QoS-Aware Cross-Layer Multicasting for Optical Packet-Switched Networks: Simulation Exploration and Test-Bed Demonstration

Caroline P. Lai[(1)], Balagangadhar G. Bathula[(2)], Vinod M. Vokkarane[(2)], and Keren Bergman[(1)]

[(1)] Dept. of Electrical Engineering, Columbia University, New York, NY; caroline@ee.columbia.edu
[(2)] Dept. of Computer and Information Science, University of Massachusetts, Dartmouth, MA

Abstract *Cross-layer quality-of-service-aware packet multicasting is investigated for optical packet-switching network fabrics. We present both a numerical simulation exploration of the cross-layer routing algorithms and an experimental demonstration on an optical switching test-bed with 10×10-Gb/s wavelength-striped packets.*

Introduction

A novel Internet architecture will be essential to accommodate the exploding bandwidth demands faced by the current infrastructure. The next-generation design should leverage innovative optical technologies to offer a more intelligent, programmable optical layer with flexible bandwidth allocation and dynamic interaction with higher network layers[1]. We envision an integrated platform for optical cross-layer (OCL) network communication and control (Fig. 1). OCL enhanced designs will facilitate the extraction of optical performance monitoring (OPM) measurements directly from the optical layer to optimize performance[2]. These OCL routing protocols must also invoke quality-of-service (QoS) classes on the optical layer. Ultimately, the OCL-optimized algorithms must provision for the data's QoS as well as for the physical-layer performance and impairments[2-5].

The future Internet should also engage emerging physical-layer technologies, such as optical packet switching (OPS)[6]. OPS networks comprise a favourable technology approach to enable the flexible high-bandwidth, low-latency interconnections required by future Internet applications. OPS fabrics may be deployed within optical network routers to support high-bandwidth multi-wavelength packet streams between line cards. Additionally, OPS fabrics may achieve a high level of programmability to transparently route wavelength-division-multiplexed (WDM) packets entirely in the optical domain. A significant application that may leverage the greater functionality and programmable flexibility is broadband packet multicasting. We define packet multicasting as the ability to simultaneously transmit broadband multi-wavelength optical messages from a single source to multiple output destinations[7]. Multicasting may be advantageous in high-bandwidth applications, such as networked gaming and real-time diagnostic telemedicine.

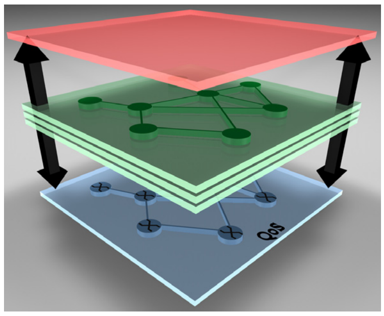

Fig. 1: Cross-layer-optimized stack, indicating the bidirectional information flow between the application (top), network and routing (middle), and optical layers (bottom) enhanced to provide QoS guarantees.

Broadband QoS-based packet multicasting constitutes an important functionality for future OPS networks. Notably, for bandwidth and latency sensitive applications, such as real-time collaboration, high-QoS packet transmission may be leveraged to provide a high-quality communication link. Here, we explore an OCL-enabled platform whereby a packet multicasting operation is realized accounting for both the message's QoS and physical-layer degradation. The concept of cross-layer QoS-aware multicasting is investigated both in simulation and with a test-bed demonstration. We first provide a simulation-based comparative analysis between shortest distance and minimum hop routing algorithms using the NSF network. We then experimentally demonstrate the OPS fabric within one NSF node, validating the error-free operation of cross-layer QoS-based multicasting with bit-error rates (BERs) less than 10^{-12} and a power penalty of 2 dB.

Simulation Validation

The proposed OCL algorithms for QoS-aware packet multicasting are first investigated in simulation. One-way signaling is used to reduce the end-to-end packet transmission latency. The 14-node NSF network topology (Fig. 2) is

978-1-4244-8534-5/10/$26.00 ©2010 IEEE

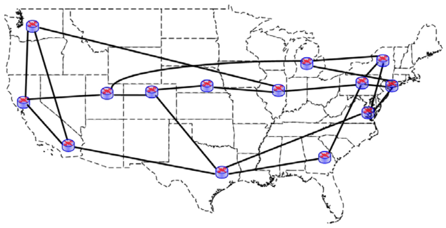

Fig. 2: NSF topology with bidirectional links between the nodes, each carrying 10×10-Gb/s packets.

Tab. 1: Simulation Parameters.

Parameter	Value
Number of Packets	10^6
BER	10^{-9}
Latency	1 ms
Input Optical Power	-10 dBm
Inline Amplifier Gain	14 dB
Switch Crosstalk Ratio	25 dB
Starting Wavelength	1537.4 nm
Wavelength Spacing	2.8 nm

numerically simulated using a global control plane to track each node's QoS performance. A centralized routing and wavelength assignment (RWA) scheme is realized. Packets are assumed wavelength-striped, using ten wavelength channels each at 10 Gb/s.

Packets are simulated as discrete events[3]. The packets follow a Poisson arrival rate and depart with exponential service times. Upon an arrival event, each packet is assigned to a request and then routed based on a minimum distance routing (MDR) or a minimum hop routing (MHR) algorithm. The necessary QoS parameters are retrieved from the application layer. Optical packets reaching the destination ensure that the threshold requirements imposed by the application layer are met[8]. The QoS is embedded in the control signal and is updated as the packet propagates through the network. The QoS parameters consist of its BER, latency, priority, and the reliability of the link. At each node, the QoS of the routed packet is computed online and compared with the threshold requirement of the application. If the QoS parameters are violated, the packet is dropped or rerouted on an alternate path if available. Multicasting is initiated as required.

An intelligent, efficient control plane acts as a middleware between the application and optical layers[8]. Based on the control plane decision, the optical packet is routed on the link.

Using the parameters in Tab. 1, the BER is estimated based on the optical-signal-to-noise-ratio (OSNR). Since the BER is a nonlinear function, we compute the link's noise factor. The overall noise factor of the lightpath is computed as a product of the individual noise factors of the links[8]. The overall latency of the packet is the sum of the individual latencies of the links. The reliability of the switch is based on the downtime and path restoration time. The priority parameter enables possible packet routing on alternate network paths.

The performance of the proposed QoS-aware cross-layer multicasting is simulated using the NSF network with the distances scaled down by a factor of ten, due to the lack of regenerators at the node's switching fabrics. In Fig. 3, we compare the performance of the NSF topology in terms of packet loss, average latency of the packet, hop count, and execution time for the routing algorithm.

The x-axis for all the plots in Fig. 3 is the offered network load in Erlang, defined as the ratio of the arrival rate to the departure rate. In Fig. 3(a), we observe that MHR offers lower loss compared to MDR at low network loads. This indicates that packets routed based on hop count show a higher probability of successfully guaranteeing the QoS imposed by the application layer. The average latency (Fig. 3(b)) of MHR is higher than MDR; this may not be problematic if the latency threshold is still satisfied. Thus, the routing layer can adopt a hop-routing at lower network loads. As the load increases, the packet loss for both algorithms converges (Fig. 3(a)). In order to optimize performance, the application layer can instruct the routing layer to switch to distance routing at higher network loads. Thus, cross-layer communication helps to achieve design trade-offs and provide the necessary QoS. We also compare the average hop count for the two routing methods. It is evident that the hop count for the MHR is lower than MDR (Fig. 3(c)). A decrease at higher loads indicates that providing QoS for optical packets that traverse longer hops is more problematic. Fig. 3(d) shows the execution time (in hours) required for the simulations using a 2.33-GHz Quad Core Xeon processor with Hyper-Threading and 8-GB RAM.

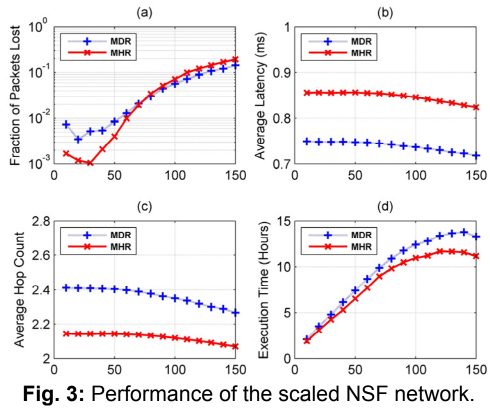

Fig. 3: Performance of the scaled NSF network.

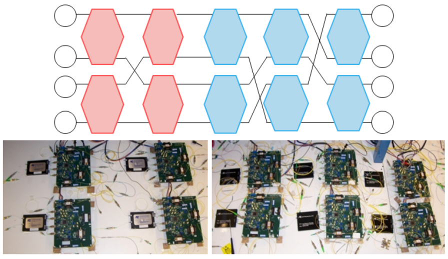

Fig. 4: Experimentally implemented multicast-capable fabric architecture and test-bed photograph.

Experimental Demonstration

The QoS-based broadband packet multicasting operation is experimentally demonstrated on a multicast-capable OPS fabric test-bed[7] (Fig.4). The fabric test-bed represents the optical switching fabric deployed within one node of the NSF network. The multistage test-bed is implemented with 2×2 photonic switches, which use semiconductor optical amplifiers (SOAs). Wavelength-striped packets are supported, with control information (e.g. frame, address, QoS) encoded on a subset of wavelengths and the payload segmented and modulated at a high data rate (e.g. 10 Gb/s) on the rest of the band. The 2×2 switches detect the control information at the packet's rising edge using filters and receivers. The packet's header bits are processed electronically at each stage. The routing control logic gates the correct SOAs to provide the desired routing. No optical buffers are used. The multicast-capable fabric[7] is realized with a multistage design, using differing packet routing (PR) and packet multicasting (PM) stages. The stages have distinct control logic that depends on the recovered header bits.

An SOA-based receiver is realized[4] whereby the real-time performance of optical packets can be monitored. Switching is triggered on the per-packet QoS and signal degradation (here, BER). Low-QoS/high-BER packets are detected by the cross-layer receiver and rerouted on an alternate path. The pseudo-BER signal is generated offline in place of an OPM, though a real-time packet OSNR monitor may be used[2].

The QoS-aware packet multicasting is validated on the 4×4 optical fabric test-bed with two PR and three PM stages. The 10×10-Gb/s wavelength-striped packets are 120-ns long, analogous to the Ethernet MTU. The 1500-B packets are modulated by a $LiNbO_3$ modulator with 2^7-1 PRBS; the payload wavelengths range from 1537.4 to 1564.0 nm. Fig. 5 depicts the pattern of optical packets injected in the multicast-capable fabric with two QoS levels (high/low priority). The QoS and packet signal quality are assessed and a real-time decision is made to forward or reroute the message on a protection path. At the output, we verify that error-free QoS-based packet multicasting is achieved. BERs<10^{-12} are obtained on all ten payload wavelengths. BER curves for the system show a 2-dB power penalty (Fig. 6).

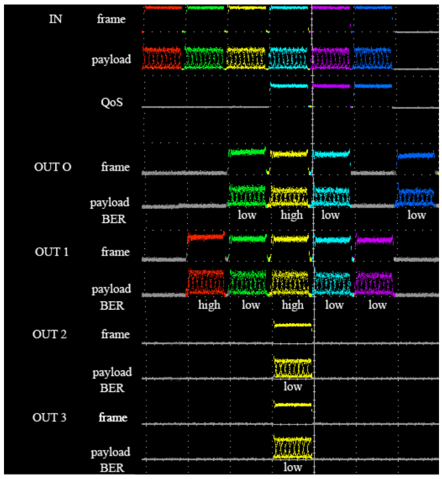

Fig. 5: Optical waveforms corresponding to the QoS-aware packet multicasting operation.

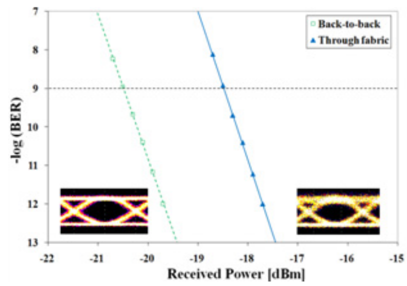

Fig. 6: Sensitivity curves with insets of the 10-Gb/s eye diagrams (input: left, output: right).

Conclusions

Future networks will require a QoS-aware cross-layer protocol stack. This work confirms that broadband packet multicasting can be realized accounting for physical-layer access in a cross-layer-optimized approach. Numerical results and a demonstration on a fabric test-bed show that packet multicasting can be performed based on QoS and signal degradation. This exploration leverages an OCL-optimized platform and novel optical technologies to achieve performance gains for next-generation networks.

We acknowledge support from the CIAN NSF ERC (subaward Y503160); BBN, GENI Project Office (agreement 1631); and the NSF SOON project (grant CNS-0626798).

References

1 CIAN, http://cian-erc.org
2 C.P.Lai et al., Proc. OFC'10, OTuM2 (2010).
3 F.Fidler et al., Proc. ECOC'09, 2.5.2 (2009).
4 C.P.Lai et al., Proc. ECOC'09, 2.5.3 (2009).
5 S.Azodolmolky et al., Proc. ONDM'09 (2009).
6 E.W.M.Wong et al., JLT **27** (14) (2009).
7 C.P.Lai et al., Proc. OFC'10, OWI4 (2010).
8 B.G.Bathula et al., TON **18** (1) (2010).