

Cross-layer reconfigurable optical network: fast failure recovery in testbed for routing algorithms

Cédric Ware*, Caroline P. Lai, Daniel Brunina, Wenjia Zhang**, Ajay S. Garg, Balagangadhar G. Bathula, and Keren Bergman

Department of Electrical Engineering, Columbia University, 500 W. 120th St., New York, NY 10027, USA

** on leave from Institut Télécom, Télécom ParisTech, CNRS LTCI, 46 rue Barrault, 75634 Paris Cedex 13, France*

*** Key Lab of Information Photonics and Optical Commun. (MoE), Beijing University of Posts and Telecom., Beijing, 100876, China*

** cedric.ware@ee.columbia.edu*

ABSTRACT

The challenges awaiting future data networks are not only the sine qua non issue of high bandwidth, but also failure resiliency and power consumption. This can be addressed by a cross-layer physical impairments awareness permeating the network, whose elements can then be reconfigured on the fly so as to optimize transmission capacity, energy efficiency, and recovery time.

We demonstrate an impairment-aware optical networking testbed, capable of quickly recovering from a router failure using optical bypasses. Its architecture aims at validating novel network and routing algorithms designed to minimize a global cost function which will take into account all relevant constraints.

Keywords: Optical communication, optical packet switching, cross-layer networking.

1. INTRODUCTION

Over the past few years, new data-consuming services have been driving forward the global communications networks at a relentless pace. Internet service providers worldwide are even now racing to upgrade their access networks with fiber-to-the-home solutions to meet the ever-increasing expectations of their customers. This includes not only the sine qua non issue of high bandwidth, but also enough flexibility for failure resiliency, while limiting power consumption.

Currently, flexibility is handled in large part by the independent routing of each IP packet: broken links can be routed around seamlessly. However, this still implies that every single packet must be detected and processed by all intermediate routers in electronic form. This leads to a large number of optical-electrical (O-E) converters in the network.

Such a multiplication of required components leads to an unsustainable increase in power consumption [1, 2]. Indeed, we are now at the point where power and heat dissipation are major cost factors in data centers and network equipment. While that problem might not be unsolvable purely in the electronic domain [3], considerable research is being devoted to keeping transmitted data in optical form for as long as possible [4]: either by establishing dynamic lightpaths bypassing routers, thus partially reverting to a circuit-switching paradigm; or by routing packets directly in the optical domain via all-optical burst switching or even packet switching [5, 6]; or both [7, 8].

The implications of such optical techniques are twofold: first, they lose the advantage of signal regeneration implicitly provided by O-E conversion, making transmission more sensitive to physical impairments. Second, circuit-switching requires a fast reactivity to equipment failures, not relying solely on independent packet routing.

We present a network node structure based on an optical packet switch capable of on-the-fly failure recovery using optical bypasses, coupled with a conventional IP router, and optical performance monitors linked to a local control plane. We demonstrate reconfiguration of the optical switching matrix in a router-failure scenario, as well as an adaptation of the application layer to better tolerate physical impairments.

2. PROPOSED NETWORK ARCHITECTURE

We focus on a network built from versatile nodes, capable of handling both IP routing and optical switching, the large bandwidth of the latter to be used for boosting the total capacity per node far beyond the capabilities of a conventional router, albeit with less flexibility.

Several optical switching solutions are available, ranging from the most mature spatial and/or wavelength switching [9] to optical burst or packet switching [5]. IP routing remains in the electrical domain, requiring that at least two ports of the switching matrix (one for input, one for output) be linked to a conventional router via O-E converters.

This work was supported in part by the National Science Foundation Engineering Research Center on Integrated Access Networks (CIAN) (subaward Y503160).

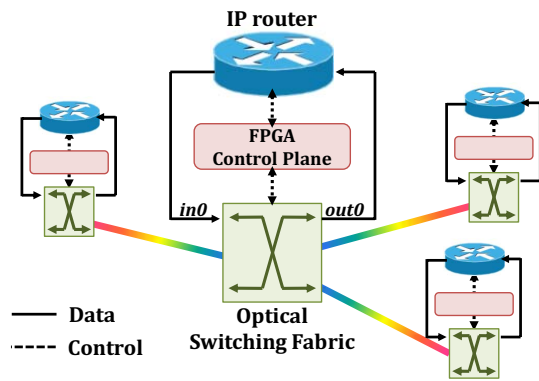


Figure 1. Proposed network architecture, with several nodes combining IP routers, optical switching fabrics, and FPGA-based local control planes.




Figure 2. Photographs of FPGA and SOA circuit boards implementing the switching fabric in this demonstration.

The node architecture we consider is based on our wavelength-striped optical packet switch [10]: each packet is routed by a 4×4 SOA matrix (4 2×2 blocks) according to its destination address, which is encoded as specific header wavelengths, distinct from the wavelengths of the WDM payload. Each payload wavelength is NRZ-OOK modulated at 10 Gbps, providing for flexibility in the total bit rate; while lower-speed header signals remain “on” or “off” for the duration of the packet, prompting control logic to configure the SOA matrix so that the packet is sent to the output corresponding to the packet’s destination address.

This switching architecture is therefore well-suited to failure recovery: a fast programmable gate array (FPGA) can serve as a local control plane, receiving input from monitoring hardware [11], and altering the address/output correspondence in reaction to hardware or link failure. Specifically:

- a failure of the local IP router could be alleviated by redirecting traffic to another node;
- optical performance monitoring systems, detecting that physical impairments exceeded a threshold, would signal the control plane which could then redirect traffic to another link.

By having precalculated configurations of the switching matrix stored in the local control plane, recovery would be immediate, even though the new configuration might not be optimal in terms of network performance. Recalculating a globally-optimized switching configuration would be performed on a longer timescale by the overall control plane.

3. RECONFIGURATION EXPERIMENT

The failure recovery scheme is demonstrated experimentally: the packet switching setup uses 120-ns long 10×10 Gbps optical packets, resulting in 1500-byte messages, analogous to the traditional Ethernet maximum transmission unit length.

As in figures 1 and 2, a FPGA realizes a control plane for the optical fabric that accepts external inputs (e.g. from a router) and then generates failure signals for the switching nodes. The routing logic within the fabric is adapted to accept these electronic failure signals to either route normally (if the router is online, optical packets should be switched accordingly), or route around the failure (if the router is offline or failed, packets are rerouted to ensure that no messages are to be transmitted to the router).

Figure 3 shows the oscilloscope traces of packets incoming from 3 input ports, correctly switched according to the destination address encoded in their header wavelengths, both in normal mode and degraded mode, where all traffic formerly intended for out0 (i.e. packets E and F) is rerouted to out1 (if the port is available; the logic prioritizes messages originally designated for the next node, and we see in Fig. 3(c) that packets C and F contend for out1, thus F is dropped). Figure 3(d) shows the eye diagrams on one of the payload wavelengths before and after the optical fabric.

This experiment is more completely described in [12].

4. ADAPTATION OF APPLICATION LAYER

In order to demonstrate a cross-layer adaptability of the application layer to the physical layer, we set up a variable-bit-rate video transmission over the switching fabric.

Two computers were linked through our optical packet switch using a pair of 10GE Ethernet interfaces customized to send traffic over our optical switching matrix, effecting a two-host private IP network. Host “Test1” displayed video originating from a Webcam physically connected to “Test2”, encoded using software based on FFmpeg [13], and streamed over the network in the form of UDP packets.

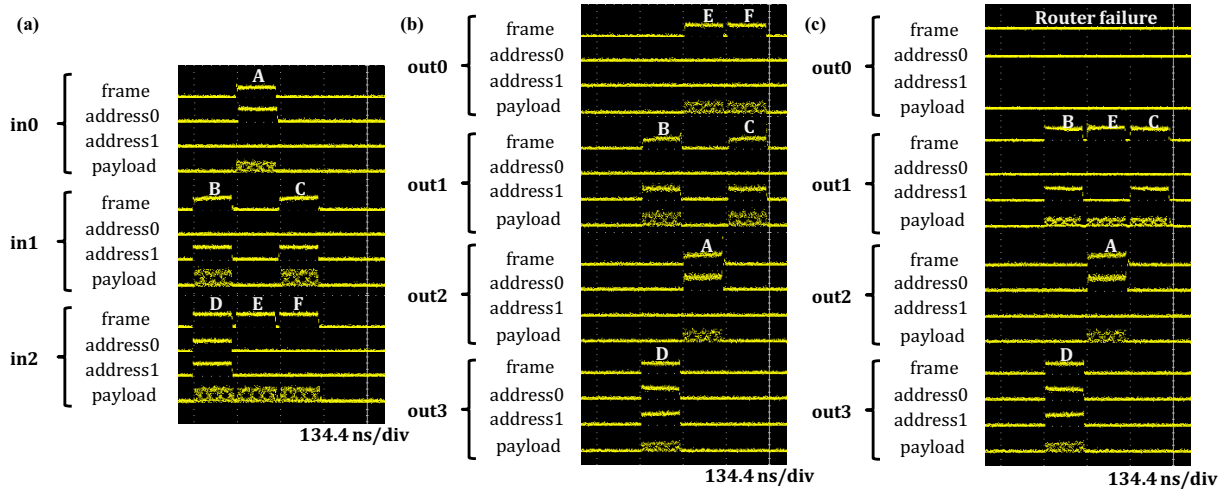


Figure 3. Experimental traces, on a sampling oscilloscope, of packets going through the optical switching fabric: (a) fabric input; (b) fabric output, all output ports functional; (c) fabric output, with link out0 disabled in a router or link failure scenario. Additionally, (d) shows the eye diagram of the 1551 nm payload wavelength at input (above) and output (below) of the switching fabric.

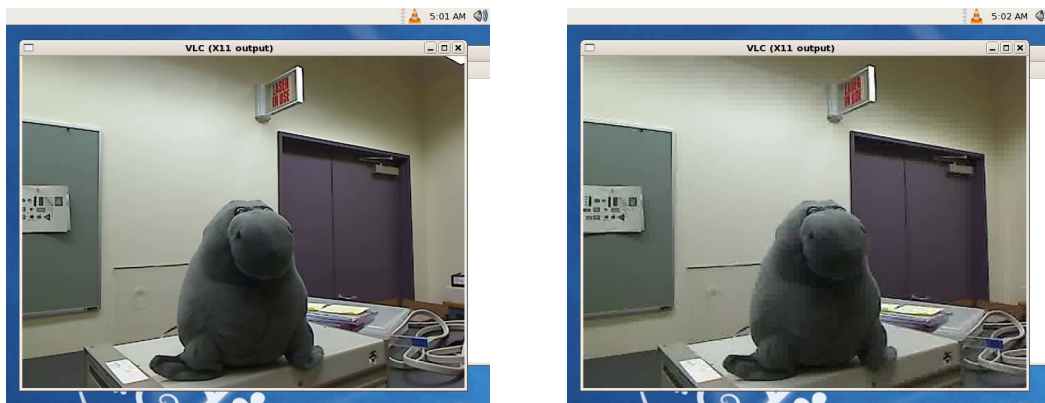



Figure 4. Screenshots of video display: high bit rate (left) and degraded bit rate (right).

The video encoder was customized so that the codec parameters could be changed on-the-fly, and the setup did switch between high and degraded bit rates upon receiving signalling commands embedded in specific UDP packets, which were sent from host “Test1” and could have been carried out of band to another network interface of “Test2”’s.

In this experiment, the signalling was manual: the control UDP packets were sent by user command. In a real-life network, physical impairment monitors would detect QoT drops and/or increases in BER on a link, signal the control plane, which would then instruct the transponders to reduce the link’s bit rate for better impairment resiliency, and inform the applications of that change to give them a chance to cope with reduced resources.

5. CONCLUSIONS

We have proposed a network node architecture consisting of a conventional IP router, an optical packet switch, and a local control plane capable of reconfiguring the switch on the fly in case of failure. The reconfiguration is demonstrated experimentally, in a router-failure scenario. We also showed the adaptation of a video-transmission application to better cope with the bandwidth reduction incurred by physical impairments.

This setup is ultimately destined to become an experimental testbed for novel routing and network algorithms suited to a global management of physical impairments and energy efficiency for future sustainable high-performance networks.

REFERENCES

- [1] D. Kilper, "Energy efficient networks (tutorial)," in *Optical Fiber Conference*, no. OW15, Los Angeles, CA, Mar. 2011.
- [2] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Energy efficiency in telecom optical networks," *IEEE Communications Surveys and Tutorials*, vol. 12, no. 4, pp. 441–458, fourth quarter 2010.
- [3] R. S. Tucker, "Optical packet switching meets Mythbusters," in *Optical Fiber Conference*, no. OTuG4, Los Angeles, CA, Mar. 2011.
- [4] J. Berthold, A. A. M. Saleh, L. Blair, and J. M. Simmons, "Optical networking: Past, present, and future," *IEEE/OSA Journal of Lightwave Technology*, vol. 26, no. 9, pp. 1104–1118, May 2008.
- [5] S. J. B. Yoo, "Energy efficiency in the future internet: The role of optical packet switching and optical-label switching," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 17, no. 2, pp. 406–418, Mar.–Apr. 2011.
- [6] H. Brahmi, M. Bougioukos, M. Menif, A. Maziotis, C. Stamatiadis, C. Kouloumentas, D. Apostolopoulos, H. Avramopoulos, and D. Erasme, "Experimental demonstration of an all-optical packet forwarding gate based on a single SOA-MZI at 40 Gb/s," in *Optical Fiber Conference*, no. OMK5, Los Angeles, CA, Mar. 2011.
- [7] H. Wang, A. S. Garg, K. Bergman, and M. Glick, "Design and demonstration of an all-optical hybrid packet and circuit switched network platform for next generation data centers," in *Optical Fiber Conference*, no. OTuP3, San Diego, CA, Mar. 2010.
- [8] H. Furukawa, N. Wada, Y. Awaji, T. Miyazawa, H. Iiduka, N. Shiga, N. Sato, and H. Harai, "Optical packet and circuit simultaneous transmission technologies for dynamic lightpath setup/release and packet traffic change," in *Optical Fiber Conference*, no. OMK2, Los Angeles, CA, Mar. 2011.
- [9] <http://www.ciena.com/products/6500/>.
- [10] C. P. Lai, D. Brunina, and K. Bergman, "Demonstration of 8x40-Gb/s wavelength-striped packet switching in a multi-terabit capacity optical network test-bed," in *IEEE Annual Photonics Society Meeting*, no. ThQ2, Denver, CO, Nov. 2010.
- [11] C. P. Lai, M. S. Wang, A. S. Garg, K. Bergman, J.-Y. Yang, M. R. Chitgarha, and A. E. Willner, "Demonstration of QoS-aware packet protection via cross-layer OSNR signaling," in *Optical Fiber Conference*, no. OTuM2, San Diego, CA, Mar. 2010.
- [12] C. P. Lai, D. Brunina, C. Ware, B. G. Bathula, and K. Bergman, "Demonstration of cross-layer failure recovery for reconfigurable optical switching fabrics," 2011, submitted to *IEEE Photonics Technology Letters*.
- [13] <http://www.ffmpeg.org/>.