European Conference on Optical Communication (ECOC) 2022 © Optica Publishing Group 2022

Exploring Point-to-Multipoint Coherent Capabilities Across Metro and Core Networks

Ashwin Gumaste [(1)], Joao Pedro [(2)], Harald Bock [(3)]

[(1)] Infinera Corporation, USA, agumaste@infinera.com [(2)] Infinera Unipessoal Lda, Carnaxide, Portugal, jpedro@infinera.com [(3)] Infinera GmbH, Munich, Germany hbock@infinera.com

Abstract *We investigate point-to-multipoint coherent capabilities for traffic grooming and provisioning across interdomain metro-edge and metro-core networks. Results highlight benefit of P2MP coherent from a transceiver count perspective. ©2022 The Author(s)*

Introduction

Typical IP over WDM architectures are evolving to meet the growth of service line-rates and corresponding availability of pluggables at rates beyond 100Gb/s into the realm of 400+Gb/s. Two innovations are at the forefront of this evolution: the availability of optical engines and pluggables at 400Gb/s and beyond and router fabrics at several Tb/s. In a typical operator environment, residential and enterprise users connected access networks as well as leased line customers routinely lead to aggregation of traffic at metro edge/access nodes of 100Gb/s and beyond, implying common usage of 100Gb/s ports at the metro edge. The metro network hence is characterized by nodes that can have multiple 100Gb/s and sometime higher rate (400Gb/s) traffic mapped on to wavelengths that are routed towards metro-to-core interconnect nodes (see Fig. 1). At the core nodes, most of the traffic is processed by a large (several Tb/s) (Provider) P-router or an OTN cross-connect (XC) for leased lines. The router/OTN XC both house pluggables of various line-rates (10, 100, 400Gb/s) in their connection to the metro edge, while in their connection to the core use either high-speed (400Gb/s) pluggables for medium reach, or embedded interfaces (for larger core networks and to preserve spectral efficiency). As can be seen in Fig. 1, metro edge nodes comprise of ring-supporting ROADMs and provider-edge (PE) routers. The PE routers usually have up to few 10Gb/s and largely 100Gb/s WDM pluggables, steadily evolving towards 200Gb/s and 400Gb/s (OpenZR+). At the metro-to-core interconnection point, traffic from the metro edge is dropped via a multi-degree ROADM on to a large P router. The P-router aggregates traffic towards other P-routers or to data-centers through the core network. The P-router houses on one side WDM pluggables to interconnect with the metro network. At the core side, the P-router must map the traffic to higher speed wavelengths (typically 400Gb/s and beyond) that can traverse a large distance. For this, the P-router connects to an optical engine (called Xponder) through grey interfaces. The Xponders aggregates multiple 100Gb/s or 200Gb/s clients from the P-router and maps these onto a high-baud rate 400-800Gb/s wavelength that expresses through the core network. Essentially, the role of the P-router is that of an aggregator, whereby lower rate channels from the metro/metro access network are aggregated onto higher rate channels in the core. In such an architecture, the key question we want answer in this paper is to determine the network-wide hardware requirements, and specifically transceiver count, which is the single biggest variable commodity in the network using multiple technology options. These multiple options are as shown in Figs 2, 3, 4. Fig. 2, 3, and 4 use two novel components, a coherent point-to-multipoint (P2MP) coherent transceiver operating at up to 400Gb/s rates and an Xponder interfacing the core at up to 800Gb/s rates. For a set of traffic mixes, we want to compute, which approach would result in the least number of hardware components (and hence cost).

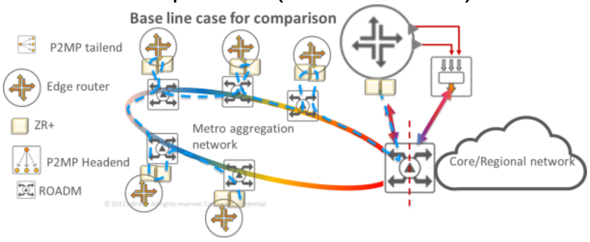

Fig. 1: Baseline case.

Network Architecture Choices

We now detail the different architectures that we consider for comparison. A variation of the baseline model is the hop-by-hop model in which from the ingress metro-edge node to the core node, the traffic is mapped on to an optical circuit, which is dropped and processed at each node along the path. In this model, we assume ZR/ZR+ pluggables at 400Gb/s with QSFP-DD form factor plugged into routers. The benefit of the hop-by-hop model is that there is no requirement of wavelength continuity, and the ability to groom lower-rate traffic onto line (wavelength) at interim nodes on the path. The problem with this approach is the mandatory requirement of transceivers and router ports at each node.

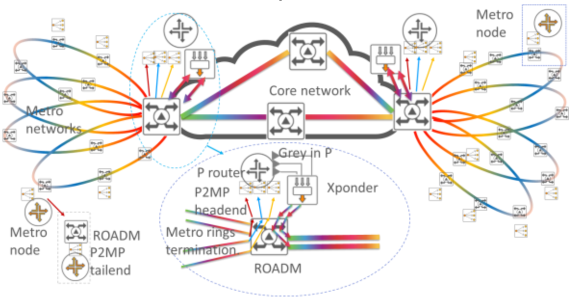

Fig. 2. Use of P2MP optics.

In Fig. 2 is a model that introduces a novel component: a point-to-multipoint (P2MP) optical module along with an Xponder. The P2MP module has a head-end and a tail-end [1]. The head-end sends a 400Gb/s channel using DSCM (Digital Sub-carrier Multiplexing) that realizes the P2MP architecture. We consider two variants of the DSCM architecture, realizing 25Gb/s per tail-end node (in 16x25Gb/s configuration), and up to 100Gb/s tail-end node (in 4x100Gb/s configuration). Metro-edge nodes support tail-end modules that latch onto one or more digital sub-carrier(s) (SCs) from the head-end in the downstream direction, while sending a similar digital sub-carrier towards the upstream (which is done by a low-speed transceiver at the metro-edge node (see tail-end in Fig. 2). A simple drop-and-continue architecture can be used: in the downstream the entire modulated composite channel (consisting of up to 16 SCs) reaches the different metro-edge receivers, each of which will only process its intended SCs; in the upstream, metro-edge transmitters generate their SCs, which are optically groomed until reaching the head-end node [7]. Importantly, this approach enables to allocate SCs (in multiples of 25Gb/s) according to traffic requirements, enabling an effective use of both tail-end and head-end node transceivers. Recent works have demonstrated the feasibility of such a system [4] for metro/metro-edge networks [7]. Hence, in Fig. 3, the requirement to connect to *N* metro-edge nodes tapers down to the use of just *N+1* transceivers (*N* metro edge node and 1 at the PE router, i.e., the P2MP head-end) for an aggregate capacity of 400Gb/s.

Shown in Fig. 3 is a further optimized version of Fig. 2 specifically adapted to content distribution networks [2]. Here, metro access networks (rings) are connected via a PE router to a local/edge data-center [6]. The local data-center is then connected to a data-center at a core node through a data-center interconnect (DCI) link. It is assumed that the edge data-center has a PE router connecting to it and the rest of the network, while at the core node is a P router, which for the sake of connecting to the PE router can potentially be replaced with the elements shown in the dotted-bubble. Inside the dotted bubble are multiple P2MP head-ends along with an Xponder. The P2MP head-end is connected to tail-ends in both the metro-edge, as well as at the PE router. The P2MP transceiver is back-connected to the Xponder which further aggregates multiples of 400Gb/s into 4x400Gb/s or 3x600Gb/s or 2x800Gb/s [3].

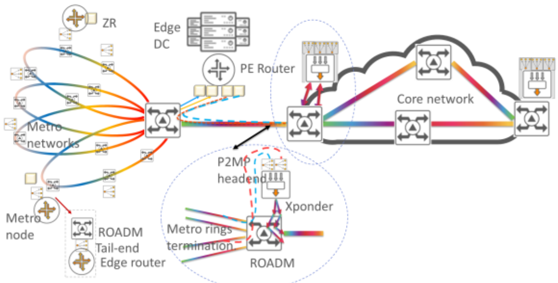

Fig. 3. Use-case of DCI with P2MP optics.

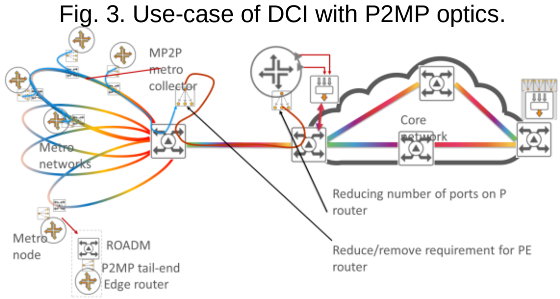

Fig. 4. DCI using P2MPoptics, router bypass.

Shown in Fig. 4 is a further elucidated version of the architecture in Fig. 3, in which, a P router now is made available at the core node, and which supports towards the metro-edge the head-end of the P2MP, while at the core end, supports another head-end of the P2MP along with Xponders. Further, at the termination point of the metro-edge (ROADM), standalone P2MP head-ends are made available thereby absolving the need for the PE router. These P2MP head-ends are also connected in a transceiver-to-transceiver fashion to the P2MP head-end at the core node (plugged into the P-router) (thereby further reducing the number of transceivers).

Modelling the Approaches

For purpose of comparison, we consider a synthetic network with 20 core nodes (average degree 2.6), at an average path length of 2125km (modelled as a US core), with the smallest link of 140km, and the largest link is of 4822km. Each core node subtends metro edge rings, each of which has between 5 and 10 nodes (average

7.25 nodes), with mean ring size of 120km. There are 240 metro rings resulting in 1740 metro-edge nodes. PE routers are of size 1.2Tb/s to 4.8Tb/s in increments of 1.2Tb/s. P routers are assumed to be of size 6.4Tb/s to 25.6Tb/s in increments of 2.4Tb/s. The Xponder [3] is of 1.6Tb/s capacity, with 2x800 or 3x600 or 4x400Gb/s output. P2MP head-ends are at 400Gb/s while tail-ends are either at 25Gb/s or 100Gb/s. Traffic is varied in 5 unequal instalments from 100Gb/s per metro-edge node to 800Gb/s per metro-edge node, resulting in combined traffic in the network at 174Tb/s to 1392Tb/s across the network. We ran the simulation for ten iterations at each load value and results are averaged.

Shown in Fig. 5 and Fig. 6 are the main results. In Fig. 5 we compare the number of transceivers required for the baseline case, the HBH case and the P2MP case as we vary the load across 5 data-points. P2MP in this case is deployed as 400Gb/s head-end and 100Gb/s tail-end. The number of P2MP transceivers required is on average 70.5% less than HBH and 57% lesser than baseline case. Even if we do not consider the impact of transceivers/Xponders in the core part of the network, there is still a 59.7% betterment compared to HBH and 51% improvement compared to the baseline case. The HBH case shows a peculiar trend. The initial transceiver count (ZR+) is very high, and then despite the (50%) increase in traffic (and thereafter increase of 33%, 40%), the comparative transceiver count does not increase proportionally. This is because, initial deployment of pluggables (and router ports) at 400Gb/s are sufficient to meet the traffic needs. Despite the initial increase in load, the comparative growth for HBH is not much due to available bandwidth in the ZR+ pluggables that are used.

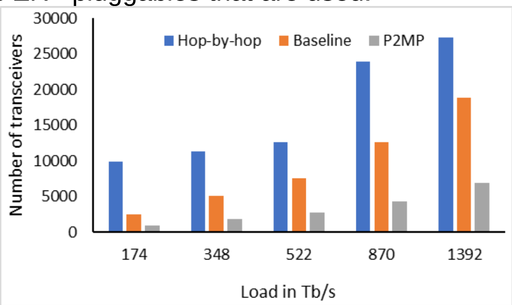

Fig. 5. Number of transceivers.

In Fig. 6 is the number of transceivers required for the CDN and CDN+DCI cases for baseline (using P2P ZR+) and P2MP. In both cases (CDN and CDN+DCI), the P2MP model has significant benefit in transceiver count compared to the baseline model. We do not consider HBH as its performance is restricted in DCI environments due to large number of router ports resulting in increased transceiver count. Compared to the baseline model for CDN, the P2MP model needs only 35% of the transceivers to provision the same traffic, while for CDN+DCI, P2MP requires 42% of the transceivers on average compared to the baseline case across the 5-load values.

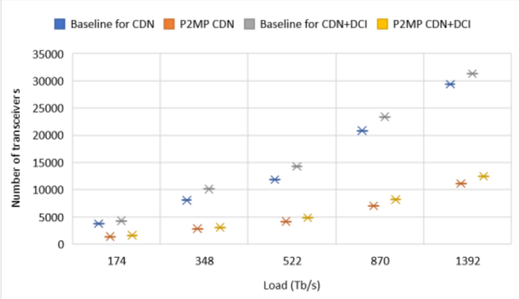

Fig. 6. Number of transceivers (CDN and DCI cases).

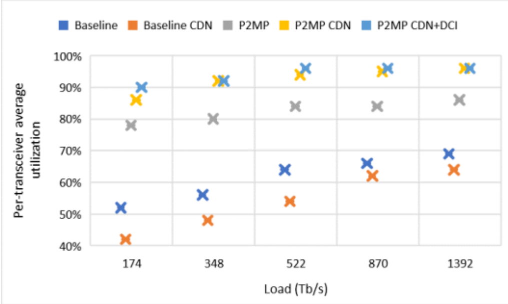

Fig. 7. Network-wide per-transceiver utilization.

In Fig. 7 we show the per transceiver utilization averaged across network-wide transceivers. The P2MP (82%), P2MP in CDN (91%) and in CDN+DCI (93%) result in far superior per-transceiver utilization compared to the baseline (ZR+) (61%) and baseline CDN cases (56%). We did not consider baseline in CDN+DCI. On average, there is 36% better utilization of transceivers when P2MP are used compared to baseline case in CDN. Note that the impact of P2MP is not just in the metro-edge to core part but also between data-centers. With baseline ZR+, the performance recedes as we consider CDN case, as we need to repeat the inclusion of upfront transceivers (and router ports) to provision such traffic. The key takeaway being that P2MP due to flexible bandwidth support (multiple ways to share 400Gb/s upstream DSCM channel) results in a more gradual model facilitating pay-as-you-grow.

Conclusions

We showcase P2MP as a solution for provisioning traffic across metro-edge and metro-core networks and compared with baseline (ZR+) solution with additional use cases of CDN and DCI. Results show significant benefit in transceiver count and utilization (also implying router port reduction) when P2MP is deployed.

References

[1] D. Welch et al., "Point-to-Multipoint Optical Networks Using Coherent Digital Subcarriers," in Journal of Lightwave Technology, vol. 39, no. 16, pp. 5232-5247, 15 Aug.15, 2021, doi: 10.1109/JLT.2021.3097163.

[2] R. Krishnan et al., "Moving Beyond End-to-End Path Information to Optimize CDN Performance," Internet Measurement Conference (IMC), ACM, Chicago, IL (2009), pp. 190-201

[3] H. Sun et al., "800G DSP ASIC Design Using Probabilistic Shaping and Digital Sub-Carrier Multiplexing," in Journal of Lightwave Technology, vol. 38, no. 17, pp. 4744-4756, 1 Sept.1, 2020, doi: 10.1109/JLT.2020.2996188.

[4] A. Napoli *et al*., "Live Network Demonstration of Point-to-Multipoint Coherent Transmission for 5G Mobile Transport over Existing Fiber Plant," *2021 European Conference on Optical Communication (ECOC)*, 2021, pp. 1-4, doi: 10.1109/ECOC52684.2021.9605901.

[5] J. Bäck et al., "CAPEX savings enabled by point-to-multipoint coherent pluggable optics using digital subcarrier multiplexing in metro aggregation networks," in Proc. Eur. Conf. Opt. Commun., 2020, pp. 1–4

[6] L. Peterson *et al*., "Central office re-architected as a data center," in *IEEE Communications Magazine*, vol. 54, no. 10, pp. 96-101, October 2016, doi: 10.1109/MCOM.2016.7588276.

[7] J. Back, et al., "A Filterless Design with Point-To-Multipoint Transceivers for Cost-Effective and Challenging Metro/Regional Aggregation Topologies", in Proc. ONDM 2022.