

Demonstration and Performance Analysis of 4 Tb/s DWDM Metro-DCI System with 100G PAM4 QSFP28 Modules

Mark Filer¹, Steven Searcy², Yang Fu³, Radhakrishnan Nagarajan³, and Sorin Tibuleac²

¹Microsoft Corporation, One Microsoft Way, Redmond, WA 98052, USA, mark.filer@microsoft.com

²ADVA Optical Networking, 5755 Peachtree Industrial Blvd, Norcross, GA 30092, USA, ssearcy@advaoptical.com

³Inphi Corporation, 2953 Bunker Hill Lane, Suite 300, Santa Clara, CA 95054, USA, rnagarajan@inphi.com

Abstract: We demonstrate a 4-Tb/s metro-DCI system with commercial QSFP28 modules (40x100G dual-wavelength 56-Gb/s PAM4). We detail system performance over 80km and quantify tolerance to chromatic dispersion and nonlinearity over a wide range of fiber types.

OCIS codes: (060.2360) Fiber optics links and subsystems; (250.3140) Integrated optoelectronic circuits.

1. Introduction and motivation

The Data Center Interconnect (DCI) space has become an area of increased focus for traditional DWDM system suppliers over the last few years. The growing bandwidth demands of cloud service providers (CSPs) offering SaaS, PaaS, and IaaS capabilities have in turn driven demand for optical solutions to connect switches and routers at the different tiers of the CSP's internal data center network. Today, this often requires solutions at 100 Gb/s, which inside a given data center, can be met with direct attach copper cabling, active optical cables, or 100G grey optics.

However, in contrast to a more traditional data center model, where all the data center facilities reside in a single campus, many CSPs have converged on distributed regional architectures to be able to scale sufficiently and provide cloud services with high availability. Interconnecting physically disparate data center buildings within the same geographic region poses its own challenge, where the spectral efficiency of grey optics is too low.

In the 100G era, this problem has typically been solved using dense wavelength-division multiplexed (DWDM) coherent QPSK transponders which offer soft-decision FEC enabled performance down to 11 dB OSNR, subsea-capable chromatic dispersion (CD) compensation of 250,000 ps/nm, power efficiencies on the order of 100 Watts per 100G, and capacities of 8 – 9.6 Tb/s per fiber pair. For DCI links that are typically latency-constrained to 80 km or less (<1600 ps/nm), line systems can easily deliver OSNRs in the low 30's. Rack space and power are typically limited and costly in these facilities, making power and space efficiency critical design goals. While fiber is not as abundant in these metro environments as within the data center itself, typically tens of fiber pairs are available at reasonable costs, relaxing ultimate spectral efficiency as a primary design criterion. With these considerations, today's coherent solutions are not an ideal fit for DCI applications.

In response, a low-power, low-footprint direct-detect solution was conceived employing PAM4 modulation format. By utilizing silicon photonics technology, a dual-carrier transceiver featuring a PAM4 ASIC, with integrated digital signal processing (DSP) and forward error correction (FEC), was developed and packaged into a QSFP28 form-factor. The resulting switch-pluggable module enables DWDM transmission over typical DCI links at 4 Tb/s per fiber pair and power consumption of 4.5 Watts per 100G.

While previous work has shown the feasibility of such PAM4 transmission for metro links [1-3], these studies have been based on laboratory systems with one or few channels. This paper seeks to demonstrate the optical performance of a commercial system when used in typical DCI deployment scenarios. The system is fully-loaded with PAM4 QSFP28 modules plugged directly into a layer-2/3 switch and paired with an appropriate line system operated in an open fashion. First, descriptions of the PAM4 module and line system are given. Next, aspects of module and line system optical performance are presented, with particular attention given to the need for appropriate dispersion compensation. Finally, a performance summary over multiple ITU-T G.652 – G.655 fiber types is presented to account for the diversity of fiber types which may be encountered in metro environments.

2. Description of the PAM4 modules

The PAM4 QSFP28 module is based on a highly integrated silicon photonics optical chip. The chip contains a pair of integrated traveling wave silicon Mach Zehnder modulators (MZM) in the output optical path, and a pair of high speed Ge photodetectors (PD) in the receive path. Both the MZM and the PD have small signal bandwidths in excess of 25 GHz. The transmit signals are combined using a 2:1 multiplexer and a similar structure is used for the demultiplexer. The receive path has polarization diversity. This is accomplished by using a low loss polarization beam splitter (PBS). Both the MZM driver amplifier and the PD transimpedance amplifier (TIA) are wire bonded to the silicon photonics chip. The DFB lasers are external, and edge coupled to the silicon photonics chip.

The module incorporates a real time PAM4 ASIC [4,5] which also provides the CAUI 4 host interface to the switch. The 4x25 Gb/s input data are converted into 2x28Gbaud PAM4 streams encoded with an Inphi proprietary FEC (IFEC) which has a coding gain of 10.5 dB. On the receive side, the output of the TIA is sampled in the PAM ASIC using 28 GS/s ADCs. The receive DSP has FFE (feed forward equalizer) and DFE (decision feedback equalizer) options to recover the PAM4 signal. The equalizer taps are automatically adapted using a least-mean-squared (LMS) algorithm. The IFEC decoder then corrects the errors and generates the original Ethernet data.

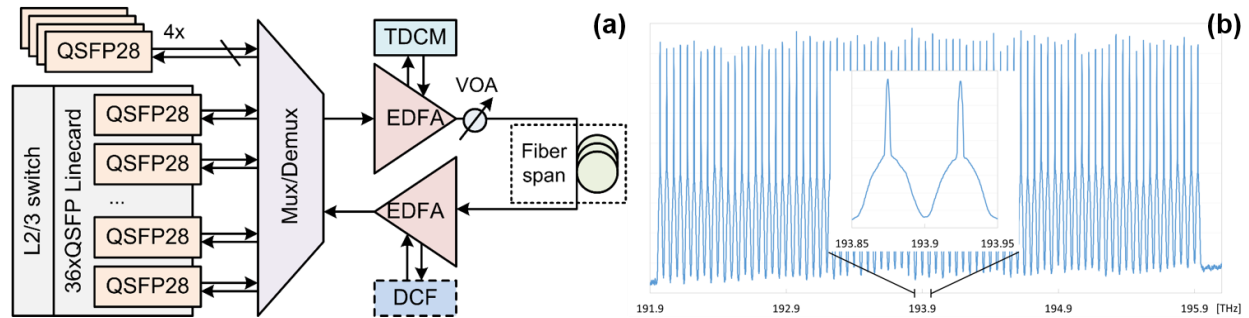


Fig. 1. (a) Experimental setup with PAM4 modules and line system; (b) Measured receive spectrum for system fully loaded with 4-Tb/s PAM4.

3. Experimental conditions

The experimental link, shown in Fig. 1(a), consisted of a single fiber span with ADVA line system, comprised of a 100-GHz arrayed-waveguide grating mux/demux, pair of EDFAs, etalon-based tunable dispersion compensation module (TDCM), and optional dispersion compensation fiber (DCF). The system was fully loaded with 40 Inphi QSFP28 transceivers; 36 modules (192.2-195.7 THz) ran on a single linecard in an Arista 7500E series L2/L3 switch, and 4 modules (192.0-192.1, 195.8-195.9 THz) ran on independent evaluation boards. The received optical spectrum with all 40 modules (80λ) is shown in Fig. 2(b). To evaluate performance, channels on the Arista linecard were simultaneously monitored for pre-FEC BER, SNR, and received optical power. In all results below, Q^2 factor is calculated from the reported pre-FEC BER. An EXFO FTB-5240BP optical spectrum analyzer (OSA) measured the received OSNR per channel (via pre-amp output tap) using the polarization-nulling method.

The system performance was tested back-to-back (no fiber) and with various fiber spans. One link consisted of 80 km (16 dB) Corning SMF-28, with 60 km DCF. Another link consisted of 40 km SMF-28 followed by a variable optical attenuator (VOA), with 40 km DCF. Tolerance to fiber nonlinearity was tested over five different fiber types: OFS AllWave, TrueWave-RS, and TrueWave-Classic, and Corning LEAF and MetroCor. These spans ranged from 50-60 km with 11-13 dB loss (including connectors), such that the average received OSNR was in the same range (~34-35 dB, depending on launch power). For AllWave, 40 km DCF was used, while for the other types no DCF was used and the TDCM performed all compensation. For each test, the TDCM setting and amplifier tilt compensation were optimized. Post-amp gain was fixed and launch power adjusted using a VOA built into the amp. Pre-amp gain was adjusted based on launch power and span loss to achieve average receive power of +2 dBm/λ.




Fig. 2. (a) Tolerance to residual CD; Distribution of per-channel (b) Q^2 factors and (c) receive OSNR values on 80 km SMF link.

4. Results and discussion

The tolerance of the transceivers to residual CD was tested by scanning the TDCM setting in a back-to-back condition with no fiber (network loopback with 10 dB loss). The average Q^2 factor penalty across all monitored channels is shown in Fig. 2(a). The average penalty is <1 dBQ for ± 100 ps/nm TDCM setting (with worst-case penalty <2 dBQ). It should be noted that these results include small amounts of residual CD from other optical components, as well as small inaccuracies and variations in CD across channels covering the full C-band.

A long-term validation test was performed over an 80 km SMF link. The pre-FEC BER was logged for 15 hours with all channels running post-FEC error-free for the duration. The distribution of average Q^2 factor for each monitored channel is shown in Fig. 2(b). Each channel has 1.5 to 4 dB Q^2 margin to the practical FEC threshold of ~ 8 dBQ (BER = $6E-3$). The corresponding distribution of per-channel OSNR values is shown in Fig. 2(c).

A series of tests was performed over five widely-deployed fiber types to quantify the impact of nonlinearity. These results are shown in Fig. 3(a), where the average Q^2 penalty across all monitored channels is plotted versus the average launch power per wavelength into the fiber span. There is a small increase of ~ 1 dB in received OSNR when moving from lowest to highest launch powers. For AllWave, the impact from fiber nonlinearity is negligible up to the maximum launch power of $+4$ dBm/ λ . This nonlinear tolerance is higher than previously reported for similar systems, despite the significantly larger number of co-propagating channels [3]. The other NZ-DSF types all exhibit some nonlinear penalty at higher launch powers. LEAF, MetroCor, and TrueWave-RS all show similar performance, with nonlinear impact becoming significant starting at launch power $+2$ dBm/ λ . TrueWave-Classic shows the strongest nonlinearity due to its small effective mode area and very low fiber dispersion. Significant impact starts at launch power 0 dBm/ λ , yet even at this condition, all channels had >1 dB Q^2 margin. Long-term tests with LEAF (15 hours) and TrueWave-RS (64 hours) confirmed stable, error-free performance over NZ-DSF.

Finally, a test was performed to determine the system performance with increasing span loss. The span consisted of 40 km SMF followed by VOA, with launch power of $+4$ dBm/ λ into the fiber. The results are summarized in Fig. 3(b), where average Q^2 and OSNR are shown for span loss ranging from 9 dB up to 22 dB, which was the maximum achievable span loss in these tests. The Q^2 factor begins to degrade significantly above 18 dB span loss, as shown in Fig. 3(b), due to the lower OSNR and the gain limit of the amplifier producing a drop in received power. At 18 dB span loss, all channels still have >1 dB Q^2 margin. At the max span loss of 22 dB, the worst-case channels have zero margin, operating around the practical FEC limit (though all channels were error-free in the short-term test). This highlights the robust performance of the system, which can accommodate typical maximum span loss for DCI links. Longer links could also be supported by using distributed Raman amplification or further system optimizations.

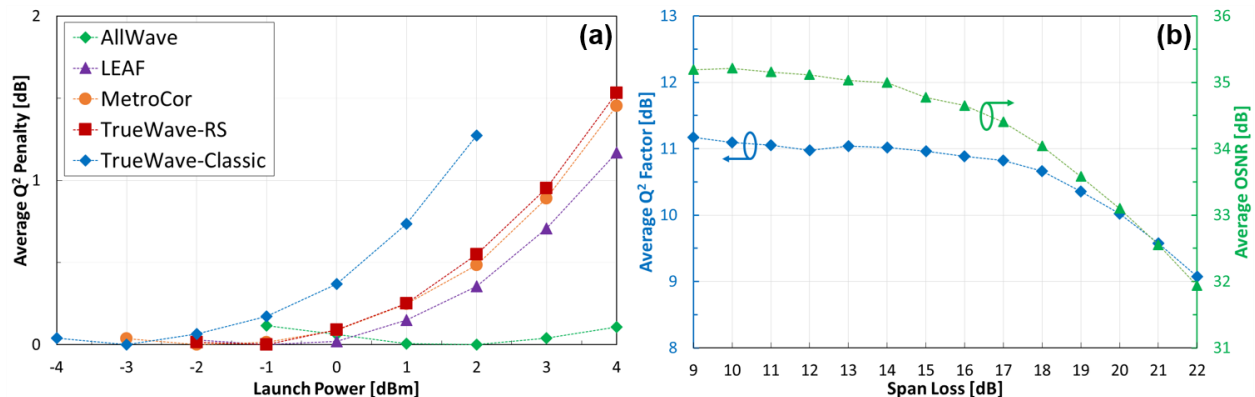


Fig. 3. (a) Nonlinear tolerance over five commonly deployed fiber types; (b) Q^2 and OSNR results for increasing span loss (40 km SMF + VOA).

5. Conclusion

We have presented the first demonstration of a fully-populated system with commercial transceivers and line system which can support 4-Tb/s capacity over 80 km using PAM4 modulation. Test results show the tolerance of the system to key impairments, including chromatic dispersion and fiber nonlinearity, and show that this type of system is robust to support large-scale deployment in DCI links.

6. References

- [1] S. Yin, et al., IEEE PTL 27(24), 2531-2534 (2015).
- [2] J. Lee, et al., ECOC'16, M.2.D.3 (2016).
- [3] N. Eiselt, et al., OFC'16, W1K.5 (2016).
- [4] K. Gopalakrishnan, et al., Proc. ISSCC, 3.4 (2016).
- [5] F. Chang, et al., OFC'16, Th1G.2 (2016).