Networking Benefits of Coherent Pluggable Optics

João Pedro[1,2]
1 - Infinera Unipessoal Lda., Rua da Garagem 1, 2790-078 Carnaxide, Portugal
2 - Instituto de Telecomunicações, Instituto Superior Técnico, Av. Rovisco Pais 1, 1049-001 Lisboa, Portugal
JPedro@infinera.com

Abstract: This paper overviews the range of applications of coherent pluggable optics. Selected simulation results highlight how deploying devices that feature both high-performance and digital subcarrier multiplexing is a key enabler of cost-effective network solutions. © 2024 The Author

1. Introduction

Coherent pluggable optics stand as one of the most impacting optical networking developments in recent years. The original driver for these devices was the growth in data center interconnect (DCI) applications, which lead to the OIF proposing the 400ZR implementation agreement (IA) for low-cost, short-reach, interoperable, and pluggable 400G coherent interfaces [1]. 400ZR laid the groundwork for improved coherent pluggable optics specifications, featuring, for example, better forward error correction (FEC) and modulation format adaptation, which enable longer transmission distances and a broader set of applications. The OpenZR+ multi-source agreement (MSA) defines optical channel formats and the appropriate mapping of Ethernet client signals (100G, 200G and 400G) [2] for implementing coherent pluggable optics in a QSFP-DD form factor, whereas the OpenROADM MSA also includes the support of OTN client signals and can be implemented in QSFP-DD and CFP2 form factors [3]. Pluggable optics compliant with these MSAs can be deployed to cover typical metro network distances and traverse ROADM nodes with different architectures [4]. Currently, the set of key MSAs is completed with the Open XR Optics Forum MSA, which targets the emergent point-to-multipoint (P2MP) coherent pluggable transceivers and the requirement to guarantee multi-vendor interoperability and an open, multi-source solution ecosystem [5]. Figure 1 highlights the bifurcation of the coherent transceiver market, which now features embedded and pluggable devices.

Importantly, standardization and implementation/multi-source agreements need to evolve to accommodate the industry requirements. For instance, the latest OpenZR+ specification (v.3.0) specifies high Tx output power modes, support for colorless add/drop structures, and a higher baud rate format for longer reach 400G. Moreover, the Open XR Optics Forum has specified a management architecture, paramount to address not only the commissioning and management of subcarriers and the abstraction of P2MP data flows, which is specific of the digital subcarrier multiplexing (DSCM) solution to realize P2MP [6], but also the industry-wide issue of network management when pluggable optics are installed at routers and switches, as illustrated in Fig. 2. Unlike in the traditional architecture, where there is a clear device-level demarcation of transport and IP domains, the IP-over-DWDM architecture requires solutions to manage a device with transport functions that is now physically present in the IP domain [5]. It should be noted that coherent pluggable optics can also be beneficial when deployed in the traditional architecture with strict domain segmentation, e.g., by designing line cards (e.g., muxponders, switchponders) to accommodate them instead of embedded optics, maximizing their utilization across the portfolio of a network solutions vendor [7].

This paper overviews the applications of coherent pluggable optics, focusing on opportunities to expand their role in different network segments and identifying the main capabilities/features it imposes on these devices. It also discusses how additional savings can be realized in some applications, particularly when P2MP optics are used.

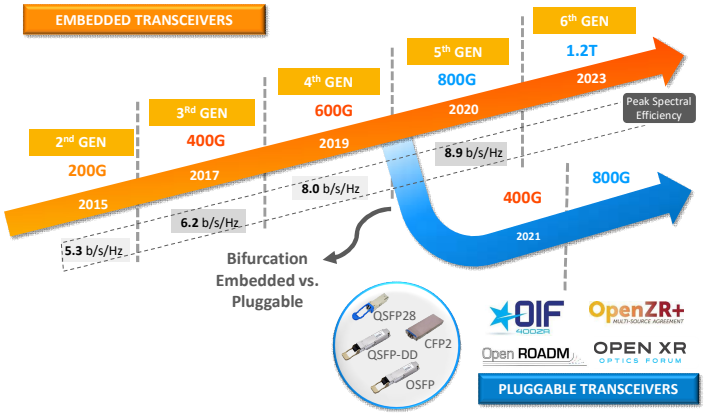

Fig. 1. Coherent transceiver evolution.

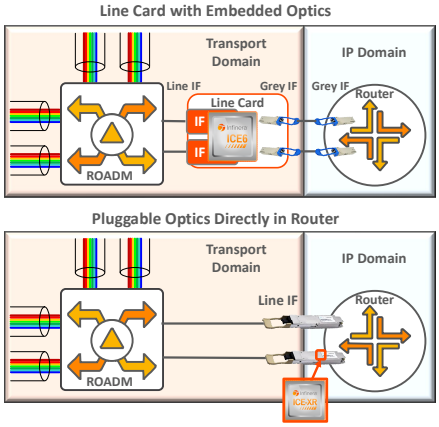

Fig. 2. Embedded vs. pluggable optics.

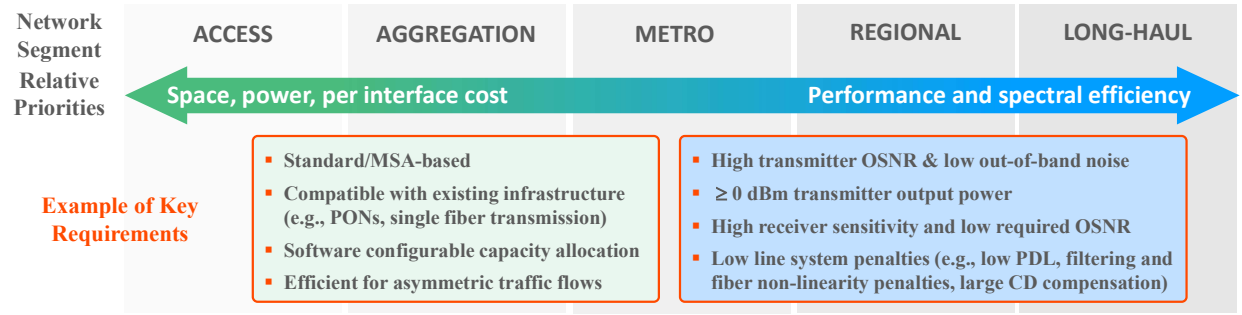

Fig. 3. Example of key requirements to expand the utilization of coherent pluggable optics.

2. Expanding the Role of Coherent Pluggable Optics

As described, the market for coherent pluggable optics has grown from DCI to metro networks, with potential to further expand the set of applications covered. Figure 3 illustrates the relative priorities set by different network segments and how it can translate into key requirements that transceiver solutions and their system-level capabilities must meet to successfully fulfil those priorities. For instance, detailed studies in [7, 8] have shown that in highly loaded regional and long-haul networks the benefits of embedded optics are still evident when compared to using 400ZR-based hop-by-hop and OpenZR+/OpenROADM-grade coherent pluggable optics with optical bypass, with [9] arguing that the lowest-cost solution can be a mix of pluggable and embedded optics. Thus, higher performance pluggable optics are required to compete with embedded ones in performance-demanding deployments, which means these devices need to exceed the performance targets set by MSAs. However, using improved capabilities, such as proprietary FEC, advanced modulation formats, DSCM, can prevent interoperability between devices from different vendors. The benefits are, nevertheless, compelling and may justify accepting bookended connections in these network segments. However, the main challenge is the feasibility of designing a high-performance pluggable device that still meets the stringent footprint and power consumption limits imposed by the small form factors.

Coherent pluggable optics face different challenges in network segments dominated by intensity-modulated with direct-detection (IM-DD) pluggable devices. Besides the lower cost targets, the compatibility with existing fiber infrastructures and the need to efficiently support sub-100G, asymmetric and often dynamic traffic flows require an architecturally disruptive solution, such as implementing P2MP connectivity in the optical domain via DSCM.

3. Leveraging the Potential of P2MP Coherent Pluggable Optics

DSCM has been used in commercially deployed 800G embedded optics to improve performance [10]. Importantly, it has been demonstrated conceptually, as well as in lab and field trials, that this technique can also be exploited to cost-effectively realize point-to-multipoint connectivity in the optical domain [6, 11]. This means DSCM holds the promise of enabling an architectural shift, by breaking the established paradigm of always using point-to-point (P2P) optical connections and imposing that the transceivers at the end nodes have necessarily the same capacity. In fact, the resulting P2MP solution offers a super-set of functionalities, when compared to the traditional approach, since it still retains the capability to support P2P connections in addition to introducing P2MP [6]. A key principle of this solution is to specify the individual subcarrier (SC) and design pluggable devices that have a capacity measured in number of SCs. For example, if a 25G SC is defined as a 4 Gbaud signal modulated with 16QAM, a 100G device supports 4 SCs (~16 Gbaud) and a 400G transmits/receives 16 SCs (~64 Gbaud). The higher capacity pluggable device can communicate with multiple lower capacity ones and dynamically allocate a certain number of SCs to each, allowing to trade-off the number of lower capacity devices reached and the capacity provided, as long as the total number of SCs involved does not exceed 16 in case the high-capacity pluggable transceiver is a 400G. This section summarizes some of the applications and the resulting benefits from leveraging the potential of DSCM as a fundamental enabler to expand the role of coherent pluggable optics in access, aggregation and metro networks.

- *Cost-effective filterless horseshoe metro-aggregation networks with P2MP transceivers*: upstream transmission, i.e., from multiple lower capacity devices at leaf nodes to a single high-capacity device at the hub node, simply consists of optical grooming of the SCs originated at the leaf nodes. Combined with the broadcast nature of the downstream transmission, mean DSCM-based P2MP transceivers are a perfect match to exploit the low-cost filterless architecture [12] in metro-aggregation networks that usually have a horseshoe topology with two hub nodes (see Fig. 4). A detailed analysis over real metro-aggregation networks from Telecom Italia Mobile (TIM) highlights that the proposed P2MP transceivers can reduce capital expenditures (CAPEX) between 25 and 37% when compared to P2P transceivers and considering different traffic scenarios [13]. Further lowering the cost of the overall network solution is possible via selectively deploying optical amplifiers and customizing the power

splitters/combiners. The optimization framework described in [14] to achieve this objective while respecting physical constraints (e.g., receiver sensitivity, maximum launch power for negligible nonlinear impairments, maximum power unbalance between SCs at the hub's receiver) shows that the number of optical amplifiers can be halved (see Fig. 5), which further adds to the benefits of replacing P2P by P2MP transceivers.

- *Long-term cost-effectiveness and environmental impact of P2MP transceivers*: in addition to the benefits against state-of-the-art P2P coherent pluggable transceivers, it is paramount to demonstrate that P2MP transceivers are expected to remain cost-effective in the long-term. The study reported in [15] considers a metro-aggregation network evolving over a 5-year period, with the introduction of higher capacity transceivers and modelling different planning strategies that account for varying degrees of confidence on earlier traffic predictions. The results provide evidence that CAPEX reductions from using P2MP instead of P2P transceivers are typically in the range of 30% and remain above 20% even in the scenarios more favorable to P2P transceivers. Moreover, an increasingly important aspect of telecommunication networks is their environmental impact. Comparing the use of 400G/100G P2MP transceivers against using 400ZR transceivers with hop-by-hop routing or 400G OpenZR+ devices with optical bypass, the analysis presented in [16] over a 5-year period of network operation highlights that in addition to the immediate benefits of device count reduction, a significant decrease in power consumption and CO_2 emissions can be observed when opting for the P2MP solution instead of the P2P ones.
- *Router bypass and converged metro-core/access networks*: communications service providers (CSPs) rely on a hierarchical architecture composed of multiple layers (e.g., Telefónica uses 5 layers from access to core). The ability to bypass intermediate layers, without compromising functionality and network operation robustness, is an appealing proposition as it would allow to reduce the number of transceivers and router ports required. The joint investigation with Telefónica reported in [17] provides evidence that CAPEX reductions above 40% can be achieved by leveraging P2MP transceivers to realize a selective router bypass. Similarly, in a different study described in [18] it is also shown that P2MP transceivers can be effectively exploited to provide connectivity across metro-core and metro-access networks with expressive savings in both transceiver and router port count.
- *Bidirectional single fiber transmission and passive optical network (PON) overlay*: in access networks it is critical to ensure compatibility with existing infrastructures, which often comprise a single fiber for bidirectional transmission. Using traditional P2P coherent pluggable optics for bidirectional transmission on the same fiber via different wavelengths for downstream and upstream is costly since it demands a second laser for the local oscillator [19]. DSCM-based transceivers can efficiently address this application by allocating different SCs to each transmission direction [20], as illustrated in Fig. 6. Hence, it becomes possible to cost-effectively provide high-capacity connectivity, such as 5G fronthaul applications, over existing PON infrastructures [21].

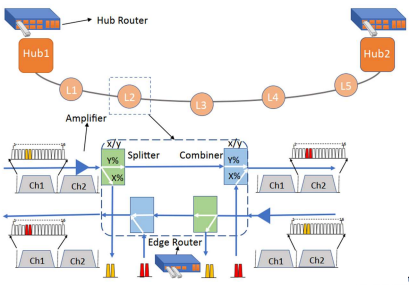

Fig. 4. Filterless horseshoe network with P2MP.

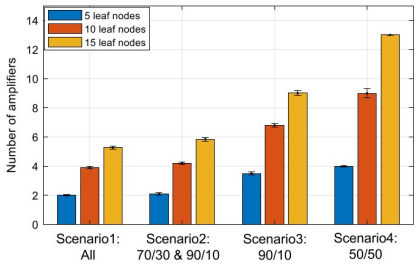

Fig. 5. Number of amplifiers required [14].

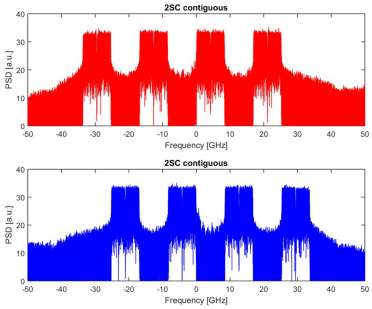

Fig. 6. Bidirectional transmission with DSCM.

Acknowledgements: This work has been partly supported by SEASON (101096120) project and by Fundação para a Ciência e Tecnologia (FCT) project UIDB/50008/2020. The author is grateful to Nelson Costa, Mohammad Hosseini, Ashwin Gumaste, Antonio Napoli, Carlos Castro, Yesica Rumaldo, Johann Bäck, and Harald Bock for insightful discussions.

References

[1] OIF 400ZR IA. https://www.oiforum.com/technical-work/hot-topics/400zr-2/
[2] OpenZR+ MSA. https://www.openzrplus.org/
[3] OpenROADM MSA. http://www.openroadm.org/
[4] J. Pedro, *et al*, *IEEE LATINCOM 2022*, TS4.2.
[5] OpenXR Optics Forum MSA. https://openxropticsforum.org/
[6] D. Welch, *et al*, *JLT*, 39(16), 5232-5247.
[7] J. Pedro, *et al*, *JOCN*, 14(2), A154-A165.
[8] A. Gumaste, *et al*, *OFC 2022*, Th1F.2.
[9] P. Wright, *et al*, *ECOC 2020*, Tu2H-2.
[10] H. Sun, *et al*, *IEEE/OSA JLT*, 38 (17), 4744-4756.
[11] J. Bäck, *et al*, *ECOC 2020*, Tu2H-1.
[12] C. Tremblay, *et al*, *IEEE FNWF 2022*.
[13] C. Castro, *et al*, *JOCN*, 15(5), B53-B66.
[14] M. Hosseini, *et al*, *JOCN*, 15(9), 569-578.
[15] M. Hosseini, *et al*, *JOCN*, 15(3), 155-162.
[16] F. Masoud, *et al*, *ONDM 2023*.
[17] A. Napoli, *et al*, *OFC 2022*, W3F.5.
[18] A. Gumaste, *et al*, *ECOC 2022*, We5.58.
[19] D. Welch, *et al*, *JLT*, 41(4), 1175-1191.
[20] A. Rashidinejad, *et al*, *OFC 2020*, W2A.30.
[21] A. Napoli, *et al*, *ECOC 2021*, Th1F.3.