Demonstration of Failure Reconfiguration via Cross-Layer Enabled Optical Switching Fabrics

Caroline P. Lai, *Student Member, IEEE*, Daniel Brunina, *Student Member, IEEE*, Cédric Ware, Balagangadhar G. Bathula, *Member, IEEE*, and Keren Bergman, *Fellow, IEEE*

***Abstract*—Growing bandwidth requirements of future Internet applications are driving the potential deployment of all-optical packet switching fabrics in next-generation routers. The switching fabric should be capable of executing a fast reconfiguration of its switching state, allowing for the dynamic management of optical packets and the seamless recovery of the fabric, in the case of IP-layer router failures and cross-layer enabled optical-layer signal degradations. We demonstrate a reconfigurable optical switching fabric architecture that uses a field-programmable gate array control plane. Based on the state of a higher-layer router, the switching fabric supports the correct routing and error-free transmission of 10 $\times$ 10-Gb/s wavelength-striped optical packets, with bit-error rates less than 10^{-12} on all payload channels. A power penalty less than 1 dB is shown.**

***Index Terms*—Optical communication, optical packet switching (OPS), photonic switching systems, restoration.**

I. Introduction

THE design of the next-generation Internet is driven by the need to address the exploding growth in bandwidth and network traffic. To achieve these high bandwidths, future networks may leverage all-optical transmission at high data rates with advanced modulation formats [1]. Enhanced physical-layer technologies, such as optical packet switching (OPS), are a promising, energy-efficient solution to realize dynamic optical-layer switching functionalities while simultaneously supporting high-throughput traffic [2]. OPS fabrics can be deployed in future routers to enable the transparent switching of multiwavelength optical messages.

In our envisioned bidirectional cross-layer optimized infrastructure, optical packets' transport and switching is affected by real-time performance monitoring metrics, which can be extracted on a packet timescale [3]. The cross-layer platform will allow for a unique means of dynamically optimizing performance and energy consumption. In addition to taking into account the introspective physical-layer awareness, it is also necessary for the switching state of the OPS fabric to be affected by higher-layer Internet Protocol (IP) parameters. In the case of costly IP router failures or when a router is in sleep mode for energy-saving benefits [4], connections between end nodes can be maintained by optically routing around the failures. Optical-layer recovery can also realize optical restoration mechanisms that protect clients against IP router failures [5]. To enable an optical bypass [6] and prevent network outage, an on-the-fly reconfiguration of the optical switching fabric will be required to mitigate failure.

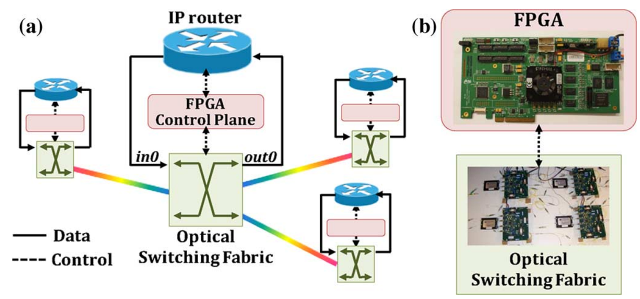

Fig. 1. (a) Envisioned network architecture with network nodes composed of IP routers, optical packet switching fabrics, and an FPGA-based control and management plane. The FPGA acts as a cross-layer interface, accepting control inputs to manage physical-layer switching. (b) Photographs of FPGA circuit board and implemented OPS fabric used in this demonstration.

This work shows an optical fabric that dynamically responds to failure by performing a reconfiguration based on external input signals (Fig. 1), with nanosecond response times. Upon the detection of a router failure or degraded optical packet streams, the fabric executes a fast, nanosecond-scale switching state reconfiguration to yield dynamic management of optical packet routing. Current state-of-the-art optical switch fabrics rely on slow technologies (e.g. micro-electro-mechanical systems (MEMS) or arrayed waveguide gratings (AWGs)), leading to high packet losses. We present a semiconductor optical amplifier (SOA) fabric with fast switching, low recovery latencies, and high bandwidths via wavelength-division-multiplexing (WDM), using a field-programmable gate array (FPGA) controller. FPGA control planes have been shown effective for SOAs [7]. Here, we demonstrate improved optical-layer performance and highlight the novel concept of cross-layer signaling to control an optical packet router, instantiating this capability at a packet rate. To our knowledge, this concept has not been raised within the scope of optical switching. The FPGA's fast processing realizes protection switching, showing an example

Manuscript received April 01, 2011; revised August 08, 2011; accepted August 22, 2011. Date of publication August 30, 2011; date of current version October 21, 2011. This work was supported in part by the National Science Foundation Engineering Research Center on Integrated Access Networks (CIAN) (subaward Y503160).

C. P. Lai was with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA, and is now with the Photonics Systems Group, Tyndall National Institute, University College Cork, Cork, Ireland (e-mail: caroline.lai@tyndall.ie).

D. Brunina, B. G. Bathula, and K. Bergman are with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: daniel@ee.columbia.edu; bbathula@ee.columbia.edu; bergman@ee.columbia.edu).

C. Ware was with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA, and is now with Institut Télécom, Télécom ParisTech, CNRS LTCI, 75634 Paris CEDEX 13, France (e-mail: cedric.ware@telecom-paristech.fr).

Color versions of one or more of the figures in this letter are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/LPT.2011.2166539

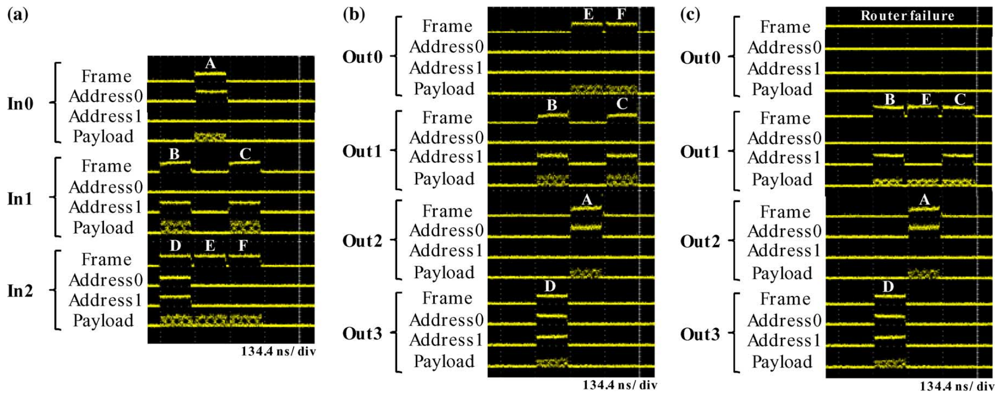

Fig. 2. Optical waveform traces of the experimental packets taken at the (a) fabric input; (b) fabric output under an online router scenario with all output ports functional; (c) fabric output in the case of an offline router, with switching fabric aware that link out0 is down/degraded.

of cross-layer feedback. Nanosecond reconfiguration will be required for future fabrics to effectively execute novel cross-layer algorithms.

The cross-layer failure recovery scheme is experimentally implemented on a 4 × 4 fabric. The FPGA acts as the fabric's control plane, enabling signals from higher layers to trigger recovery and reroute messages. FPGA control, in conjunction with nanosecond switching devices, enables fast cross-layer fabric recovery. We demonstrate the successful routing of 10 × 10-Gb/s wavelength-striped optical packets, for both cases of an online router (i.e. packets are transmitted correctly) and an offline router (i.e. the router or subsequent optical link is down, thus packets are rerouted according to the recovery switching logic). The fabric operates error-free, attaining bit-error rates (BERs) less than 10^{-12} on all payload channels.

II. Optical Switching Fabric

The experiment uses a 4 × 4 synchronous OPS fabric test-bed (Fig. 1) [8]. The two-stage design uses four 2 × 2 photonic switching nodes, which are comprised of four SOAs, providing wideband, packet-granular switching. The fabric supports 10 × 10-Gb/s wavelength-striped packets, with low-speed control headers (i.e. frame and two fabric address bits), which use dedicated wavelengths, and ten payload channels, each modulated at 10 Gb/s. The header bits are set constant for the message duration (one bit/timeslot/wavelength). The packets have an aggregate bandwidth of 100 Gb/s, with each payload channel being 120 ns in length, resulting in 1500-B messages, analogous to the Ethernet maximum transmission unit (MTU).

The wavelength-striped optical packets traverse the entire OPS fabric, with all wavelengths routed cohesively end-to-end. At each stage, the node extracts the control information using fixed wavelength filters and low-speed 155-Mb/s p-i-n photodetectors. The frame and address bits are processed by high-speed complex programmable logic devices (CPLDs). The appropriate SOAs are gated on by the CPLD, routing the optical packets to their desired destination (or dropped if they contend for the same 2 × 2 node output.) The CPLDs could also help with load balancing in the network by sending packets to multiple ports of the switching node. In this implementation, one control header is required per fabric stage (i.e. SOA hop); WDM optical designs have been shown to be scalable to large port counts with many stages while maintaining low BERs [9].

III. Failure Recovery Scheme

The failure recovery scheme allows the fabric's 2 × 2 switching nodes to account for router failures. When a failed/degraded link is detected, the fabric reconfigures its state to route around the failure, yielding dynamic switching. The Altera Stratix II GX FPGA controller accepts external inputs (e.g. from a router) and generates failure signals for the nodes. The CPLDs' routing logic is adapted to accept these electronic failure signals to either act normally (for an online router, packets are switched accordingly), or route around the failure (for the offline/failed router, packets are rerouted to ensure that no messages are transmitted to the link). As in Fig. 1, if the router is offline, packets that would have been transmitted to the router are instead rerouted to another output port if there is no contention; otherwise, they are dropped. The recovery scheme simply deflects packets to an alternate port on the same switching node to avoid failure. Upon a failure, the following sequence of events must occur: generating and processing the failure message, creating the fabric's control signal, reconfiguring the fabric on-the-fly, and the subsequent full fabric operation. The focus of this letter is on the fabric reconfiguration, providing a straightforward demonstration of the fabric's nanosecond scale recovery time. Enabling the entire feedback mechanism allows this sequence to occur dynamically according to higher-layer cross-layer algorithms.

IV. Experimental Demonstration and Results

We use a FPGA controller and a 4 × 4 fabric built with commercial parts, including four SOAs, digital electronics, optical receivers, and passive optics. Due to the lack of mature optical buffers, contending packets are dropped in the implemented fabric, which may result in high packet loss rates. An optical interface buffer may be realized to store copies of packets to mitigate contention-based packet loss [10]. 10 × 10-Gb/s wavelength-striped packets are switched (Fig. 2(a)) based on the failure state denoted by the controller. We use a FPGA circuit board, containing eight flip switches and 28 general purpose input/output (GPIO) pins. The FPGA is programmed to receive input from the flip switches, indicating a router failure, and signal the fabric using the GPIO pins. In the envisioned realization, the FPGA will be synchronized with the packet transmission; here, we simply show the fabric's reconfiguration time in the event of a failure.

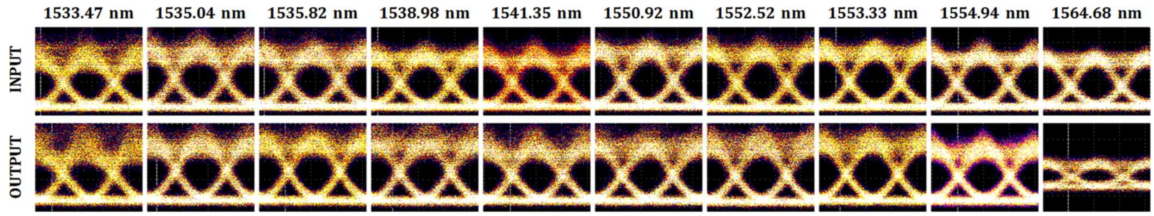

Fig. 3. A 10-Gb/s optical eye diagrams for the input and output of the switching fabric, for all ten payload wavelength channels.

The optical packets are generated using 13 separate continuous-wave (CW) distributed feedback (DFB) lasers: three DFBs for the control (frame: 1555.75 nm; addresses: 1531.12 nm, 1543.73 nm), and ten DFBs, ranging from 1533.47 nm to 1564.68 nm, for the payload wavelength channels. A single $LiNbO_3$ modulator encodes a $2^7 - 1$ nonreturn-to-zero (NRZ) pseudorandom bit sequence (PRBS) on each of the payload channels. The CW control and PRBS payload wavelengths create a multiwavelength signal, which is sent to a 1:3 passive splitter to form three independent input streams. Each stream is gated into a pattern of 1500-B packets using external SOAs (one stream per fabric input) (Fig. 2(a)).

Fig. 2 provides the optical waveforms of the experimental packet sequence, denoting the frame, address, and one 10-Gb/s payload channel for all input (Fig. 2(a)) and output (Fig. 2(b), (c)) ports. We show two explicit cases. First, Fig. 2(b) gives the output traces for an online router scenario, where all output ports are available and packets are transmitted correctly. Second, we show an offline router scenario (Fig. 2(c)), in which the output port corresponding to the router is designated as failed (or sleeping). The FPGA informs the fabric of the failure, so the fabric can reconfigure its state to reroute traffic and avoid sending packets to the failed node. Fig. 2 shows that no packets are transmitted to out0 (i.e. router); messages formerly intended for out0 (i.e. packets E and F) are rerouted to out1 (if the port is available). Packets C and F contend for out1, thus F is dropped; the logic prioritizes messages originally designated for the next node. In future implementations, a higher-radix switching node can be realized so that the control plane allows the fabric to deflect these packets to any of the available output ports, accounting for possible contention with other packets. This new switching node will also mitigate concerns of fabric scalability.

All 10×10-Gb/s packets are transmitted error-free. Signal integrity is verified using a DC-coupled 10-Gb/s p-i-n receiver, and a bit-error-rate tester (BERT) that is synchronized with the packet gating signals. BERs less than 10^{-12} are attained on all ten payload channels; Fig. 3 shows the 10-Gb/s input and output eye diagrams for all the payload channels. BER sensitivity curves for the channel exhibiting the most degraded optical eye are given (Fig. 4). The two-stage fabric has a 0.9-dB power penalty (0.45 dB/SOA hop), taken at a BER of 10^{-9}.

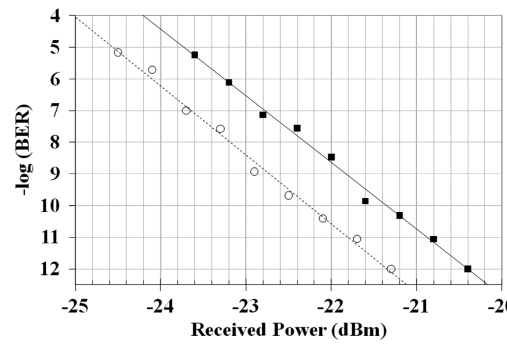

Fig. 4. Sensitivity curves for one payload channel ($\lambda = 1564.68$ nm). The dashed line with open points corresponds to the back-to-back measurements, while the solid line with filled points refers to the data at the fabric's output.

V. Conclusion

We show an agile OPS fabric that can be seamlessly reconfigured in nanoseconds via a cross-layer control plane to realize enhanced optical switching functionalities. The fabric is reconfigured on-the-fly using FPGA control signals, providing a means of protecting packet transmission and realizing optical bypasses that allow traffic to bypass failed nodes. The fabric supports the error-free transmission of wavelength-striped optical packets. This work is fundamental to creating a cross-layer node that uses a reconfigurable fabric with dynamic failure response, a FPGA control plane, and monitoring devices, to enable innovative technologies for future access/aggregation networks.

References

[1] R. W. Tkach, "Scaling optical communications for the next decade and beyond," *Bell Labs Tech. J.*, vol. 14, no. 4, pp. 3–10, Feb. 2010.

[2] S. J. B. Yoo, "Energy efficiency in the future internet: The role of optical packet switching and optical-label switching," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 406–418, Mar. 2011.

[3] C. P. Lai and K. Bergman, "Cross-layer communications for high-bandwidth optical networks," presented at the 12th Int. Conf. Transparent Optical Networks (ICTON) 2010, Munich, Germany, Jun. 2010, Paper Tu.C3.2.

[4] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Energy efficiency in telecom optical networks," *Commun. Surveys Tuts.*, vol. 12, no. 4, pp. 441–458, Fourth Quarter, 2010.

[5] A. L. Chiu and J. Strand, "An agile optical layer restoration method for router failures," *IEEE Network*, vol. 17, no. 2, pp. 38–42, Mar./Apr. 2003.

[6] G. Shen and R. S. Tucker, "Energy-minimized design for IP over WDM networks," *J. Opt. Commun. Netw.*, vol. 1, no. 1, pp. 176–186, Jun. 2009.

[7] T. Lin, K. A. Williams, R. V. Penty, I. H. White, and M. Glick, "Capacity scaling in a multihost wavelength-striped SOA-based switch fabric," *J. Lightw. Technol.*, vol. 25, no. 3, pp. 655–663, Mar. 2007.

[8] C. P. Lai, D. Brunina, and K. Bergman, "Demonstration of 8×40-Gb/s wavelength-striped packet switching in a multi-terabit capacity optical network test-bed," presented at the 23rd Annu. Meeting IEEE Photonics Society, Denver, CO, Nov. 2010, Paper ThQ2.

[9] O. Liboiron-Ladouceur, B. A. Small, and K. Bergman, "Physical layer scalability of WDM optical packet interconnection networks," *J. Lightw. Technol.*, vol. 24, no. 1, pp. 262–270, Jan. 2006.

[10] C. P. Lai, H. Wang, and K. Bergman, "Cross-layer communication with an optical packet switched network via a message injection control interface," *IEEE Photon. Technol. Lett.*, vol. 20, no. 12, pp. 967–969, Jun. 2008.