Cost-effective IP-over-DWDM Aggregation and Improved Router Bypass using P2MP Optics

Ashwin Gumaste[(1)], Antonio Napoli[(2)], João Pedro[(3)], Walid Wakim[(1)], Harald Bock[(2)]

[(1)] Infinera Corporation USA, agumaste@infinera.com; [(2)] Infinera, Germany, [(3)] Infinera, Portugal

Abstract *We explore point-to-multipoint (P2MP) based solution to reduce hardware requirements in high-bandwidth access and supporting metro IP over DWDM networks. We evaluate over hardware count, power and latency and show comparative benefit. ©2023 The Author(s)*

Introduction

P2MP in access networks is the de facto method of implementing broadband access to residential and small/medium enterprises. The OLT terminates at the edge of the metro network, and then the backhauled metro traffic is taken to Provider (P) routers or data-centers across the metro network using point-to-point high density circuits. These high-density circuits in the metro require coherent optics to transport data all the way to the access. A typical operator network consists of a core network, usually in the form of a sparse mesh. Each node in the core subtends one or more metro rings, with each ring having 5 to 8 nodes. Each metro node further subtends access rings, such that an access ring of up-to15 nodes is connected to a metro ring in either a single or dual homed fashion. It is advantageous to have a dual homed connection from a protection perspective, so that a single node loss in the metro does not have impact on the 10s of access nodes (through multiple rings connected at the node). Typically, the metro node has a Provider Edge (PE) router, equipped with coherent pluggables to facilitate communication to the edge access nodes. The granularity of access traffic doubles every other year, implying that the technology choice in the access and at the metro nodes is crucial. Conventional passive optical networks with OLTs at the metro edge and ONUs at access nodes are now being replaced due to the bandwidth explosion with higher end PON technology such as 50G PON. However, even with these technologies, the requirement of almost 100Gb/s per access node (in order to support 100Mb/s per dwelling), implies that there is a need for a higher end coherent optical access solution [1]. One approach is to use point-to-point (P2P) coherent optics embedded in edge routers at metro nodes, and edge routing/Ethernet aggregation infrastructure at access nodes. Recent developments in coherent pluggables, such as the 100G ZR, facilitate this approach, which has gathered significant attention under the gamut of IPoDWDM to the edge [7]. Though coherent pluggables are available in multiple line-rates such as 100Gbs/s, 200Gb/s and 400Gb/s (with 800Gb/s on the horizon) [8], for concerns of cost, it makes sense to only put in 100Gb/s coherent modules (such as 100Gb/s ZR) from the access to the metro nodes. The idea is to have PE routers at metro nodes aggregate traffic from access rings in 100Gb/s slices and map them on to 400Gb/s and above line rates within the metro networks, as illustrated in Fig. 1 towards core data-centers.

A second competing approach, depicted in Fig. 2, is based on a recent advance in point-to-multipoint (P2MP) technology that leverages the efficient multiplexing of digital subcarriers (as opposed to TDM muxing in PONs). This digital subcarrier multiplexing (DSCM) based approach is characterized by the ability to have a head end as high as 400Gb/s and tail ends that can support 25Gb/s slices (for a total of 16 subcarriers) [4]. The idea of aggregating traffic using digital subcarriers is quite promising – can lead to reduction in IP router ports, as well as reduction in the overall hardware (transceiver count). Importantly, the feasibility of DCSM-based P2MP technology has been demonstrated experimentally and in field trials [5]. Using P2MP, it is possible to run connections between the metro-core network and access nodes using optical grooming through DSCM technology.

This paper investigates and compares the utilization of P2MP pluggables versus low cost P2P pluggables to fulfil the capacity requirements in metro and access networks. The results obtained in two realistic networks suggest that P2MP optics can lead to a reduction in hardware count, power and cost.

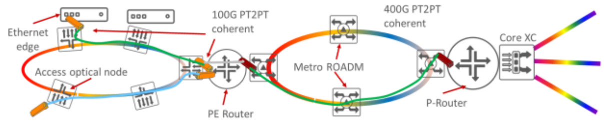

Fig. 1. Architecture using point-to-point pluggable in IP over DWDM access and metro networks.

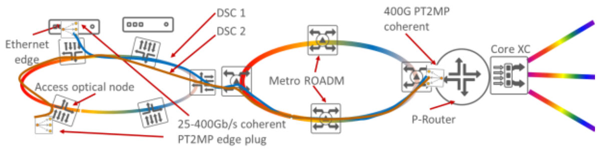

Fig. 2. Architecture using point-to-multi-point pluggable in IP over DWDM access and metro networks.

System Design and Optimization

We now describe the architecture of the two approaches that we consider and delve on the optimization techniques to reduce overall hardware count of both solutions. Shown in Fig. 1 is the P2P solution that uses coherent 100G (in Ethernet aggregation switches) in the access, a PE router at the access-metro edge and a P router at the core sites. Shown in Fig. 2 is the P2MP solution with edge based coherent (DSCM supporting) nodes in Ethernet aggregation switches, optical equipment in the metro (same as the earlier solution) and P routers at the edge of metro and core networks, that have P2MP head end pluggables. Our goal is to compute the hardware count, power and utilization values of both solutions. To this end, we built an optimization algorithm that facilitates efficient grooming of traffic in both network types as to determine the lowest hardware count in each solution.

We assume traffic demands of different granularity at the edge of the network consisting of 1G, 10Gs at the edge, few 100Gs at the metro sites, TDM leased lines and backhaul traffic consisting of both STM-x, fiber channel (over Ethernet), and IP VPNs.

We developed a constrained optimization model whose objective function is to minimize hardware count (transceivers) in the network based on the two approaches. The P2P approach is in some sense analogous to hop by hop routing [], in IP over DWDM networks, where at a node along the path, where there is a router, the router is leveraged to act as a traffic aggregator as well as regenerator. The aggregation function is critical in the case of the P2P coherent optical pluggable based solution as it reduces the number of wavelengths in the metro ring. While the objective function is to map the traffic with the minimum number of transceivers, constrained to the following: The number of wavelengths per fiber in the access and metro are fixed (48); the capacity of wavelengths can support either 100G, 200G or 400G through pluggables. The traffic is given, and all of it must be provisioned on wavelengths with fixed capacity. When a router is deployed, it adds latency to the solution, and the latency is proportional to the load at the interfaces of the router. If the solution is all-optical i.e. from the access to the metro-core P-router, then there is the constraint of wavelength continuity across the access and core networks. The total capacity of the pluggables at a router must not exceed that of the router, with the assumption that the router is deployed in a non-blocking cross-connect function. Demands are to be aggregated in an efficient way (combinatorial choices), such that they are either co-routed to the destination, or fill up wavelengths of fixed capacity (100/200/400Gb/s). Regenerators are placed based on a reach table presented in [2] and [6] for the P2P and P2MP cases. The reach depends on the average span lengths and number of interim ROADMs traversed.

Numerical Evaluation

We evaluated the P2P and P2MP approaches over two networks – a US (network 1) and an European one (network 2) (metro part as shown in Fig. 3). We developed a Python-based optimization model using *networkx* and convex optimization (*CVXPY* and *numpy*) libraries. Network specifics including traffic details are as per the table below. We increased the traffic in the two networks based on load ranges shown in row 2 in the table, and then normalized this load on to 17 discrete steps represented as 0-0.85 (in 0.05 increments). To provision the traffic, we assumed Ethernet edge switches and P routers for both solutions, as well as PE routers for the P2P solution (at metro-access interface). We assumed multi-terabit P and PE routers that are capable to support 400Gb/s ports and consume 50W of power per port. The pluggable 100G P2P and the 400G P2MP modules consume 5W and 18W, respectively. Our goal is to compute overall power of the IP over DWDM solution. We varied the traffic as shown below and made measurements that are presented in Fig 4-7.

Table 1:Network Specifics.

Metric	Network 1	Network 2
Load range (Gb/s)	1200-4400	700-2300
Number of access nodes	64	52
Average path length (km)	112	84
Nodal degree	2.14	2.18
Average line rate (Gb/s)	44.8	32.1

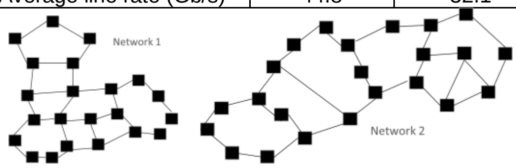

Fig. 3: Metro part of the two network architectures.

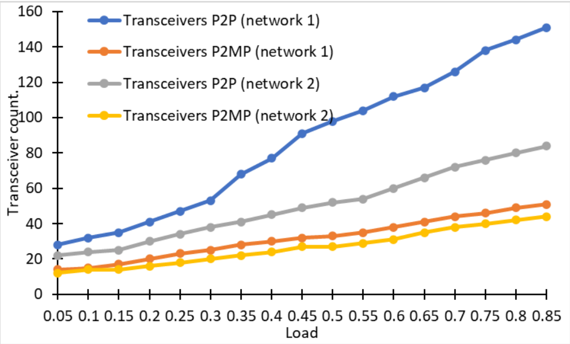

Fig. 4: Transceiver count for the two networks.

Depicted in Fig. 4 is the number of transceivers for the two solutions across the two networks. This is the main result – the P2MP solution requires 62% less transceivers than the P2P solution in network 1, and 49% less in network 2. This enormous transceiver saving is resultant of optical layer grooming due to sub-carriers as opposed to electronic grooming in P2P. An interesting observation is that the reach of the P2P solution is not a factor as lightpaths between metro edge and access have to be terminated at the PE router, to avoid wavelength exhaustion in the metro. This result not only quantifies the benefit of P2MP solution, but stresses upon the fact that for supporting high bandwidth access (with coherent technology), the only low cost way to do so is using P2MP and optical grooming.

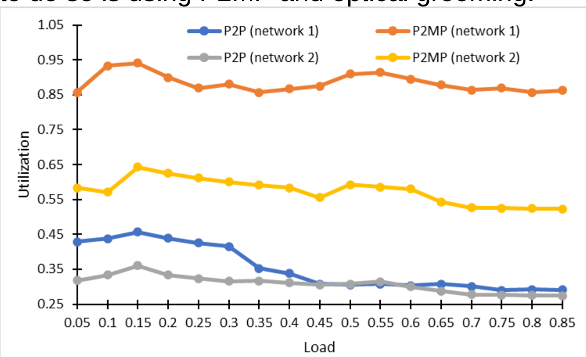

Fig. 5: Average utilization of transceivers.

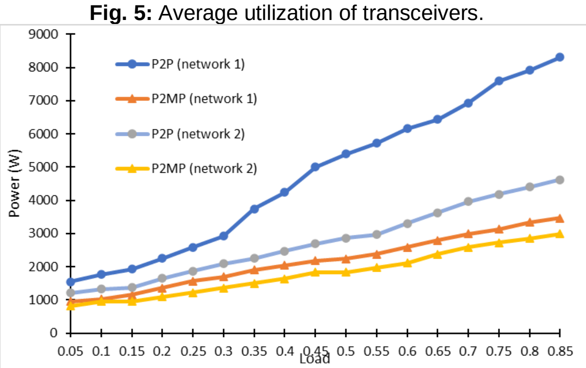

Fig. 6: Power of the IP over DWDM solution.

Shown in Fig. 5 is the average utilization of the transceivers. Due to significant reduction in transceiver count, the utilization of P2MP transceivers is higher than P2P transceivers. While one may view this as opportunity for more growth in the P2P case, most access networks have higher over provisioning implying limited scope for further traffic engineering.

Fig. 6 shows the power consumption of the overall IP over DWDM solution. In network 1, the power savings of the overall P2MP solution (less the common parts such as line system, edge Ethernet switches) is 54% versus the P2P solution, while in network 2, the power saving is 34%. The key aspect to note is that despite the P2MP plugs consuming more power than the P2P plugs, due to the fact that they are efficient in optical layer grooming, the overall power consumption reduces. Reduction is due to: (1) less plugs required and (2) less router ports required to house these plugs.

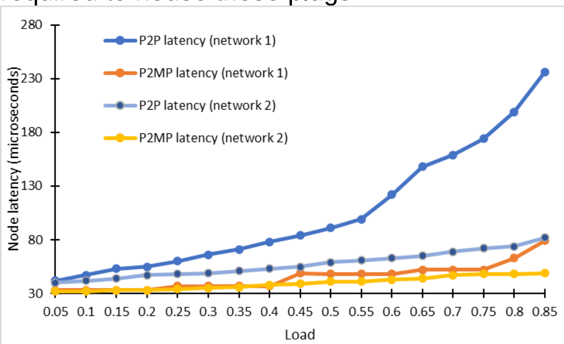

Fig. 7: Edge to edge latency through the router.

Since much of the access traffic is video or backhaul, latency through the network becomes important. While the propagation delay is common to both solutions and hence the only variable latency component is through the plugs and routers. To this end, we plotted average per-packet latency of mid-sized packets of length 250 bytes. Latency is computed by adding up P2P and P2MP pluggable values in [9] and [3, 10] and router queueing impact as a function of load using the base values in [11]. What we observe from Fig. 7 is that for lower loads, the difference in latency across the two technologies is not much, but as load increases, the use of the router has a pronounced impact and skews the results against the P2P approach. Key to note is that the almost exponential increase latency in the P2P case results in inordinate service delay at the network edge, which can be allievated with P2MP optical grooming.

Conclusions

We have shown benefits of P2MP in metro and access networks over two network topologies, and across measures of transceiver count, power consumption, transceiver utilization and latency. An almost 45.5% betterment in transceiver count, and 44% betterment in overall power usage is observed, implying a clear case for P2MP in next-generation coherent metro-access networks.

References

[1] "IEEE Draft Standard for Ethernet - Amendment: Physical Layers and Management Parameters for 100 Gb/s Operation over DWDM (dense wavelength division multiplexing) systems," in *IEEE P802.3ct/D3.1, January 2021* , vol., no., pp.1-147, 19 Jan. 2021.

[2] A. Gumaste, M. Sosa, H. Bock and P. Kandappan, "Optimized IP-over-WDM core networks using ZR+ and flexible muxponders for 400 Gb/s and beyond," in *Journal of Optical Communications and Networking*, vol. 14, no. 3, pp. 127-139, March 2022, doi: 10.1364/JOCN.444943.

[3] R. R. Kompella, K. Levchenko, A. C. Snoeren and G. Varghese, "Router Support for Fine-Grained Latency Measurements," in *IEEE/ACM Transactions on Networking*, vol. 20, no. 3, pp. 811-824, June 2012, doi: 10.1109/TNET.2012.2188905.

[4] D. Welch, et al., "Digital subcarrier multiplexing: enabling software-configurable optical networks", IEEE/OPTICA Journal of Lightwave Technology, vol. 41, no. 4, pp. 1175-1191, February 2023.

[5] A. Napoli, et al., "Live network demonstration of point-to-multipoint coherent transmission for 5G mobile transport over existing fiber plant", in Proc. of ECOC 2021, Bordeaux, France, September 2021, paper Th1F.3

[6] J. Bäck *et al*., "A Filterless Design with Point-to-multipoint Transceivers for Cost-Effective and Challenging Metro/Regional Aggregation Topologies," *2022 International Conference on Optical Network Design and Modeling (ONDM)*, Warsaw, Poland, 2022, pp. 1-6, doi: 10.23919/ONDM54585.2022.9782860.

[7] D. Tauber *et al*., "Role of Coherent Systems in the Next DCI Generation," in *Journal of Lightwave Technology*, vol. 41, no. 4, pp. 1139-1151, 15 Feb.15, 2023, doi: 10.1109/JLT.2023.3235820.

[8] OpenZR+ MSA (online): https://www.openzrplus.org/site/assets/files/1075/openzrplus_1p0.pdf

[9] Specifications sheet: https://www.cisco.com/c/en/us/solutions/collateral/silicon-one/datasheet-c78-744311.html

[10] M. A. Kinsy and S. Devadas, "Low-overhead hard real-time aware interconnect network router," *2014 IEEE High Performance Extreme Computing Conference (HPEC)*, Waltham, MA, USA, 2014, pp. 1-6, doi: 10.1109/HPEC.2014.7040976.

[11] G. Appenzeller, I. Keslassy and N. McKeown, "Sizing router buffers," ACM SIGCOMM Computer Communication Review Volume 34Issue 430 August 2004pp 281–292https://doi.org/10.1145/1030194.1015499