

Interoperation of Layer-2/3 Modular Switches with 8QAM/16QAM Integrated Coherent Optics over 2000 km Open Line System

Mark Filer¹, Hacene Chaouch², Xiaoxia Wu³, Jamie Gaudette¹, Jeffrey L. Cox¹

¹Microsoft Corp, Redmond, WA; ²Arista Networks, Santa Clara CA, ³Juniper Networks, Sunnyvale, CA
mark.filer@microsoft.com

Abstract: Arista, Cisco, and Juniper's layer-2/3 modular switches with integrated coherent optics are interoperated over 2000 km at 150G 8QAM and 1000 km at 200G 16QAM on Microsoft's open line system.

OCIS codes: (060.2330) Fiber optics communications; (060.1660) Coherent communications; (060.4250) Networks

1. Introduction

Long-haul and metro bandwidth demand continue to increase for large-scale cloud providers, driven by phenomenal growth of cloud services and the globally distributed nature of data centers. To efficiently meet demand, cloud providers have driven the Open Line System (OLS) concept. OLS disaggregates line system components (amplifiers, gain equalizers, ROADMs, etc.) from coherent optical transceivers associated with the line system. Motivation for this disaggregation is multifold, but fundamentally rests on the premise that OLS components are relatively long-lived, technologically stable, and largely undifferentiated parts of an optical transmission system compared to transceivers, which are still rapidly evolving.

Because of this rapid evolution, coherent transceivers have made little progress towards line-side signal interoperability [1]-[3]. Each coherent DSP supplier has improved performance and feature sets with independent efforts and no interoperability standards. Recently, spectral efficiency of the various DSPs has begun to asymptotically approach channel capacities predicted by Shannon and the various DSPs are now realizing ever smaller relative performance gains while simultaneously chasing ever more esoteric features to claim differentiation.

However, cloud providers generally do not require esoteric features. Specifically, long-haul and metro transmissions systems are largely viewed as a means to connect the highest possible bandwidths between clusters of routers or switches in a point-to-point fashion at maximal interface speeds. It is already common for cloud providers to maximally light multiple, entire optical systems in metro areas to interconnect data centers. Accelerating bandwidth demands in the long-haul are driving a need to view long-haul systems in a similar manner. To realize this, simplicity, efficiency, and supply-chain diversity are major focus areas. Emphasis on squeezing the final ten or twenty percent out of the Shannon limit or adding a plethora of obscure features hinders acceleration of mass deployment. Ultimately, the industry must agree on line-side interoperability to create a robust supply chain and accelerate the deployment of capacity.

This paper illustrates a partial step towards transport ecosystem interoperability by demonstrating interoperability over an OLS between packet switch systems provided by three different suppliers, each with their own version of Integrated Coherent Optics (ICO), where the coherent transponders are integrated directly into the switch platform. Across the three suppliers, at least two fundamentally different packet chipsets and two different electro-optic suppliers are shown to interoperate.

2. Configuration

The three packet switching systems were Arista 7500, Cisco NCS5500 and Juniper QFX 10000. Each switch was equipped with industry standard 100G/40G QSFP "grey" cards, as well as an ICO line card providing coherent





Figure 1. (a) Interop system configuration, (b) spectrum of 104 channels spaced at 37.5 GHz, with coherent linecard channels in center.

TABLE 1 – CHANNEL SCHEME FOR INTEROP

freq.	lane 0 [THz]	lane 1 [THz]	West	East
f_1	193.5375	193.6875	Juniper	Arista
f_2	193.5750	193.7250	Cisco	Juniper
f_3	193.6125	193.7625	Arista	Cisco
f_4	193.6500	193.8000	Ref.	Ref.

Figure 2. Optical spectrum in the east to west direction of 8QAM signals after multiplexing into open line system. Each channel was configured for RRC = 0.2, non-differential encoding with a baud rate of approximately 34 Gbaud



DWDM channels. The ICO cards represent different implementations, with one supplier integrating a complete and generally available board-mounted MSA, and the two others using discrete coherent DSP and faceplate-pluggable analog optics (CFP2-ACO). However, all ICOs integrated the same coherent DSP making line-side interoperation possible. All ICOs provide selectable QPSK/8QAM/16QAM modulation with a line-rate of 100G/150G/200G respectively, 25%-overhead SD-FEC and Nyquist pulse shaping. The electro-optics used in this demonstration include both silicon photonic (SiP) and indium phosphide (InP) processing technologies, and all interfaces support grid-less tuning with 6.25 GHz resolution and 37.5 GHz channel spacing at full line-rate.

The test system is shown in Figure 1a. An ICO line card from each switch supplier was placed at both ends of the system (“east” and “west”) and provided 6 bi-directional test channels (2 per supplier). Each ICO port was configured with identical encoding, and with a root-raised cosine (RRC) pulse-shape at a roll-off of 0.2. A generally available stand-alone integrated transceiver was added as a reference (labeled “ref”) [4]. A total of 32 bulk-modulated sources were evenly distributed on either side of the devices under test (DUT), modulated with the same modulation format, baud rate, and pulse-shape as the DUTs. Another 64 shaped-ASE “channels” were added to load the remainder of the C-band. The ASE channels were shaped with a programmable WSS to emulate RRC = 0.2 mQAM. In total, 104 channels spaced 37.5 GHz were propagated for all results reported (Figure 1b). Spectral dips seen between channel groupings are intentionally provisioned 6.25 GHz “dead-bands,” which avoid the impact of WSS rolloff as channel groups are muxed or demuxed at terminal ROADMs. Measured performance is expected to be conservative because 1) no specific measures were taken to time-decorrelate the bulk-modulated channels and 2) ASE channels produce more harmful cross-phase modulation than data-modulated carriers.

The transmission testbed consisted of the same OLS and fiber plant reported in [5]. The OLS features a colorless, directional architecture with full support for alien wavelengths. The mostly NZ-DSF fiber plant spans 2000 km bidirectionally, and is roughly made up of 85% Corning LEAF and 15% Corning SMF-28e LL. The optical ports of the line cards were patched into the OLS colorless multiplexers on each end, along with the bulk-modulated and ASE channels. Interop between the different manufacturer’s switch ports was achieved optically via software, utilizing the colorless aspects of the OLS, by simply tuning the lasers of the coherent ports on the east and west ends of the system to achieve the desired combinations of transmit/receive pairs. The laser frequency settings shown in Table 1, which correspond to the color-coded lines in Figure 1a, demonstrate how this was done for these measurements. This configuration ensures that each switch supplier is transmitting receiving to/from each of the other switch suppliers when considering both east and west paths. Figure 2 details the transmitted spectra from each




Figure 3. (a) 8QAM Q vs launch power into G.655 (G.652 3 dB higher), (b) 8QAM and 16QAM constellations before and after transmission




Figure 4. (a) 48-hr BER logs for 8QAM over 2000 km, (b) 2-hr BER logs for 16QAM over 1000 km

of the line cards in the east to west direction, which correspond to the “east” column of Table 1.

Once the line cards and line system were properly configured, traffic generator test-set ports were connected to 100G QSFP-28 grey ports of the switch residing at the beginning and end of the path. The test-set generated traffic at 100% line rate, i.e., 100 Gb/s UDP traffic with 1000-byte packet size to a fixed IP address. Traffic was weaved through all available 100G client interfaces by setting up virtual routing and forwarding paths (VRFs) and static routes on all three connected switches. A similar traffic flow was set up in the opposite direction which resulted in bidirectional traffic at maximum link utilization.

3. Measurement results and discussion

With all ICOs configured for 8QAM transmission at 150 Gb/s line rate, and layer 3 test-set traffic flowing, a sweep of fiber launch power was performed for all 104 channels. The result (for the east to west direction) is shown in Figure 3a. The x-axis references launch power into G.655 LEAF spans; launch powers into G.652 SMF-28 fibers were 3 dB greater than those shown. At the optimum launch power of -3 dBm/channel into LEAF (0 dBm/channel into SMF-28), the Q-factor was more-than 1dB above the FEC limit, and OSNR margin between 2 and 2.5 dB was measured. The system ran for 48 hours with full-rate test-set traffic at optimal launch power, and the pre-FEC BER was logged during the entire period in both the east and west directions (Figure 4a) for stability analysis. The peak-to-peak Q-factor varied by ~0.1 dB, and over 2.5×10^{12} packets were transmitted during the logging period with 0 packets lost. After the logging period expired, a capture of the transmit and received constellations was performed, as shown on the left half of Figure 3b.

Next, all ICOs were reconfigured for 16QAM at 200 Gb/s line rate. Client traffic VRFs and routes had to be reconfigured to account for the increase from 3 to 4 100G streams on each coherent interface, but optically the interfaces were unchanged. The line system distance was reduced from 2000 km to 1000 km (optimal launch power assumed the same), and once end-to-end traffic was verified, the test set and pre-FEC BER logs were restarted. Between 0.5 – 1.5 dB Q margin was measured (Figure 4b) and no packets were lost over a 2-hour traffic test. Transmit and received constellations are shown in the right half of Figure 3b.

4. Conclusions

Interoperation was demonstrated between three different switch suppliers using ICO and two different electro-optic suppliers. Over primarily NZ-DSF, 8QAM at 2000km and 16QAM at 1000km was demonstrated without trading performance or spectral efficiency for interoperability. The electro-optics used in this demonstration comprise both SiP and InP processing technologies, and the results presented demonstrate the suitability of both platforms for long-haul 8/16QAM transport and interoperation. Constructing global network infrastructures at cloud scales already leverages the extensive interoperability that exists in the switching, routing, and grey optical technologies. Additionally, the OLS is essentially a stand-alone system not requiring interop at this time. The only missing interoperable component in the entire network ecosystem is the coherent DSP that drives the electro-optics, interconnecting switches or routers over OLS.

This work took a tremendous amount of coordination between engineering teams at Microsoft, Arista, Cisco, and Juniper. The authors would like to thank the extended teams at Arista, Cisco, and Juniper for their support.

- [1] H.-J. Schmidtke and R. Marcocchia, “DWDM Optical Line-side Interoperability” Proc. OFC, paper OM3G.7 (2012)
- [2] M. Tomizawa, “100G DWDM Transport Systems: Driving the Technologies and Deployment,” Proc. COIN, paper TuE.2 (2012).
- [3] M. Gunket et al., “Vendor-Interoperable Elastic Optical Interfaces,” J. Opt. Commun. Netw., vol. 7, no. 12, pp.B184-B193 (2015)
- [4] C.R.S. Doerr et al., “O, E, S, C, and L Band Silicon Photonics Coherent Modulator/Receiver,” Proc. OFC, paper Th5C.4 (2016)
- [5] M. Filer et al., “Elastic Optical Networking in the Microsoft Cloud,” J. Opt. Commun. Netw., vol. 8, no. 7, pp.A45-A54 (2016)