

Priority Encoding Scheme for Contention Resolution in Optical Packet-Switched Networks

Caroline P. Lai (1), Howard Wang (1), Assaf Shacham (2), and Keren Bergman (1)

1 : Department of Electrical Engineering, Columbia University, New York, NY 10027; caroline@ee.columbia.edu

2 : Aprius Inc., Sunnyvale, CA 94085

Abstract

A priority encoding scheme is introduced for an optical switching fabric testbed, resolving contention by dropping low priority messages. Optical packets with 6x10 Gb/s payloads are routed error-free.

Introduction

Optical packet switching (OPS) is a valuable, scalable approach for obtaining high throughput traffic routing while satisfying the low power and latency requirements of next generation routers and high performance computing systems [1]. OPS switching fabric architectures offer a low latency networking infrastructure for transmitting high bandwidth multi-wavelength optical packets. One important functional feature for OPS networks is the ability to route prioritized messages. Applications that leverage OPS networks may support a vast number of higher quality connection-oriented services [2]. The network will then be required to offer differentiated services and protocols on the physical layer through varying levels of quality-of-service (QoS) and priority. Thus, functionality and performance will be greatly enhanced by implementing different packet priorities, ultimately leading to the realization of a diverse number of classes of services in OPS networks.

Additionally, a significant challenge in the practical implementation of OPS networks is contention resolution, which may be necessary at a network switching node as multiple packets attempt to leave simultaneously on the same output link. Due to the lack of practical optical random-access memory (RAM) technology, contention resolution in OPS networks is not as straightforward as electronic store-and-forward methods. Current methods involve costly packet-dropping and retransmission, which greatly degrades performance. Thus, different packet priority classes may further alleviate the high cost of contention resolution in the optical domain [3].

Here, we introduce a priority encoding scheme to address contention resolution in OPS networks. The implemented 4x4 3-stage OPS fabric testbed utilizes the SPINet architecture [4]. The switching nodes do not employ optical buffering; instead, contending messages are dropped and retransmitted at a later time. The prioritized routing scheme prevents a high-priority packet from being dropped, thus yielding an overall reduction of retransmission penalty for critical data paths. Correct routing and error-free transmission ($BER < 10^{-12}$) are verified for wavelength-striped messages.

Switching Fabric Architecture

The network testbed is realized in the SPINet architecture, which is composed of 2x2 non-blocking wideband photonic switching nodes organized in the multistage Omega topology, with messages establishing end-to-end lightpaths that extend across the network (Fig. 1). The network is designed to be integrated on a PIC, utilizing semiconductor optical amplifier (SOA) gates to switch messages, yielding wideband transmission, data transparency, and packet-rate granularity. The packets have a wavelength-striped format (Fig. 1), with the control information such as frame and address encoded on dedicated wavelengths. The payload is segmented and modulated at a high data rate (here, 10 Gb/s per wavelength) on the rest of the available band [4].

The switching nodes decode the control information immediately upon reception of the packets' leading edges using a wavelength filter and low-speed optical receiver. The nodes' switching elements are four SOAs arranged in a gate-matrix structure. When the framing and address signals are recovered from the ingressing packet and processed by the high-speed electronic circuitry, the appropriate SOAs are switched on and the optical messages are either routed to their desired destination or dropped upon contention. A physical layer acknowledgment (*ack*) protocol is realized via short optical pulses that are sent in the reverse direction from the receiving port upon successful transmission along the lightpaths created by the successfully routed messages. Sources of dropped messages do not receive *acks*, which can then attempt retransmission at a later time.

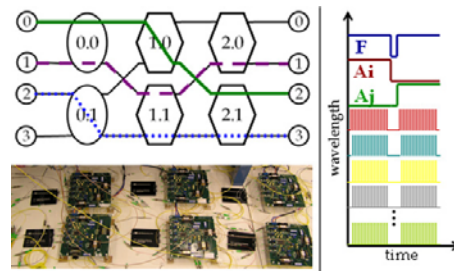


Figure 1: Network topology and photograph (left) and wavelength-striped optical message format (right)




Figure 2: Priority-encoded packet structure

Priority Encoding Scheme

Owing to the high level of reprogrammability of the switching nodes, the existing network testbed can be easily adapted for prioritized packet transmission. The priority encoding scheme is supported through straightforward modification of the switching node logic and optical packet format. The frame header signal is modified to incorporate a one bit priority and is sampled using a low-duty pseudo-clock, which has two pulses per timeslot to determine the presence of a packet and its priority (Fig. 2). According to the priorities encoded in the frame, a high or low class of service for the packet is determined. In the case of a contention, the modified routing logic will gate the SOA associated with the high priority message, and drop the contending low priority packet.

Experimental Results

The 6-node 3-stage experimental network testbed is implemented using individually-packaged elements (SOAs, passive optics, optical receivers, and digital electronics). In order to demonstrate the prioritized packet routing scheme, a pattern of wavelength-stripped packets is injected via three independent input ports with a combination of high and low priority encoding. The experimental demonstration supports 57.6 ns timeslots, containing 51.2 ns duration packets with data modulated at 10 Gb/s on six payload wavelengths ranging from 1540.1 nm to 1558.3 nm. The packets are modulated by a single LiNbO₃ modulator with 2⁷-1 PRBS NRZ. The optical header of each packet is comprised of a frame signal encoding a one bit priority, a distribution address (A0, selecting one of two possible paths), and a routing address (A1, A2). Fig. 3 shows correct routing via the input and output optical waveforms of the pseudo-clock and optical packets: the frame with priority bit, indicated address information, and 6×10 Gb/s payload. Faded waveforms in Fig. 3 refer to contending low priority packets that are dropped and retransmitted later.

The experimental packet sequence demonstrates the functionality of the prioritized packet scheme. The exploration exemplifies one high-priority source (in1), one low-priority source (in2), and one source whose retransmitted packets are given higher priority (in0). When messages contend at a given switching node, the lower-priority packet is dropped, and no ack pulse

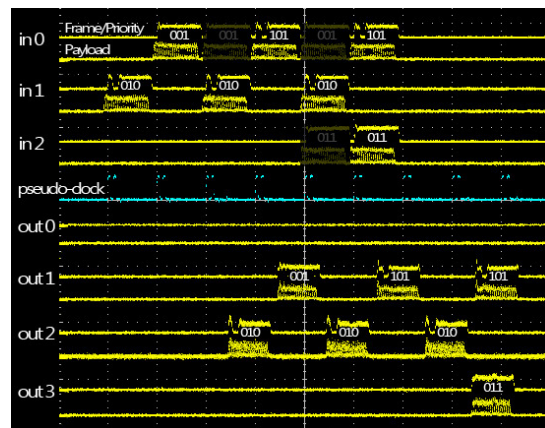


Figure 3: Experimental optical waveforms




Figure 4: 10 Gb/s eye diagrams of the input (left) and output (right) signals ($\lambda = 1558.3 \text{ nm}$); BERs $< 10^{-12}$ are obtained for all 6 payload channels

is received at its input port. Here, in the macro-scale testbed realization, acks are not used due to the large round-trip latency of the testbed (270 ns), compared to the envisioned integrated realization. Thus, acks are assumed to be received after 10 ns. If the source does not receive an ack, retransmission occurs on a different path during the following timeslot by modifying the packet's distribution address. The retransmitted packet can have equal or higher priority. Error-free transmission of all high and low priority packet payloads is confirmed for the network testbed. Bit-error rates (BERs) less than 10^{-12} are verified for all six payload wavelengths. The input and output eye diagrams at 10 Gb/s are provided in Fig. 4.

Conclusions

A priority encoding scheme for an OPS network is successfully demonstrated, with two classes of frame-encoded packet priority. The network supports high and low priority levels, as well as prioritized retransmission upon message-dropping. Correctly routed wavelength-stripped packets with 6×10 Gb/s payload are transmitted error-free ($\text{BER} < 10^{-12}$). Thus, this investigation demonstrates the potential for implementing QoS protocols in OPS networks.

References

1. Ramaswami *et al.*, *Optical Networks: A Practical Perspective*, Morgan Kaufmann, 2002.
2. Yao *et al.*, OFC (2001), Paper TuK2.
3. Rahbar *et al.*, *JON* 5 (2006), pp. 1056-1070.
4. Shacham *et al.*, *IEEE Micro* 27 (2007), pp. 6-20.