Design Tradeoffs for Coherent Pluggable Optics At 800G and Beyond

Eric Maniloff
Ciena Corporation, 383 Terry Fox Dr, Ottawa, ON K2K 2P5, Canada
emanilof@ciena.com

Abstract: As coherent pluggable optics scale to 800Gb/s, new applications will be addressed. This paper provides an overview of these applications, and the details of how module implementations differ to address different applications' unique requirements.

1. Introduction

Since the introduction of CFP Digital Coherent Optical (DCO) modules, DCO modules have evolved to play an increasingly important role in optical communications. CFP modules allowed DWDM 100Gb/s coherent interfaces to be deployed in the same pluggable form factors as single-wavelength client-side optical modules. Following this, 200Gb/s interfaces were deployed in CFP2 form factors, and more recently rates of up to 400Gb/s have been supported in CFP2, QSFP-DD, and OSFP form factors. The term DCO is now broadly applied to coherent modules in any pluggable form factor. New coherent line rates arrive first in discrete line card solutions, with pluggable modules following around five years later enabling lower power solutions and expanding the market reach.

As rates have increased, coherent interfaces have expanded their market coverage, with each increasing line rate being deployed in shorter reach higher volume markets. At 100Gb/s, coherent solutions were deployed in long haul and submarine markets, while at rates of 200Gb/s to 400Gb/s metro and Data Center Interconnect (DCI) transitioned almost entirely to coherent. The next generation of pluggable modules will continue this trend of evolution to new applications with increasing rates.

At rates up to 400Gb/s, coherent interfaces and applicable standards were almost entirely focused on amplified DWDM applications [1-3], while IMDD solutions using parallel lanes of up to 100Gb/s had sufficient reach for client side reaches of up to 10km [4]. IEEE is standardizing on client-side rates using 200Gb/s signaling for 800G Ethernet in 802.3dj, and at these rates fiber impairments such as Chromatic Dispersion and Polarization Mode Dispersion are an increasing challenge. Work is underway to use statistical analysis of fibers to reduce the penalty allocations for CD and PMD, to allow IMDD evolution for reaches up to 10km at 200Gb/s lane rates [5]. In addition, coherent standards for these 10km single-channel interconnections at 800Gb/s are being developed in both 802.3dj and OIF, with no constraints on the fiber CD and PMD parameters.

This paper highlights the different requirements for the key applications that 800 Gb/s coherent DCO modules will address, and details on how DCO designs are optimized to address these.

2. Coherent Evolution

As line rates increase to 800Gb/s modulation in the 120-130Gbaud range will be required. This will require DSP implementations using advanced process nodes and modifications to electro-optical designs.

800Gb/s ASICs will be introduced in 5 nm and 3 nm silicon, resulting in ~20-50% power savings respectively compared to the 7 nm process used at 400G. These new process nodes will allow the use of new tools for 800G tailored to the specific application requirements.

When evaluating the Required SNR (RSNR) at the receiver of various modulation formats, accounting for the imperfect nature of the DSP and optics implementation is important. The RSNR deviation from the theoretical RSNR ($RSNR_{th}$) of a modulation format can be approximated using Eye Closure (EC) and Implementation Noise (IN), where EC represents the deviation of the mean modulation points from their theoretical perfect locations and IN represents white Gaussian noise added to the signal. The relationship between RSNR and $RSNR_{th}$ is given by:

$$1/RSNR = 1/(EC * RSNR_{th}) - 1/IN \qquad (1)$$

This equation provides a way to evaluate the practical performance of various implementations. In addition, standards require the inclusion of an additional noise term to account for the performance reduction encountered due to interoperation between different module suppliers.

3. 800G Applications

3.1 Intra Datacenter: 10 km: 800G LR

Optimized coherent solutions for 2-10 km single channel 800Gb/s are being developed in both OIF and IEEE 802.3dj. At these shorter reaches, latency and power consumption are key considerations. Ethernet has used Reed-Solomon FEC at rates up to 400GbE, with typical latencies under ~100ns. A coherent implementation that addresses these requirements has been designed using a concatenated FEC scheme [6]. In previous IEEE 802.3 standards, RS-FEC was implemented end-to-end, with being generated and terminated on host ASICs. In 800LR, to meet the required power, latency, and SNR requirements, an inner BCH(126,110) FEC code is added to the end-to-end RS FEC. This concatenated FEC has a 10.4dB Net Coding Gain (NCG). This type of concatenated FEC scheme is illustrated in Fig. 1.

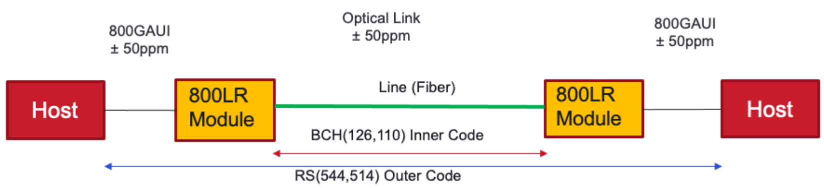

Fig 1. Illustration of concatenated FEC designed for 10 km 800G Ethernet

In DWDM applications FEC is terminated across each interface, with a higher coding gain FEC used on the optical side. For 800G LR, the coherent specification uses a dual-polarization 16QAM modulation format, with a modulated rate of 123.6 Gbaud. The addition of the inner FEC results in a latency of ~250 ns. Although higher than RS FEC, this latency is equivalent to 50m of fiber propagation, which meets the application requirements.

In most coherent modules a single laser is shared between the transmitter and local-oscillator. For DWDM applications, laser frequency accuracy is typically specified to be ≤ ± 1.8GHz, and fully tunable lasers are used. The tight frequency accuracy helps minimize the spacing between channels in order to increase spectral efficiency, and allows the receiver's DSP to acquire its signal for worst-case offsets between the two ends' lasers. For single channel applications, tight frequency accuracy is not required and unlocked lasers can be used with potentially larger frequency offsets, in the range of ± 12.5GHz. Using separate lasers for transmit and receive improves optical power budgets and allows the receive laser to tune its frequency to match the transmit frequency.

3.2 Data Center Interconnects: 800ZR

DCI applications are point to point links covering 80-120 km maximum reach. These applications use filtering at the Multiplexer and Demultiplexer and are designed to provide high capacity over the desired reach. An illustration of a DCI application is shown in Fig 2.

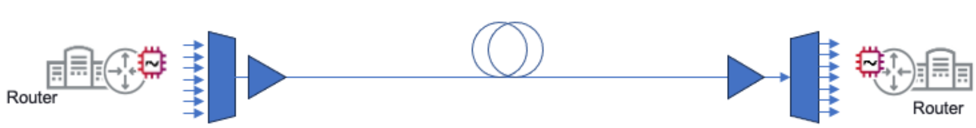

Fig 2. DCI architecture. DCO modules are connected point-to-point through an amplified line with filtered optics.

At 400G, OIF introduced an implementation agreement for DCI applications, based on a concatenated Hamming/Staircase FEC code with 10.8 dB net coding gain, allowing 64 channels at 59.8 Gbaud modulation spaced at 75GHz [2]. To improve the capacity per channel to 800Gb/s, a higher gain FEC will be used along with a baud increase, to improve the SNR tolerance. This is accomplished by using oFEC on a 118 Gbaud DP-16QAM modulation, with an 11.6 dB NCG. The higher coding gain is intended to allow comparable reach as was achieved in 400ZR, for 32 150 GHz-spaced channels. While 400ZR was adopted with a 26dB OSNR requirement, 800ZR increases this requirement to 27dB, assuming the same allocations as 400ZR: 26dB IN, 0.4dB EC, and 25dB additional noise for interoperability between suppliers. For these DWDM applications, a spectral mask is adopted to manage optical crosstalk between adjacent channels. Since these are optically filtered applications, a relatively wide spectral mask can be used, based on a root-raised cosine shape with a roll-off factor of 0.4.

3.3 Metro Networks: OpenROADM and ITU

Metro networks require longer reach, and often use colorless combiners into Reconfigurable Add Drop Multiplexers (ROADM). Reaches in these applications are typically in the 500km range, traversing OADM nodes separated by 10's of km. To achieve this reach, transmission schemes that are more tolerant to SNR are needed. Colorless combiners require control of out-of-band OSNR to control the add-noise at colorless combiners, and increased transmit optical power. Meeting the Tx OSNR and power requirements is typically achieved using tunable optical filters and optical amplification built into the DCO transmitter, as well as DSP designs for sharper spectral roll-off.

For 400G, meeting these reach requirements was accomplished by using oFEC. While this FEC scheme will be carried over to higher data-rates, further improvement is required to maintain the same reach targets. At 800G an interoperable Probabilistic Constellation Shaping (PCS) scheme has been developed in OpenROADM to further improve the SNR tolerance [7]. The addition of PCS improves the theoretical SNR requirement by 0.8 dB. Including 16dB IN and 0.4dB EC terms increases this improvement to ~1.5dB as compared to an unshaped 800G DP-16QAM. This additional PCS gain will allow the added noise tolerance needed for these more demanding applications.

In addition to these interoperable specifications, DCO modules are also designed to address longer reach applications, using proprietary schemes, such as more advanced PCS modulations. In these applications there are fewer ROADMs, but the SNR requirements demand higher performance and maximizing fiber capacity requires tight spectral shaping.

4. Conclusion: Scaling to 1.6 Tb/s

At 800Gb/s, coherent solutions have been adapted to single channel, DCI, and Metro/ROADM applications, using differing bauds, FEC, and PCS implementations.

Coherent solutions will continue to evolve in the solutions mentioned above to support coming 1.6 Tb Ethernet rates. Both 2x800G solutions as well as baud increases to a 200-240Gbaud range are potential solutions to address these applications.

The trend of addressing shorter reaches as rates increase makes the extensions of coherent 500m-2km intra datacenter reaches practical. This will require new approaches, to minimize latency and power. The improved sensitivity of QPSK has been proposed as a potential solution to allow RS FEC end-to-end for intra DCI coherent [8]. Tradeoffs of power and latency for these applications will be key topics in tailoring coherent solutions at 1.6T.

5. References

[1] ITU-T G698.2 (11/2018)
[2] "Implementation Agreement 400ZR", OIF-400ZR-2.0, November, 2022
[3] OpenROADM MSA optical specifications (v5.1)
[4] IEEE Standard for Ethernet, IEEE 802.3-2022
[5] C. Cole et al, "Update to Modern SMF Parameters for Calculating PMD (Physical Medium Dependent) Penalties", IEEE 802.3df, https://www.ieee802.org/3/df/public/22_11/cole_3df_01a_2211.pdf
[6] K. Kota et al, "Logic Baseline proposal for 800G single-wavelength coherent PHY with concatenated FEC", IEEE 802.3dj, https://www.ieee802.org/3/dj/public/23_07/kota_3dj_01b_2307.pdf
[7] G. David Forney et al, "Multidimensional Constellations-Part I: Introduction, Figures of Merit, and Generalized Cross Constellations", IEEE Journal on Selected Areas in Communications, Vol. 7, No. 6 (1989).
[8] Hongbin Zhang, "Power Efficient Coherent Detection for Short-Reach System", OFC 2023, M1E.1.