

# Demonstration of Programmable Broadband Packet Multicasting in an Optical Switching Fabric Test-bed

Caroline P. Lai and Keren Bergman

Department of Electrical Engineering, Columbia University, 500 W. 120<sup>th</sup> St., New York, New York 10027  
caroline@ee.columbia.edu

**Abstract:** Multicasting of broadband optical messages is experimentally demonstrated in a programmable optical switching test-bed. Wavelength-stripped optical packets with  $8 \times 10$  Gb/s payloads are multicasted error-free through the fabric, with confirmed scalability to 40 Gb/s data rates.

©2008 Optical Society of America

**OCIS Codes:** (060.4250) Networks; (060.4255) Networks, multicast; (060.4510) Optical communications

## 1. Introduction

Optical packet switching (OPS) constitutes a significant technology for achieving enhanced bandwidth and latency performance in next generation routers and networks [1]. OPS networks can support the low latency transmission of high bandwidth wavelength-stripped optical packets by creating transparent lightpaths across the network. In order for OPS networks to realize practical performance, greater functionality and programmable flexibility must be exhibited to optimize network performance with respect to the physical layer. Indeed, an important application to be realized is the ability to perform multicasting of broadband optical messages from a single source to multiple output destinations [2]. Multicasting will drive emerging high-bandwidth user-demand applications such as distributed computing, streaming video, and networked gaming.

In this paper, we experimentally demonstrate programmable multicasting of high-bandwidth multi-wavelength optical packets in a switching fabric architecture. The switching fabric is based on an earlier demonstrated OPS network test-bed [3]. We have previously shown the flexibility of the architecture by implementing cross-layer information exchange between the OPS network and an interface buffer, thus mitigating unsuccessful transmission and packet loss through the network [4]. The distributed nature of the fabric's control routing logic yields a high level of reprogrammability and reconfigurability, which is further leveraged in this work to demonstrate a straightforward enhancement to support the multicast operation. We propose multicasting in a switching fabric that is internally comprised of  $M$  OPS networks placed in parallel (Fig. 1a). In order to achieve the multicast-enabled capability, each source input is connected to  $M$  networks and thus to each destination output port by  $M$  separate networks that operate in parallel. One advantage of this architecture is that the switching fabric may either unicast on a single network or multicast on several networks. The switching fabric supports the simultaneous transmission of multiple broadband messages to multiple outputs without the need for optical buffers or fiber delay lines (FDLs). Also, due to the important ability to route multi-wavelength packets, the need for wavelength converters is mitigated. The architecture also provides increased path diversity in the case of contention and packet loss. The architecture is further specifically designed to support cross-layer communication capabilities to enhance the performance of the switching fabric by extracting physical layer monitoring measurements to reconfigure the network; in this way, we may leverage the multiple networks to optimize path routing through the fabric and offer dynamic quality-of-service (QoS) classes.

For the purpose of this experimental demonstration, we show multicasting in an implemented  $4 \times 4$  switching fabric that is comprised of two networks. Multi-wavelength optical packets containing  $8 \times 10$  Gb/s wavelength-stripped payloads are correctly routed and transmitted with confirmed error-free performance through the switching fabric with bit-error rates (BERs) less than  $10^{-12}$ . Due to the transparency of the coupled OPS networks, the data rates of the payload channels are easily scalable; we also show the successful transmission of 250 Gb/s aggregate bandwidth through a combination of 10 Gb/s and 40 Gb/s payload data rates.

## 2. Switching Fabric Architecture

The complete  $4 \times 4$  switching fabric architecture is composed of two OPS networks connected in parallel (Fig. 1b). The first parallel network is based on the 3-stage SPINet architecture, which is organized in the Omega topology [3]; the second parallel network is a 2-stage Banyan network. Multistage Banyan networks are advantageous as they provide interconnection for  $N$  ports with only  $\log_2 N$  stages of  $N/2$  switching nodes. Both networks are comprised of non-blocking programmable wideband  $2 \times 2$  photonic switching nodes, where messages are switched using semiconductor optical amplifier (SOA) gates, yielding broadband transmission, data transparency, and packet-rate

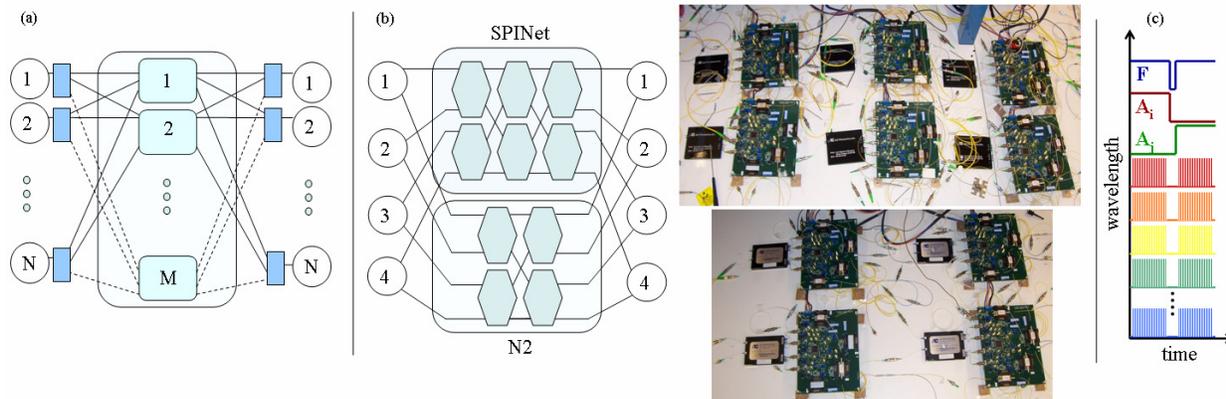


Fig. 1. (a) Diagram of multicast-enabled switching fabric with  $M$  parallel internal networks connecting  $N$  nodes; (b) Implemented switching fabric (block diagram and photograph); (c) Wavelength-striped optical packet format

granularity. No optical buffering is implemented within the switching nodes, thus messages are dropped upon contention. The switching fabric's packet format leverages wavelength-division multiplexing (WDM) to offer high transmission bandwidths. The messages have a wavelength-striped structure (Fig. 1c) such that the control information (frame and address) is encoded on a subset of dedicated wavelengths (with a single bit per wavelength), while the payload is segmented and modulated at a high data rate on the rest of the available band. The switching nodes decode the control information immediately at the leading edge of the optical message by extracting the two header bits using fixed wavelength filters and a low speed optical receiver. The four SOAs of each switching node are gated on according to the recovered header information, thus routing the appropriate optical messages to their desired destinations (or dropped upon contention). The SOAs also provide optical gain to compensate for the insertion loss of the passive devices; thus, each switching node contributes no net power gain or loss.

Messages which are successfully routed through the switching fabric create end-to-end transparent lightpaths. A physical layer acknowledgement (ack) protocol is realized using the bidirectional transparency of the switching fabric architecture. When the leading edge of the optical messages reach their desired destination output ports, short optical control ack pulses can be sent backwards to notify the source of successful transmission; these pulses can also be envisioned to be utilized for cross-layer control signals. In future applications, the packet-level physical layer performance, such as the BER at the receiver, can be encoded in the ack pulses and extracted by higher layers for path rerouting. Sources of dropped messages do not receive ack pulses and can then retransmit at a later time; retransmissions can occur within minimal time, resulting in reduced latency penalties.

The switching nodes utilize distributed control logic, yielding straightforward reconfigurability and flexibility. Each switching node leverages simple logic; no central network control plane management is required. Considering the distributed nature of the node routing, implementing multiple OPS networks in parallel is scalable as it does not grow a required control plane. This simple transition realizes the multicasting operation with limited additional complexity. Optical buffers are not necessary to facilitate the multicast capability in this fabric architecture design.

### 3. Experimental Demonstration

The multicast-capable 4x4 optical switching fabric, comprising of two networks operating in parallel, is implemented using ten 2x2 photonic switching nodes (Fig. 1b). The electronic decision logic is synthesized in a high speed Xilinx complex programmable logic device (CPLD). Each switching node is realized with discrete components: the CPLD, four SOAs, four 155 Mb/s p-i-n photodetectors, passive optics, and electronic circuitry.

A pattern of optical packets is injected into both networks simultaneously. The experimental system here supports 128 ns timeslots, incorporating 115.2 ns duration packets. The packet format of the initial setup uses 10 Gb/s modulated data on eight payload wavelength channels (8x10 Gb/s). The packets are modulated at 10 Gb/s with a LiNbO<sub>3</sub> modulator with a 2<sup>7</sup>-1 pseudorandom bit sequence (PRBS) in non-return-to-zero (NRZ) format, and gated into packets using a SOA. The optical waveforms associated with the injected optical packet sequence and subsequent network output packets are shown in Fig. 2a. All multi-wavelength optical messages containing 8x10 Gb/s payloads are shown correctly routed through the switching fabric, with the multicast operation clearly validated. The switching fabric straightforwardly supports both unicast and multicast operations. In the first active timeslot, the sources transmit only through the first network (SPINet). During the second timeslot, the sources unicast simply with the second network (N2). In the third timeslot, a single source (in0) multicasts to two different destinations (out2 and out3) simultaneously using both networks. During the fourth active timeslot, multiple sources perform broadband multicasting to multiple outputs through both networks. Here, acks are not implemented since

the round-trip time of the fabric is large (270 ns), compared to the envisioned integrated realization. We assume that acks are received (or not received) after 10 ns. Packets that are dropped are retransmitted upon not receiving an ack.

For the  $8 \times 10$  Gb/s setup, the packets are received by a p-i-n photodiode with a transimpedance amplifier and limited amplifier pair at the output of the switching fabric. The received electronic signals are then sent to the bit-error rate tester (BERT) that is synchronized with the packet gating signals. The pulse pattern generator (PPG) drives the bit pattern. BER measurements verify error-free transmission through both networks, with BERs  $< 10^{-12}$  confirmed on all eight payload wavelengths. Fig. 2b provides the sensitivity curves related to transmission through N2 ( $\lambda=1541.05$  nm); insets portray 10 Gb/s optical eye diagrams for the fabric input and output.

Further, we demonstrate the switching fabric's bandwidth transparency, attesting to the scalability of the initial 10 Gb/s payload data rates to higher modulation data rates. We show successful transmission of a combination of  $6 \times 40 + 1 \times 10$  Gb/s (250 Gb/s, modulated with  $2^7-1$  PRBS) through the fabric. The input (after one gating SOA) and output (one gating SOA and SPINet network) optical eye diagrams of one 40 Gb/s payload channel ( $\lambda=1558.24$  nm) is shown in Fig. 2c. Error-free transmission is attained on the 10 Gb/s payload channel (BER  $< 10^{-12}$ ).

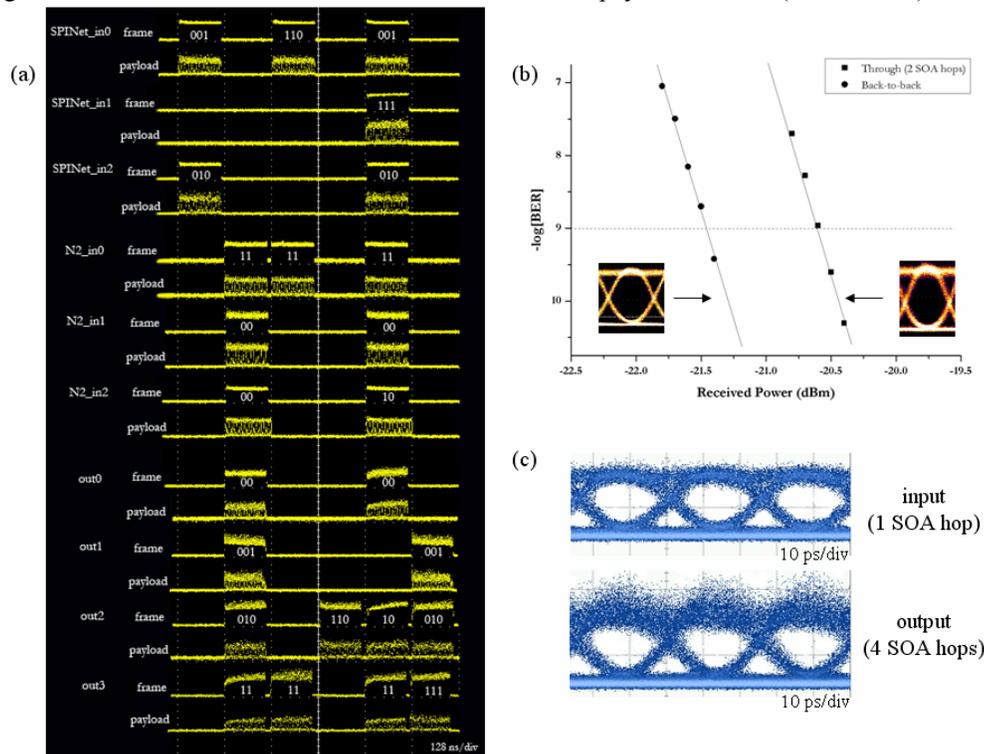


Fig. 2. (a) Optical waveforms pertaining to the packet sequence injected in both networks (SPINet and N2) with resultant output waveforms (labels refer to encoded address information); (b) BER curves for one 10 Gb/s payload of N2, inset: 10 Gb/s input and output optical eye diagrams; (c) 40 Gb/s input and output optical eye diagrams for SPINet

#### 4. Conclusions

Multicasting of broadband multi-wavelength optical packets is successfully demonstrated on an experimentally implemented programmable switching fabric test-bed consisting of two OPS networks operating in parallel. Wavelength-striped optical packets with  $8 \times 10$  Gb/s payloads are correctly routed, with error-free transmission verified in both networks for all payload wavelengths (BERs  $< 10^{-12}$ ). Data rate scalability is also shown with 40 Gb/s payloads. This work demonstrates the feasibility of straightforwardly implementing a broadband multicasting application on the optical physical layer and will constitute an important stepping stone for future endeavors of leveraging cross-layer communication to provide path rerouting based on real-time physical layer measurements. We gratefully acknowledge support for this work by the Intel Corporation and by the National Science Foundation under grant CNS-837995.

#### 5. References

- [1] R. Ramaswami and K. N. Sivarajan, *Optical Networks-A Practical Perspective*, 2nd ed. San Francisco, CA: Morgan Kaufmann, 2002, ch.12.
- [2] G. N. Rouskas, "Optical layer multicast: rationale, building blocks, and challenges," *IEEE Network* 17 (1) 60-65 (Jan 2003)
- [3] A. Shacham and K. Bergman, "Building Ultralow-Latency Interconnection Networks Using Photonic Integration," *IEEE Micro* 27 (4) 6-20 (Jul/Aug 2007)
- [4] C. P. Lai, H. Wang, K. Bergman, "Cross-Layer Communication With an Optical Packet Switched Network via a Message Injection Control Interface," *IEEE Photon. Technol. Lett.* 20 (12) 967-969 (Jun 2008)