End-to-End Asynchronous Optical Packet Transmission, Scheduling, and Buffering

John P. Mack, John M. Garcia, Henrik N. Poulsen, Emily F. Burmeister, Biljana Stamenic, Geza Kurczveil, John E. Bowers, and Daniel J. Blumenthal
Electrical and Computer Engineering Department, University of California, Santa Barbara
jmack@ece.ucsb.edu

Abstract: The first demonstration of asynchronous optical packet transmission, scheduling, and buffering is presented using asynchronous/autonomous transmitters, lookup, buffer and a burst mode receiver. Contention resolution is shown with greater than 99.8% packet recovery at 10Gb/s.

OCIS codes: (060.1810) Buffers, couplers, routers, switches, and multiplexers; (230.4480) Optical amplifiers

1. Introduction

Optical packet switching (OPS) is a potential solution to achieving high capacity routing while addressing the impending power and footprint limitations of ultra-high capacity electronic packet routing systems [1]. Optical packet switches must operate asynchronously and autonomously for the scalability of optical networks [2]. This requires that transmitters, routers, and receivers operate on independent plesio-synchronous clocks which will not have the exact same frequency and phase. Individual packet arrival times at any given moment will be independent and random across all router input ports. Packet networks employ statistical time division multiplexing (TDM) on the packet level leading to packets within each stream arriving at a router port asynchronous to one another. Optical packet switches must be able to handle contention resolution of asynchronous packets simultaneously destined for the same output port. Optical buffering is necessary to avoid temporal collisions of contending packets, handle congestion and manage link utilization at the router output ports [3].

In recent years, there have been many advances in asynchronous optical packet switching and buffering. An asynchronous labeled swapped optical packet switch has been proposed and demonstrated at 40 Gb/s [4]. Synchronous optical buffering of 40 Gb/s optical packets has been shown using integrated buffer technologies [5]. Optical buffering at 40 Gb/s has been demonstrated with the buffer and lookup operated asynchronously to a synchronously operated transmitter and receiver [6]. An optically buffered packet switch has also been demonstrated at 40 Gb/s with the transmitter, lookup, and receiver operated asynchronously [7].

In this paper, we present the first demonstration of truly asynchronous optical buffering where custom transmitters, control circuitry, buffer, and receiver operate on independent clocks and are completely asynchronous to one another. In the experiments referenced above, contention was pre-engineered by using a single transmission source. The next critical step, demonstrated here is where optical packet switches and buffers handle asynchronous packets coming from multiple independent sources. Contention can occur for packets destined for the same output port regardless of load. Contention resolution using optical buffering is shown for 10 Gb/s packets originating from multiple independent transmission sources with greater than 99.8% packet recovery. Optical buffering is demonstrated using a packaged integrated InP switch matrix with a fiber delay line. A custom 10 Gb/s burst mode receiver is used where clock and data are successfully recovered for TDM packets from the buffer and contention path. Packet recovery measurements are presented based on real-time data analysis of the recovered asynchronous packets.

2. Asynchronous transmitters

The custom transmitters used in this paper are FPGA based running on a 625 MHz local clock which is internally divided down to a 156.25 MHz clock for the main processing. Each transmitter has its own independent clock labeled as clock 1 and 2 which is shown in Fig. 1. Packet streams are generated using custom software on a PC and loaded into 64 bit registers in the FPGA through a USB interface. The 64 bit wide registers are multiplexed from 156.25 Mbps to 16 channels at 625 Mbps in the FPGA. The data is further multiplexed to a single serial 10 Gb/s line using an external 16:1 serializer circuit. In this experiment, both transmitters are loaded with the same packet stream which consists of a 40 byte NRZ

packet at 10 Gb/s with a total period of 128 ns. The 40 byte packets contain a 32 bit idler, 64 bit unique packet identifier, and repeated PRBS 2^7-1.

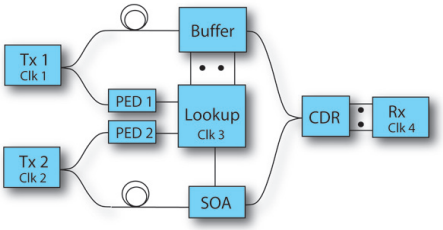

Fig. 1. Asynchronous transmitters, PEDs, lookup, buffer, SOA, CDR, and receiver

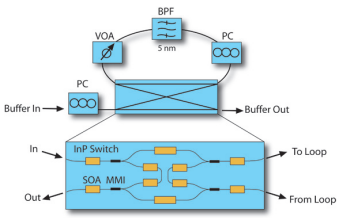

Fig. 2. Re-circulating optical packet buffer with inset of 2x2 InP switch

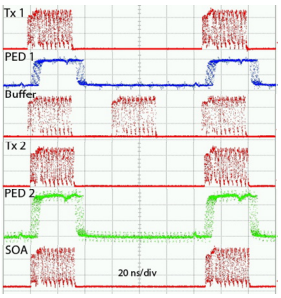

Fig. 3. Transmitter 1, PED 1, buffer, transmitter 2, PED 2, and SOA

3. Asynchronous contention resolution

Successful contention resolution is demonstrated using asynchronous PEDs, lookup, buffer, and a contention path as shown in Fig. 1. The PEDs provide a precise timing reference of the packets coming from the two transmitters. The electronic lookup is FPGA based running on a 156.25 MHz local clock, labeled clock 3. The lookup latches incoming PEDs into the FPGA using D flip-flops running on the local clock and are shown in Fig. 3. The timing uncertainties of when asynchronous packets are detected can be seen on the rising and falling edges of the PEDs. Furthermore, the PEDs of the two sources will wander from one another due to the plesio-synchronous nature of the transmitters' clocks. The buffer used in this experiment is a packaged InP switch with a fiber delay line which is depicted in Fig. 2. The delay line consists of a variable optical attenuator, band pass filter, and a polarization controller. A contention path was used on the second transmission source to force contention to occur. The total length of the delay line was chosen to be 64 ns which is twice the size of a 40 byte packet at 10 Gb/s. The delay line was chosen to be twice the packet size because only one buffer was used in this experiment and must be able to delay packets without temporal collisions regardless of the relative timing differences of the two PEDs. The lookup compares the two asynchronous PEDs and determines if contention is present. If no contention is present, then the packet from transmitter 1 is gated through the buffer. If the two PEDs overlap then the packet from transmitter 1 is loaded into the buffer. The packets from transmitter 2 are simply gated by an SOA to avoid intraband crosstalk. Screenshots were taken for packets from the buffer and the contention path by using a trigger signal from the corresponding transmitter as shown in Fig. 3. As can be seen, there are two states that occur over time for the buffer; either a packet passes through or circulates once in the buffer. How often these states change is based on the difference in frequencies of clock 1 and 2 of the transmission sources which manifests the wander of the two packet streams.

2. Asynchronous clock and data recovery

A key functionality of optical packet switching is the ability to recover data from asynchronous optical packets. The nature of asynchronous OPS does not allow for the use of commercially available clock and data recovery (CDR). This is due to the fact that commercially available CDR circuits are designed to operate on packet streams where data is synchronous on the bit level although asynchronous to the receiver. In OPS, packets within the same stream are asynchronous on both the bit and packet level. This requires custom clock recovery on a per packet basis which needs to be on the order of ns for high bit rates.

The custom clock and data recovery circuit used in this experiment is depicted in Fig. 4. The packet stream is converted to the electrical domain using a 10 GHz photo detector followed by a 10 GHz limiting amplifier. The recovered data is split using a 2x2 cross point switch where part of the data is fed through and the other injected into a frequency doubler to convert the 5 GHz base bandwidth of the NRZ data to a 10 GHz signal. The doubled signal is then filtered using a narrow band pass filter with a center frequency of 9.95328 GHz and bandwidth of 400 MHz which gives out a 10 GHz tone when data is present with 800 ps rise and fall times. The recovered clock travels through a tunable delay to finely align the recovered data and clock. The recovered clock is amplified and injected into the positive input of a 10 GHz differential limiting amplifier. Since the recovered clock is only present when data is present, an attenuated 10 GHz dummy clock derived from the receiver FPGA is injected into the negative input of the limiting amplifier. This insures that a clock is always present regardless if data is present. The dummy clock is attenuated so that when data is present the recovered clock gains dominance and only the recovered clock exits the output of the amplifier. When data is not present, the dummy local clock gains dominance and

outputs from the amplifier. The mixed clock is filtered and the data is coarsely delayed to match the clock recovery circuit. The recovered data, clock, and mixed clock of the SOA path are shown in Fig. 5 as an example of the CDR circuit operation. The TDM packets from the buffer and contention path triggered from transmitter 2 are shown in the bottom of Fig. 5. Since the trigger from transmitter 2 was used the packets from the contention path can clearly be seen while the packets from the buffer appear blurred as the trigger used was asynchronous to transmitter 1. Both the recovered data and clock are used as inputs of a 1:16 de-serializer that converts the 10 Gbps asynchronous data to 625 Mbps parallel data channels and the recovered clock to 625 MHz. The recovered data and clock enter a receiver FPGA based board operating on 625 MHz and divided 156.25 MHz clocks, labeled clock 4. The data is further de-multiplexed down to 64 bit parallel channels from 625 Mbps to 156.25 Mbps and the recovered clock is divided from 625MHz to 156.25 MHz. The FPGA implements a 64 bit wide by 512 bit deep asynchronous first in first out RAM block to transfer the recovered data to the local clock domain of the FPGA for performance analysis.

4. Asynchronous real-time performance analysis

The receiver FPGA performs real-time packet analysis by looking for the 64 bit unique packet identifier within the recovered data. If a packet is identified successfully then the FPGA updates a counter of packets recovered. The counter is written into a register that is read through a USB interface using custom software. A GUI is used on a PC to display packet recovery in real-time which can also change the capture length and expected packet recovery on-the-fly. The recovered counter is compared to the expected number of packets received to obtain a real-time packet recovery measurement. Packet recovery results are shown in Fig. 6 for the asynchronous back-to-back from transmitter 1, transmitter 2 and TDM packets from the buffer and contention path. Greater than 99.8% packet recovery is achieved for the TDM packet stream. Since the system was operated asynchronously, two states of the buffer will occur over time; either a packet is buffered or not, which is indistinguishable at the receiver. Packet recovery of the TDM stream is a measurement of the superposition of states in the system where each state results in recovered packets with different insertion losses and optical signal to noise ratios.

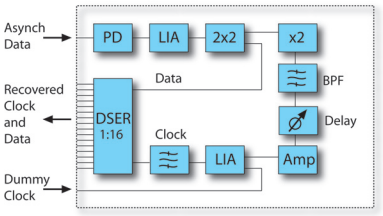

Fig. 4. Schematic of clock and data recovery (CDR) circuit

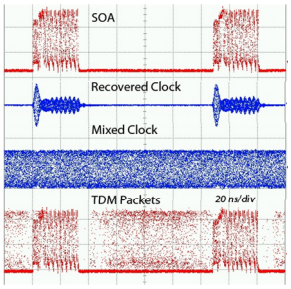

Fig. 5. SOA packets, recovered clock, mixed clock, and TDM packets

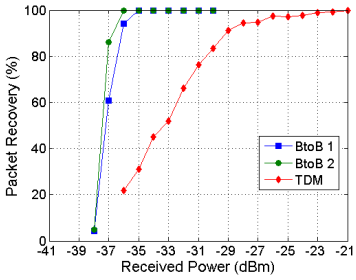

Fig. 6. Packet recovery measurements for back-to-back Tx 1, Tx 2, and TDM packets

5. Summary

Completely asynchronous and autonomous optical buffering is demonstrated using custom independent transmitters, control circuitry, buffer, and a burst mode receiver operating on different clocks. Successful contention resolution of optical packets is shown with greater than 99.8% packet recovery

6. Acknowledgement

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

7. References

[1] S.J.B. Yoo, "Optical packet and burst switching technologies for the future photonic internet," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4468-4492, Dec. 2006.

[2] D. Wolfson *et al*., "Synchronizing optical data and electrical control planes in asynchronous optical packet switches," in *Proc. OFC 2006*, Paper OThM3.

[3] D. K. Hunter, M.C. Chia, and I. Andonovic, "Buffering in optical packet switches," *J. Lightw. Technol*., vol.16, pp. 2081-2094, Dec. 1998.

[4] D. Wolfson *et al*., "All-optical asynchronous variable-length optically labeled 40 Gb/s switch," in *Proc. ECOC 2005*, Paper Th. 4.5.1.

[5] E.F. Burmeister *et al*, "SOA gate array recirculating buffer for optical packet switching," in *Proc. OFC 2008*, Paper OWE4.

[6] J.P. Mack, H.N. Poulsen, E.F. Burmeister, J.E. Bowers, and D.J. Blumenthal, "A 40 Gb/s asynchronous optical packet buffer based on an SOA gate matrix for contention resolution," in *Proc. OFC 2007*, Paper Th. 4.5.1.

[7] N. Wada, H. Furukawa, T. Miyazaki, "Prototype 160-Gbit/s/port optical packet switch based on optical code label processing and related technologies," *JSTQE*, vol.13, no. 5, pp. 1551-1559, Sept/Oct. 2007.