Experimental Demonstration of QoS-Aware Cross-Layer Packet Protection Switching

Caroline P. Lai, Franz Fidler, and Keren Bergman

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

Abstract *A cross-layer based, quality-of-service aware packet protection mechanism is experimentally demonstrated on an optical switching test-bed with 8×10-Gb/s wavelength-striped payloads. The scheme leverages optical performance monitoring to reroute degraded data streams before packets are lost.*

Introduction

Facing today's explosive growth of bandwidth demand, a clean-slate architectural design of the network protocol stack is essential for successfully developing the next-generation internet and its network routing applications[1,2]. To avoid excessive capacity overprovisioning while accommodating flexible bandwidth allocation and dynamic network routing and protection capabilities, the ultimate design must endeavour to engage emerging physical layer technologies in a cross-layer optimized fashion[3,4]. An integrated cross-layer communications networking environment would enable flexible broadband information flow across the otherwise rigidly layered network architecture (Fig. 1). Introspective access to the physical layer can then be used to extract optical performance monitoring (OPM) parameters to achieve significant network performance gains.

Physical layer performance may be accessed by performance monitors directly embedded in the networking equipment, e.g. forward error correction (FEC) modules with bit-error rate (BER) read-out[5], or by dedicated OPM devices[6], measuring optical signal noise ratio (OSNR), packet loss, or link failure. Bidirectional information exchange is then realized to use the measurements in a cross-layer way to reconfigure and optimize packet routing.

Within our scope of cross-layer optimization, a proactive packet protection switching mechanism[5] may be implemented whereby a degradation of high-priority packets is proactively detected at the receiving end using the OPMs' read-out. A control signal is then sent to the transmitting node to allow the data stream to be switched and rerouted to an alternate protection path to ensure no packet loss. Thus, packet routing decisions can be optimized dynamically with respect to varying quality-of-service (QoS) requirements and optical substrate signal degradation on a packet-by-packet basis.

We report on an experimental demonstration of the packet protection scheme on an optical packet switching (OPS) fabric test-bed. The proactive protection scheme utilizes a customized receiver in which optical packets are monitored and, depending on their QoS (high/low priority) and BER, rerouted on a packet-by-packet basis. Multi-wavelength optical packet transmission with 8×10-Gb/s wavelength-striped payloads is shown with correct routing, error-free transmission, and no packet loss.

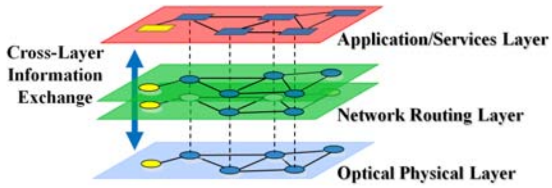

Fig. 1: Cross-layer optimized protocol stack

Switching Fabric Overview

As in Fig. 2, the basic 4×4 switching fabric network test-bed is composed of two OPS network entities placed in parallel, providing increased protection path diversity and broadband packet multicasting capabilities[7]. The design is based on a previously demonstrated SPINet[8] architecture which creates end-to-end transparent lightpaths across the network. The two parallel networks are organized in multistage Banyan topologies and comprised of wideband 2×2 photonic switching nodes. The nodes' electronic control logic is realized in a distributed manner to provide a high level of programmability. The nodes leverage four semiconductor optical amplifier (SOA) gates, offering broadband transmission and data format transparency. The optical packets' wavelength-striped structure[8] contains control information (e.g. frame, priority, address) encoded on dedicated wavelengths, with the payload segmented and modulated at a high data rate (e.g. 10 Gb/s per wavelength) on the rest of the available band. The photonic switching nodes filter and decode the control signals instantaneously on reception of the packets' leading edges using a low-speed optical receiver. Based on the recovered headers, the appropriate SOAs are gated; optical messages are either routed to their desired destination or dropped upon contention. An optical layer acknowledgement (ack) protocol is realized via short optical pulses sent in the reverse direction from the receiving node to notify the source of successful reception. Sources which do not receive acks can retransmit at the next timeslot and/or reroute the data stream to a protection path.

In the experimental realization of our proactive packet protection scheme, contending messages are dropped. Depending on their QoS class, packets with high BER are intentionally discarded after reception, in order to suppress the ack and to trigger rerouting.

978-3-8007-3173-2

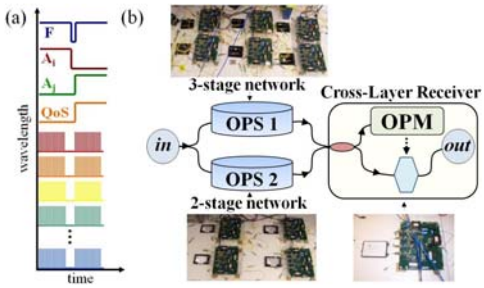

Fig. 2: (a) Packet format; (b) Block diagram and photographs of fabric test-bed with cross-layer scheme

Cross-Layer Packet Protection Scheme

Within our vision of cross-layer optimization, the proposed packet protection mechanism leverages potential signal introspection as offered by a 10-Gb/s FEC. These measurements can feedback to higher network layers to provide packet rerouting along an alternate network protection path. In this way, we can capitalize on the two parallel networks within the fabric test-bed to offer two independent paths between network ports. Proactive protection detects a degrading BER and uses a predefined BER threshold above which packet rerouting is triggered to avoid packet loss. In our experimental demonstration, the loss of a degraded message is mitigated by a cross-layer control signal and subsequent protection path transmission on a parallel network route.

The system uses a modified receiver where optical packets may be monitored, proactively discarded if the signal is degraded, or forwarded to the destination port. The proactive switching mechanism is triggered on per-packet QoS and BER metrics. Data streams with high-priority/high-BER optical messages are rerouted on an alternate protection path, while low-priority (regardless of BER) and high-priority/low-BER messages are forwarded. Low BER denotes a signal quality below the predefined threshold for proactive protection switching, and high BER indicates a quality above the predefined limit.

Experimental Validation

In the test-bed, the electronic logic is synthesized in a high-speed Xilinx complex programmable logic device (CPLD). The system supports 128-ns timeslots with 115.2-ns duration packets with 10-Gb/s data on eight payload wavelengths. The 1152-bit packets are modulated by a $LiNbO_3$ modulator with 2^{15}-1 PRBS.

A pattern of optical packets is injected in both OPS network entities with two differing (high or low priority) QoS classes (Fig. 2). To exemplify the scheme, an additional cross-layer switching node is incorporated in a modified cross-layer receiver design (Fig. 2). In future work, we envision egressing messages to be dynamically and simultaneously monitored by an innovative OPM device (e.g. OSNR monitor[9]), which then signals the cross-layer node to proactively reroute degraded high-priority packets. In this initial cross-layer exploration, a pseudo-BER is generated in lieu of an OPM device, denoting signal quality.

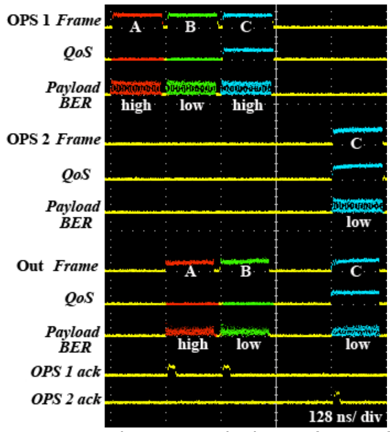

Fig. 3: Input and output optical waveforms of the experimental packet sequence (colours refer to different time-division multiplexed packet streams)

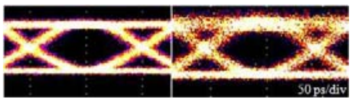

Fig. 4: 10-Gb/s input (left) and output (right) eyes

Fig. 3 validates correct network routing via the optical waveforms, providing packets' frame, QoS priority (high/low), BER (high/low), and payload. The packet sequence demonstrates the functionality of the scheme, whereby proactively protected packets (e.g. packet C) are rerouted on an alternate path within a parallel OPS network. The method is experimentally verified for packets from in1 to out0, attesting to the feasibility of implementing the cross-layer scheme for one output port. We confirm no packet loss and error-free transmission of all received packets; BERs<10^{-12} are verified for all eight payload wavelengths. Fig. 4 provides the 10-Gb/s input and output eye diagrams.

Conclusions

A proactive packet protection switching scheme is demonstrated for an optical switching fabric test-bed. Wavelength-striped optical messages are switched and rerouted on a protection path based on their QoS and signal degradation deduced from a modified monitoring receiver. The work represents a significant step in realizing advanced cross-layer network control based on emerging physical layer OPM devices and varying QoS protocols for next-generation networks.

We acknowledge support from the NSF Engineering Research Center on Integrated Access Networks (CIAN) (subaward Y503160) and NSF FIND (grant CNS-837995).

References

1 GENI, http://www.geni.net
2 CIAN, http://cian-erc.org
3 X.Lin et al., JSAC **24** (8) (2006).
4 C.P.Lai et al., PTL **20** (12) (2008).
5 O.Gerstel et al., OFC, NWD4 (2008).
6 A.E.Willner, OPN **17** (3) (2006).
7 C.P.Lai et al., OFC, OTuA5 (2009).
8 A.Shacham et al., IEEE Micro **27** (4) (2007).
9 Y.K.Lize et al., OFC, OThN2 (2007).