OThI5.pdf

Interface Optical Buffer and Packet-Switched Network Cross-Layer Signaling Demonstration

Caroline P. Lai, Howard Wang, and Keren Bergman
Department of Electrical Engineering, Columbia University, 500 W. 120th St., New York, New York 10027
caroline@ee.columbia.edu

Abstract: Queue management in an interface optical buffer is demonstrated via interoperability with an optical packet switched network. Cross-layer signaling is employed between the input buffer and network to dynamically and error-free reroute dropped wavelength-striped packets.

OCIS Codes: (060.4250) Networks; (200.4490) Optical buffers

1. Introduction

Optical packet switched (OPS) networks represent a promising technology approach for dynamically managing the immense growth of high bandwidth data-driven traffic in next generation routers and performance computing systems [1]. These networks can facilitate the low latency transmission of high bandwidth multi-wavelength optical packets by establishing end-to-end transparent lightpaths from the source to the desired destination. A significant challenge to the implementation of OPS networks is resolving contentions, which may occur at a given switching node as multiple packets attempt to leave on the same output link. In electronic network counterparts, contentions are straightforwardly addressed by storing packets and forwarding them when the path is freed; however, contention resolution in OPS networks cannot be as readily achieved due to the lack of practical optical random access memory elements. Small capacity optical packet buffers at the network input interface can partially mitigate this limitation by accepting backpressure and controlling the traffic injected into the network [1,2]. These buffers consist of queues that store a copy of the input packets prior to injection into the network. Once a packet obtains a contention-free path through the network and is successfully transmitted, it is discarded from the queue. In the case of an unsuccessful message transmission, the input buffer can reattempt injection with the copy of the message stored in the queue [1].

We have recently presented a modular optical packet buffer architecture and have demonstrated its reconfigurability and flexibility with respect to network performance through the implementation of active queue management [3-5]. In this work, the basic buffer architecture is adapted to realize the functionality of a network interface packet injection control to an implemented OPS network. The OPS network test-bed is based on the SPINet architecture [6], a switching fabric comprised of wideband 2x2 photonic switching nodes organized as a multistage interconnection network in the Omega topology. Messages are switched using semiconductor optical amplifier (SOA) gates, yielding wideband transmission, data transparency, and packet-rate granularity. SPINet does not employ optical buffering within the switching nodes and messages are subsequently dropped upon contention. A physical layer acknowledgement protocol provides a drop-detection mechanism in which an optical *ack* pulse is sent to the source following successful transmission. Re-transmission can then occur with minimal latency and reduced penalty due to message dropping.

In this paper, we experimentally demonstrate the interoperability between the implemented network interface packet buffer and a 4x4 SPINet OPS network test-bed. The network interface buffer actively queues packets preceding injection to the OPS network (Fig. 1). Cross-layer signaling via the SPINet physical layer *ack* pulses is employed at the buffer to ascertain successful or unsuccessful transmission through the network: correctly routed packets are discarded by the buffer, while packets dropped due to contention are dynamically retransmitted. Multi-wavelength optical packets containing 6x10 Gb/s wavelength-striped payloads are transparently processed at the interface buffer and correctly routed through the OPS network with confirmed error-free transmission on all payload wavelengths (BER<10^{-12}).

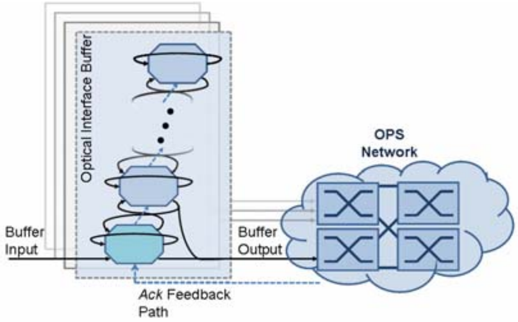

Fig. 1. Interface buffer connected to OPS network. Packets are first copied to the buffer, then injected into the network. If network transmission is successful, a control signaling *ack* is sent to the buffer and the copy is discarded. If no ack pulse is received, the packet is retransmitted in the next timeslot.

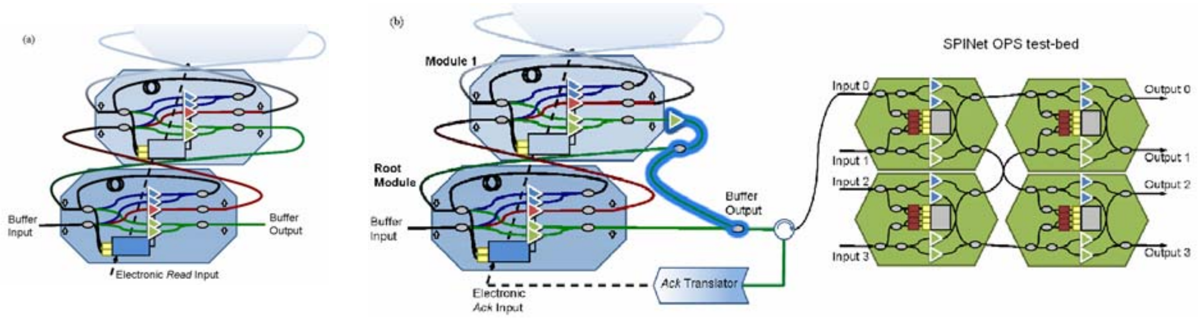

Fig. 2. (a) Original buffer architecture with two modules; (b) Implemented two-module buffer configuration with SPINet test-bed: SOAs are triangles, couplers are ovals, CPLDs are blue and gray rectangles, optical filters are brown rectangles, photodetectors are yellow rectangles.

2. Buffer Architecture & Functionality

The basic optical packet buffer architecture is comprised of identical, independent building-block modules arranged end-to-end and organized in a cascaded hierarchical structure (Fig. 2a). Each module provides the capacity to store a single optical packet of fixed length in an internal fiber delay line (FDL). Packets ingress into the buffer in a time-slotted manner via the root module and are stored in the module's FDL if the buffer is empty. Otherwise, packets are propagated upward along the height of the cascade until an unoccupied module is encountered. In a typical implementation of this packet buffer, first-in first-out (FIFO) packet ordering is simply maintained by ensuring that the age of a given packet corresponds to its position in the module stack. Further, reading packets from the buffer is performed independently of the write process. An electronic *read* signal is transmitted to the root module, which will subsequently forward the contents of its FDL to the output of the buffer. The *read* signal is then regenerated and retransmitted from module to module in the cascade, advancing the buffered packets incrementally towards the root module. Each module is a self-contained unit requiring no central management; this yields easy scalability as increased buffer capacity is feasible by simply connecting additional buffer modules to the end of the cascade.

To implement the network interface buffer for this investigation, we modify the behavior of the buffer modules via the programmable modules to provide interoperability with the SPINet OPS network (Fig. 2b). The modified architecture stores and transmits packets at each timeslot until an acknowledgement is sent by the network. *Ack* pulses are received within the same timeslot and replace the electronic *read* signal in the prototypical design. Upon reception of an *ack*, the currently transmitting packet is discarded and the next packet in the buffer is immediately forwarded to the network. To provide a direct egression path for the next available packet, the buffer architecture is modified to provide an additional output pathway from the module immediately above the root module.

3. Experimental Validation

We demonstrate the interoperable performance of the interface buffer with the SPINet OPS test-bed. Optical packets are first injected into the buffer and are subsequently stored in a queue until successful network transmission. To demonstrate the system's broadband transparency, multi-wavelength packets maintain a wavelength-striped structure such that the control information (framing and address) is encoded on dedicated wavelengths, while the payload is segmented and modulated at 10 Gb/s on six additional wavelengths across the C-band.

A two-module experimental prototype of the network interface buffer (Fig. 2b) is implemented for the demonstration experiments. The decision logic is synthesized in a high speed Xilinx complex programmable logic device (CPLD), with two distinct logic truth tables: one pertaining to the root module and another for subsequent, higher-order modules. Each building block module is comprised of commercially available components: the aforementioned CPLD, five SOA switching elements, and two 155 Mb/s p-i-n photodetectors; no optical filters are necessary. The 4x4 SPINet experimental test-bed (Fig. 2b) is composed of four 2x2 wideband switching nodes that are also realized with discrete components. The nodes decode the control information by extracting the two header bits (frame and address) using fixed wavelength filters; the CPLD then gates the four SOAs within the switching node in order to route the packets to their desired destinations (or drop them if necessary).

Both the buffer and network implementations utilize SOAs as the switching elements, which allow for insertion loss compensation. In this manner, each SOA hop contributes no net gain or loss, and packet longevity is maintained. For the interface buffer implementation, an additional SOA is required at the output of the second module; this SOA acts as a simple amplifier for the purpose of gain equalization between the first and second buffer modules. The implemented system supports 128 ns timeslots, each containing 115.2 ns duration packets with 10 Gb/s modulated data on six payload wavelengths. The packets are modulated by a single $LiNbO_3$ modulator with a

2^7-1 pseudorandom bit sequence (PRBS) at 10 Gb/s in non-return-to-zero (NRZ) format. Cross-layer signaling *ack* pulses are generated to inject packets from the buffer into the network.

Fig. 3a depicts the optical packet sequence for the experiment; Fig. 3b shows the input and resultant output waveforms for packets emerging from the integrated operation of the interface buffer and network. All packets received at the network output are shown to be correctly routed. Packet A is first stored in the buffer and injected into the network; it is then successful transmitted and discarded from the buffer. Two timeslots later, Packet B is injected into the buffer and network. Simultaneously, another network port injects Packet D and contention occurs between Packets B and D. Packet D is received at the network output, while Packet B is dropped. Due to this, no *ack* is received, and the buffer then re-injects the stored copy of Packet B. Also, Packet C appears at the buffer input and is thus stored in the second buffer module. The second transmission of Packet B is successful, and in accordance with FIFO ordering, Packet C is then injected into the network by the buffer and is successfully transmitted.

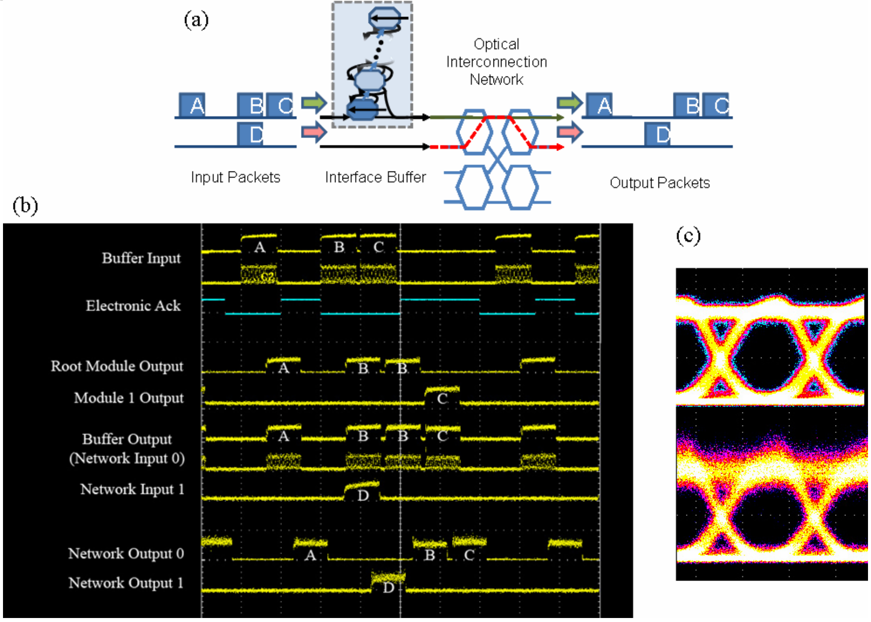

Fig. 3. (a) Diagram depicting experimental packet sequence. Contention occurs between Packets B and D; thus, Packet B is retransmitted at a later timeslot; (b) Optical waveforms pertaining to the buffer and network input and output signals; (c) 10 Gb/s input and output optical eye diagrams of Packet C (experiencing 6 SOA hops) at 1558.31 nm.

Fig. 3c portrays the input and output eye diagrams for one of the payload wavelengths at 10 Gb/s in Packet C, which undergoes six SOA hops: the greatest number of SOAs experienced by any packet in the experiment. At the network output, the packet is received by a p-i-n photodiode with a transimpedance amplifier and limited amplifier pair. The received electronic signals are then sent to a bit-error rate tester (BERT) that is synchronized with the packet gating signal and the bit pattern is driven by a pulse pattern generator (PPG). Bit-error rate measurements demonstrate that packets emerge from the output of the network error-free with BERs less than 10^{-12} on all six payload wavelengths.

4. Conclusions

Cross-layer signaling is demonstrated through the joint operation of a packet injection control interface buffer and an OPS network. Multi-wavelength optical packets are transparently processed at the interface buffer queue and dynamically rerouted through the OPS network. The wavelength-striped optical packets containing 6x10 Gb/s payloads are shown to be correctly routed through the network, with error-free transmission verified.

We acknowledge support for this work from the U.S. Department of Defense under subcontract B-12-664.

5. References

[1] R. Ramaswami and K. N. Sivarajan, *Optical Networks - A Practical Perspective, 2nd ed.* (Morgan Kaufmann, 2002).
[2] M. Enachescu, Y. Ganjali, A. Goel, N. McKeown, T. Roughgarden, “Routers with Very Small Buffers,” IEEE INFOCOM 2006.
[3] B. A. Small, A. Shacham, K. Bergman, “A Modular, Scalable, Extensible, and Transparent Optical Packet Buffer,” J. Lightwave Technol. 25 (4) 978-985.
[4] A. Shacham, K. Bergman, “Optical Packet Buffers with Active Queue Management [invited],” ONDM 2007.
[5] H. Wang, C. P. Lai, A. Shacham, K. Bergman, “Experimental Demonstration of Network Congestion Control with a Programmable Optical Packet Injection Buffer,” LEOS 2007 ThG2.
[6] A. Shacham, H. Wang, K. Bergman, “Experimental Demonstration of a Complete SPINet Optical Packet Switched Interconnection Network,” OFC 2007, OThF7.