

Lessons Learned from CFP2-ACO System Integrations, Interoperability Testing and Deployments

Hacene Chaouch¹, Mark Filer², and Andreas Bechtolsheim¹

¹Arista Networks, Santa Clara CA, ²Microsoft Corp, Redmond, WA
 hacene@arista.com; mark.filer@microsoft.com

Abstract: We discuss the key metrics of analog coherent interfaces for today's 200G 16QAM and future 400-600G 64QAM pluggable systems. A cloud service provider perspective on next generation DCI requirements is also discussed.

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

1. Introduction

There is a tremendous demand in the market from hyper-scale cloud data centers to increase datacenter interconnection (DCI) bandwidth in the most spectral efficient, power efficient, and economical fashion. The arrival of merchant silicon for DWDM DSPs in conjunction with open line systems (OLS) has enabled switch-router system vendors to build high-density DWDM line cards with integrated DSPs and pluggable analog coherent optics (ACO) modules that offer long-haul transmission capability with significantly improved density and power efficiency compared to traditional transport solutions. This paper discusses the experiences learned from building such a line card.

2. The modular coherent switch

A coherent line card for a modular L2/L3 switch-router is shown in Figure. 1. On the line side, four embedded merchant silicon DSP chips feed eight adjacent CFP2-ACO ports. A total transport capacity of 1.6 Tbps / RU is enabled with 200G PM-16QAM. Additionally, the card offers 100GbE wire speed MACSec encryption as well as layer 2/3 switching capability with deep buffer architecture.



Figure 1. Arista 1.6 Tbps coherent line card

3. The analog interface key metrics

A. IQ skew and Nyquist pulses

The time delay mismatch between the in-phase and quadrature tributaries in each polarization is the cause of an OSNR sensitivity reduction at the receiver. While 100G QPSK is relatively tolerant, the OSNR penalty increases rapidly with 8QAM and 16QAM as shown in Figure 2(a), where we varied the absolute IQ delay in H and V polarizations within ± 5 ps in 500fs steps. It can be seen that an IQ time error of 3 ps, in both polarizations, for QPSK leads to a ~ 0.15 dB OSNR penalty at FEC limit ($\sim 3.2 \times 10^{-2}$), but in 16QAM the same skew leads to ~ 0.4 dB penalty. Color bands in Figure 2(a) correspond to 0.1 dB equi-penalty regions. Notice that the intrinsic IQ skew of this ACO was about $(-1.7, -0.7)$ ps in H and V respectively.

More important is the effect of a large (> 3 ps) IQ timing error when Nyquist shaped 0.2 RRC 16QAM signals are used. Figures 2(b-c), with 0.2 dB color band regions, indicate that a ~ 2 dB OSNR penalty at FEC threshold is measured experimentally when a ± 5 ps residual skew is present in H and V simultaneously. This is due to the significantly reduced horizontal eye opening of 0.2 RRC PAM4 signals as show in Figure 2(d). The current migration towards 37.5 GHz-spaced 8/16QAM signals in OSNR-limited terrestrial and subsea networks [1] necessitates tight control of the IQ delay on the transmitter and receiver. The effect of residual IQ skew on next generation 64 Gbaud-64 QAM systems will become even more pronounced. ACO vendors will need to improve their calibration accuracy (typically within ± 3 ps today) to provide sub-picosecond skew resolution (~ 500 fs ideally) in a repeatable manner. As illustrated in Figure 2(c), it is worth noticing that the OSNR penalty is only function of the radial skew distance between H and V inner IQ skews. For instance, a $(-2, 3)$ ps IQ skew in H and V will yield the same penalty as a $(0, -3.6)$ IQ skew when referenced to an ideal case of $(0, 0)$ ps.




Figure 2. (a) OSNR penalty at FEC vs skew for 1.0 RRC QPSK, 8QAM, and 16QAM. (b) OSNR penalty at FEC vs skew for 1.0 RRC vs. 0.2 RRC 16QAM. (c) Top view of the 0.2 RRC 16QAM case with 500 fs gridlines. (d) 1.0 RRC vs. 0.2 RRC PAM4 eye for 16QAM

B. Why S21 matters

The wide variety of photonic technologies (i.e., InP, Si-ph, LNb) used in today's CFP2-ACOs and the underlying E/O driving electrodes' architectures, imposes on the host system to accurately compensate for the overall channel loss in the [0, Baud-Rate] window on a port to port basis depending on the specific ACO that is plugged in. As currently defined in the CFP2-ACO OIF agreement [2], in its class 2/3 variants, the module should provide the host with an accurate, mask-compliant, S21 profile per lane. However, it has been our experience that many ACO vendors do not provide reliable and accurate S21 E/O responses nor comply to the OIF's lower S21 mask limit. Nyquist shaped 0.2 RRC 16QAM signals in particular, once equalized for large channel losses, suffer a reduction in OSNR sensitivity if relying on these inaccurate S21 responses. For example in Figure 3(a), we purposely under-equalized and over-equalized the analog channel loss of a CFP2-ACO module plugged in the line card shown in Figure 1. Defining a normalized S21 slope alpha as $\alpha = \frac{S21_{test}}{S21_{ideal}}$ we measure in Figure 3(b) a ~0.5 dB penalty at FEC threshold when $0.6 < \alpha < 1.4$.

Looking forward, enabling ACO interfaces in 64 Gbaud, 64QAM systems will largely depend on module and connector vendors' ability to provide high bandwidth S21 responses while drastically limiting reflections.




Figure 3. (a) Under-equalized and over-equalized 0.2 RRC 16QAM optical spectra. (b) OSNR sensitivity for the different equalization cases

C. Auto-equalizing for next gen 64-64 ?

As shown in the two previous sections, the task of accurately calibrating ACO modules for next generation subsystems becomes crucial and does not offer much room for error. The repercussions of having to run every optics through a meticulous sub-picosecond skew calibration setup, in addition to the manual high-speed probing of the four differential pairs for S21 acquisition, could turn out to be a major disadvantage for a wide adoption of ACO-based approaches in next generation metro and long-haul networks. One solution that could address this concern, while saving time and cost to module vendors, would be to give the host system ownership of the analog channel tuning when the pluggable coherent optics is first inserted in any port. Of course, this could be achieved with an external loopback but would ideally be met by an internal optical feedback path that would be suppressed during operation. Such a design should be possible today in highly integrated Si-Ph platforms that have the transmitter and receiver collocated on the same PIC.

3. CFP2-ACO for the long haul

To exploit the full benefits of in-router integrated coherent optics (ICO) technology such as pluggable CFP2-ACOs in a layer 2/3 switch (Figure. 1), the ICOs must be paired with the appropriate open line system (OLS), in which the transponders are disaggregated from the line system common components such as amplifiers, ROADMs, and multiplexing/demultiplexing structures. The concept of the OLS is not new [3], but has recently become more practical to implement with advances in software-defined networking (SDN) and standardized ways of representing network elements to northbound interfaces (e.g., REST protocol with YANG data model, OpenConfig [4]). Typically for cloud service providers (CSP), layer 2/3 devices are already under the CSP's SDN control. With ICO line cards residing in those devices, it's logical to eliminate the DWDM line system's stand-alone, proprietary element-managing SW and bring the provisioning and monitoring capabilities of the line system under the same controller. This gives the CSP the most flexibility in tailoring the network configuration to their specific requirements, while providing a complete end-to-end view of network traffic flow. It also allows the line system, if properly architected, to be a long-term asset which can be used over multiple technology refreshes of router and ICO technology.

4. The 100 km DCI battleground, coherent vs. direct detect

For DCI applications less than 100 km, the economics and scale of a direct-detect solution which can fit into a QSFP28 form-factor are impossible to overlook. Rack space and power are typically limited and costly in data center facilities, and tens of inter-DC fiber pairs are typically easily procured, making density and cost more important than spectral efficiency. In addition, the density of state of the art CFP2-ACO-based integrated ICO line cards (8 ports) can't match that of the QSFP28 (36 ports), and switch port density in the tier in which DCI optics are used cannot be traded off.

In the near term, CFP2-based ICOs aren't yet the best fit for metro-reach (<100km) hyper scale DCI ecosystems for the density and cost reasons mentioned above. Longer term, if the port density of coherent solutions can reach parity with the equivalent grey port density, the prospect of in-router ICOs for DCI becomes much more interesting. In fact, recent advances in heterogeneous integration of InP within Si platforms on one side, and the advent of 7 nm CMOS DSP process on the other, indicate that a highly integrated DCI-focused tunable coherent solution targeting a single baud rate, modulation format, and very small amounts of dispersion compensation (< 120 km), will be able to achieve this parity within the next few years.

An example target specification for such DCI solution would be a 400 Gb/s single-carrier DWDM signal dissipating less than 4W/100G. There is currently wide support for the concept and, as of this writing, a project-start effort is underway within the OIF to ensure multiple vendor support and line-side SD-FEC interoperability. Such an interface should ideally meet requirements of the 400G Ethernet ER and ZR reaches.

5. Conclusions

The performance impact of IQ skew and S21 response on Nyquist shaped 16QAM signals is presented in the context of an 8-port coherent line card. Sub-picosecond error skew reporting and high bandwidth, accurate, S21 responses will help enable analog interfaces for future 64Gbaud-64QAM systems. The density imperative of next generation coherent metro DCI is also discussed in view of the ongoing advances in 100 km direct detect approaches.

6. References

- [1] M. Filer, J. Opt. Comm. Netw., v.8 (7), pp. A45-A54 (2016)
- [2] Implementation Agreement for CFP2-Analog Coherent Optics Module (OIF-CFP2-ACO-01.0, January 22, 2016). <http://www.oiforum.com/wp-content/uploads/OIF-CFP2-ACO-01.0.pdf>
- [3] M. Gunkel et al., "Vendor-interoperable elastic optical interfaces: Standards, experiments, and challenges [Invited]," J. Opt. Commun. Netw., vol. 7, pp. B184-B193, 2015.
- [4] <http://www.openconfig.net>, retrieved Oct. 30, 2016.