tasks *simultaneously*: i) wavelength conversion of 40Gbps RZ payloads, ii) re-write of NRZ 10Gbps headers and iii) switched the wavelength on a per-packet basis.

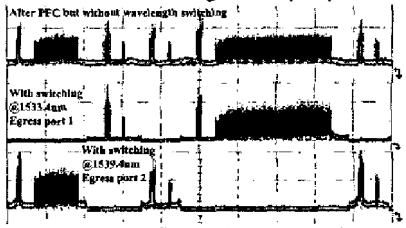

Figure 5: Pulse traces after payload conversion and header rewrite but without λ-switching (top) and after switching alternating packets to separate egress ports.

Layer-2 measurements are shown in Figure 6. Recovered headers and packets after the switch are measured as a function of the input power to the packet receiver. Back-to-back measurements without the packet switch and with the packet switch in static mode are included in the graphs. A penalty is incurred when the optical packets are forwarded dynamically. Operation of the switch with the receiver at the proper operating point allows more than 90% of the packets to be recovered after switching.

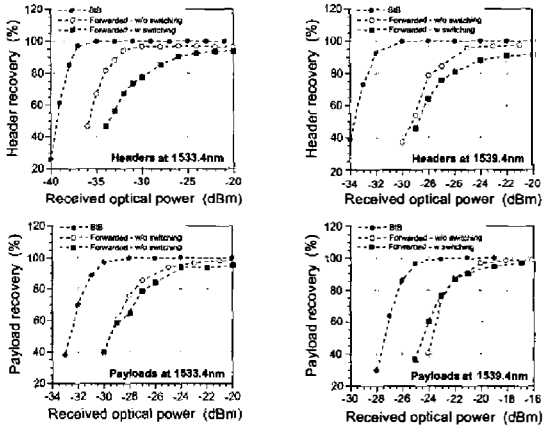

Figure 6: Optical header and payload recovery for dynamically switched optical packets.

Conclusions

In this paper we have demonstrated a 40Gbps asynchronous optical packet switch. We have shown dynamic forwarding of variable-length packets and evaluated the packet switch with respect to Layer 1 and Layer 2 performance metrics. Header recovery loss and packet switch throughput are characterized with over 90% throughput at Layer-2 and error free operation in Layer-1. We believe this is the first time Layer-2 measurements have been reported for a 40Gbps optical packet switch. The switch employs an integrated InP packet forwarding chip that performed 40Gbps wavelength conversion of RZ payloads, re-write of 10Gbps NRZ headers and λ-switching on a per-packet basis.

Acknowledgements

This work is supported by LASOR award #W911NF-04-9-0001 under the DARPA/MTO DoD-N program.

References

1 V. Lal et al, Submitted to ECOC PDP, 2005.
2 Z. Hu et al., PTL, Volume 17 (2005), pp. 1537-1539.