ARE YOU READY FOR SUBMARINE 400 GBIT/S?

Jean-Sébastien Tassé (EXFO), Brian Lavallée (Ciena)
Email: jean-sebastien.tasse@exfo.com

EXFO, 2500 Alfred-Nobel, Montréal (Québec), Canada, H4S 0A9

Abstract: 400 Gbit/s roll outs in subsea networks are about to become common, following recent deployments of 200 Gbit/s on subsea cables such as MONET in the Atlantic Ocean and ATISA in the Pacific. With the first single wavelength 400 Gbit/s already deployed in terrestrial networks in 2017, subsea networks are next.
400 Gbit/s promises added capacity on subsea routes as well as terrestrial backhaul, on the client side. Like any new technology, 400 Gbit/s raises a number of questions: will existing cables support it, what are the key design considerations and key limitations for 400G, etc. This presentation will explore the new challenges pertaining to 400G deployments as compared to 100G or 200G, with an emphasis on the key performance indicators to take into account, both at the optical layer level (signal-to-noise ratio for instance), as well as the transport level (e.g., key Ethernet performance indicators, Flex Ethernet applications). We will also use our key findings from working with 400G terrestrial networks to discuss the impact of this new technology on subsea networks.

1. INTRODUCTION

To maintain pace with ongoing growth in bandwidth demand, the network industry is quickly moving towards switches and routers supporting 400GbE interfaces. This is leading to a demand for universal optical transport of 400GbE client rates over a single-wave 400G across any distance, including subsea cases.

2. 400G SUBSEA DEPLOYMENT UPDATE

At the time of this paper's writing, there exists no single wave 400G commercial subsea deployments. Ongoing field trials show great promise, but do not include real-world deployment requirements. As demand for data center interconnection applications and associated growth rates continue unabated, the need for single wave 400G transport across the Atlantic and Pacific Oceans will be required.

3. 400G SUBSEA CABLE SUPPORT

Older dispersion compensated submarine cables have challenges resulting in single wave 400G transport not being economically feasible. Newer, uncompensated subsea cables will support single-wave 400G with near future solutions. The proposed Spatial Division Multiplexing (SDM) subsea cables could limit the ability to support singe-wave 400G, as decisions made when selecting an SDM design will ultimately dictate its maximum spectral efficiency and will also drive what baud rates would be required to achieve 400G.

4. KEY DESIGN CONSIDERATIONS FOR 400G

It is well-known that the data rate, also called bit rate, is the product of the baud rate by the number of bits per symbol of the chosen modulation format and the number of polarization states used. Although there are different combinations of baud rate and modulation formats used to achieve 100G transmission, a very common one is dual polarization quadrature phase shift keying (DP-QPSK) modulation format with 4 bits/symbol and a baud rate of about 28 to 32 GBd. 400 Gbit/s will typically be achieved by using a combination of 56 GBd or higher, and higher order modulation formats such as dual polarization 16 quadrature amplitude modulation (DP-16-QAM) that offer more bits per symbol (figure 1). For instance, DP-16-QAM provides 8 bits/symbol, while DP-32-QAM offers 10 bits/symbol.

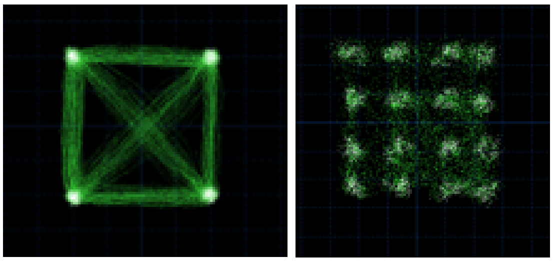

Figure 1: Constellation diagrams of a QPSK signal on left and 16-QAM signal on right

Although solutions will vary in the modulation scheme and constellation shaping techniques used to achieve better performance, the faster baud and higher order modulation formats used to achieve 400G have two significant impacts:

- A higher Optical Signal-to-Noise Ratio (OSNR) requirement, reducing the maximum reach
- A larger signal width required to accommodate the larger data throughput, resulting in fewer total channels deployed thus requiring less hardware for reduced overall cost economics

OSNR is a key design consideration for any subsea network, because it quantifies the signal quality. More precisely, it is according to standard IEC 61282-12 the "ratio of time-averaged optical signal power to the maximal spectral power density of the amplified spontaneous emission (ASE) noise within the wavelength range of the total signal spectrum, normalized to a chosen reference bandwidth" [1], the reference bandwidth being typically 0.1 nm [1].

The noise comes predominantly from optical amplifiers, so the longer the subsea link, the more optical amplifiers there are and the more challenging it becomes to meet the required SNR (RSNR) or required OSNR (ROSNR) of the receiver. If the RSNR/ROSNR is not met, errors will ensue. Moreover, the SNR/OSNR requirement increases with higher modulation formats, as shown in Table 1. Concretely, this means that higher order modulation formats can travel shorter distances than lower order modulation formats, assuming techniques such as probabilistic constellation shaping are not applied. For instance, without Probablistic Constellation Shaping (PCS), QPSK signals support transpacific distances, while 16-QAM signals are limited to transatlantic distances.

Modulation Format	ROSNR	RSNR @ 35Gbaud
QPSK	~10-14dB	~ 5.4-9.4dB
16-QAM	~16-20dB	~11.4-15.4dB
64-QAM	~22-26dB	~17.4-21.4dB

Table 1. ROSNR and RSNR threshold at receiver for different modulation formats (no PCS, with FEC, ASE dominated noise)

From a subsea cable owner or operator perspective, it is therefore important to know the ROSNR/RSNR of the receiver, so that the appropriate pass/fail evaluation can be performed using OSNR/SNR measurements taken during maintenance.

The second impact of moving from 100G to 400G is the larger channel width, which is a direct consequence of the faster baud that is required to achieve the higher capacity wavelength. This means that more capacity, including emerging 400GbE client rates, can be transported across subsea distances over a single wavelength, requiring less hardware to deploy and manage. Because of the wider spectrum that is used per channel, it is important that the new 400G channels provide better spectral efficiency than the 100G/200G channels they are replacing to justify their use.

Accordingly, designing a 400G subsea link is a delicate balance between modulation format and baud to maximize the spectral efficency, taking into account the link reach. The coherent solution must support the baud required to achieve single 400G wavelength transmission across all, or the majority of subsea cables. Each subsea cable is therefore unique.

5. 100GbE VERSUS 400GbE ON THE CLIENT SIDE

So far, we discussed line side transmission, that is the portion of the network between the landing stations. Equally important is the client side, i.e. the terrestrial portion of the network from the landing station, using other types of transceivers. It should be noted that the notion of client side and line side is rapidly evolving, with subsea line side signals being sometimes backhauled directly in data centers at tens of km from the beach.

The first innovation brought about by 400G is FlexEthernet is an implementation agreement by the Optical Internetworking Forum that describes the mechanism to transport a range of different Ethernet MAC (i.e. protocol layer) rates not based on any current Ethernet PHY (i.e. physical layer) rate. An objective of this new implementation agreement is to improve and maximize the interconnection between network elements (routers) and transport gear. A second objective is to help network elements reach the bandwidth being currently handled today by transport elements that use coherent technologies.

FlexE offers three main use cases between the router/switch and the transport network : bonding, sub-rating and channelization as exhibited in Figure 2.

Bonding: Bonding offers the capability to combine several PHYs to transport a MAC of higher capacity.

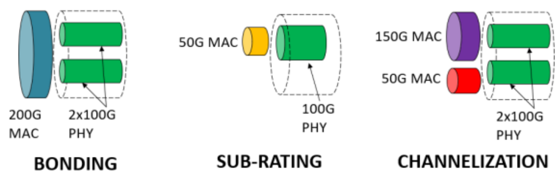

Figure 2: Applications offered by FlexE

Sub-rating: Allows network elements to match the MAC rate to the capacity of a line-side interface that is not a multiple of the PHY rate. It also enables the transport of multiple Ethernet Macs on a single PHY.

Channelization: Provides the capability to aggregate multiple Ethernet MACs over a PHY or bonded set of PHYs and use the transport equipment to independently transport each MAC to the desired destination. Channelization allows the consolidation of variable MAC rates over a set of PHYs all operating at a common rate.

Given the large diversity of modulation formats discussed above, FlexE becomes very appealing to transport non-standardized Ethernet rates. For instance, 8-QAM modulation, which typically handles 150G per wavelength, a non-standardized rate, can be transported on the client side with sub-rating.

A second difference of 400G versus 100G is the introduction of a mandatory and improved Reed-Solomon FEC (forward error correction) at 400G, while 100G FEC was optional.

Along with FlexE, the third difference is probably the most important: the availability of new form factor pluggable devices at 400G. Currently, three major form factors exist for 400GbE, and it is yet unsure which ones will prevail: QSFP-DD, OSFP and COBO.

Form factor	Pros and cons
QSFP-DD	Higher port density High power consumption
OSFP	Medium port density High power consumption
COBO	Optimal port density Lower power consumption

Table 2: Pros and cons of 400GbE form factors

6. KEY 400GbE PERFORMANCE INDICATORS AT THE TRANSPORT LEVEL

RFC 2544 has been the most widely used Ethernet service testing methodology. This series of subtests provides a methodology to measure throughput, round-trip latency, burst and frame loss. It was initially introduced as a benchmarking methodology for network interconnect devices in the lab. While this testing methodology provides useful information to qualify the network, it is no longer complete enough to test today Ethernet services. More precisely, RFC 2544 does not feature all required measurements, such as packet jitter, QoS measurement, and multiple concurrent service levels. In addition, since RFC 2544 requires the performance of multiple, sequential tests to validate complete Service Level Agreements (SLAs), this test method takes several hours, a time-consuming process. The ITU-T Y.1564 standard has been introduced to precisely address those issues: EtherSAM (the test implementation of Y.1564) supports new multiservice offerings, because it can simulate all types of services that will run on the network, and simultaneously qualify all key SLA parameters.

7. CONCLUSION

This paper states 400 Gbit/s deployments are coming very soon, enabled by modems with faster baud rates and higher order modulation formats. Two key design considerations are the higher OSNR/SNR requirements of higher order modulations, as well as larger 400G signal widths, which offer better spectral efficiency than 100G signals due to improvements in FEC and signal processing. On the client side, FlexEthernet is a major advance of 400G, because it offers channelization, sub-rating and bonding. New 400G form factors are also available on the client side, such as QSFP-DD and OSFP. Finally, although RFC 2544 Ethernet service method can be applied to 400GbE signals, the more recent ITU-T Y.1564 is a better choice to qualify SLA.

8. REFERENCES

[1] IEC technical report 61282-12, edition 1.0, “Fibre optic communication design guides – Part 12: in-band optical signal-to-noise ratio”

