Hardware Comparison of Xponders and ZR+ in Metro and Core Networks with Mixed IP and OTN Traffic

Ashwin Gumaste[1], Joao Pedro[2], Paul Momtahan[3] and Harald Bock[4]
[1]Infinera Corporation San Jose, CA USA, [2]Optical Systems Group, Infinera Unipessoal Lda, Portugal,
[3]Infinera London, UK and [4]Infinera GmbH Munich Germany
agumaste@infinera.com, jpedro@infinera.com, pmomtahan@infinera.com hbock@infinera.com

Abstract: We evaluate the role of transponders-muxponders (Xponders) and ZR+-pluggable interfaces in metro and core networks beyond 400Gb/s across 3-metro and 3-core topologies. The stochastic study computes hardware-count and overbuilds with Xponders resulting in lowest counts. © 2022 The Author(s).

1. Introduction

With the advent of long-reach pluggable interfaces, such as the OpenZR+ and OpenROADM MSAs, there are new possibilities at 400Gb/s+ rates in metro and core networks. Now, OpenZR+ interfaces, henceforth referred as ZR+, can be plugged into router ports, enabling efficient realization of IP-over-WDM network architectures. The question we want to answer in this paper is how do ZR+ pluggable devices compare with embedded optics, i. e. multi-rate transponder-muxponder (Xponder) solutions in applications ranging from metro to core networks. To this end, we consider three architectural options, two of which use the ZR/ZR+ standard, while the third uses Xponders. We study the impact of these architectures and technologies using a constrained optimization model over 6 different network types – 3 core networks, and 3 metro regional networks with stochastic traffic mixes that constitutes a mix of IP and private leased lines (OTN) traffic. The three architectures we consider are: Hop-by-Hop (HBH) [1], the ZR+ based optical ROADM (ZR+ROADM) [2] and an Xponder based ROADM solution [4]. In the HBH scheme, a two-layer network has routers connected to an optical mux/demux (typically AWG-based to minimize line system cost), with ZR/ZR+ plugged into a router, and a path *A-B-C* is characterized by a lightpath being dropped and processed by a router at intermediate node *B*, and then sent back into the fiber to node *C*. HBH, does not require adherence to wavelength continuity, as the wavelength in any two adjacent links of a lightpath can be different due to electrical regeneration at every interim node. This implies that the optical budget is also less of a concern, as long as no two nodes are beyond the reach of a ZR+ (typically 800km depending on fiber type and intra-link span-length). In the ZR+ROADM case, the ZR+ is plugged into an ingress router, which is further connected to a ROADM such that at an intermediate node, if optical performance allows, then the lightpath expresses (optical bypasses) through the interim node. In the Xponder case, routers are used only at the end-nodes of a lightpath, and traffic from a router is connected to an external Xponder, via grey optics at any convenient rate, thus reducing the stress on routers to support 400Gb/s pluggables. The Xponders are connected to ROADMs at ingress and egress with ROADMs used for optical bypass at intermediate nodes. When the lightpath requires intermediate regeneration, it is optically dropped onto an Xponder, which regenerates the signal, without the need of an IP router. In this paper, we assume the ZR+ and Xponders to support multiple line-rates via the use of various modulation schemes and symbol rates. Note related works [1-3] while focusing on ZR+ and Xponders, do not consider the holistic impact of stochastic traffic, impact of regenerators, and OTN traffic, that we do in this paper.

2. Architectural Aspects and Traffic Modelling

The three architectures are shown in Fig. 1a, b and c and a reach-table is shown in Fig. 1d. The reach-table values were obtained assuming typical fiber properties, commercial device specifications, EDFA compensated spans, and using a simplified optical performance model fitted with suitable system and operational margins. In the HBH case, traffic arrives from aggregation nodes to a router via grey optics. The router maps traffic onto appropriate ports at the ingress node. The router supports ZR or ZR+ (ZR for sub-120km, else ZR+). The ZR/ZR+ interface is connected to an AWG from which the ingress traffic enters the fiber plant for transmission to the next node. Note that at each node, every lightpath is *mandatorily* dropped, processed through the ZR/ZR+-router-ZR/ZR+ set of elements and then fed back into the network. In this case, a ROADM is not required and simpler (i.e., lower cost) AWG-based fixed filters can be deployed for multiplexing/demultiplexing optical channels. In the case of ZR+ROADM, traffic at an ingress node enters a router through grey optics, and the router maps these to an appropriate ZR+ interface that is plugged into the router. The lightpath is then transmitted by the ZR+ into a multi-degree ROADM. At intermediate nodes, if the lightpath does not have to be regenerated, then the ROADM expresses it to its downstream neighbour, and the process continues. However, if the lightpath is to be regenerated at an intermediate node, then the ROADM at that node drops the lightpath onto a router, which uses a ZR/ZR+ interface to receive the signal. The router at this intermediate site then regenerates the signal and uses another

ZR/ZR+ interface to insert the signal back into the fiber plant. In the Xponder case, traffic arrives at an ingress router, through grey optics at low-rates. This is then mapped onto grey 100/400Gb/s optics which are connected to an Xponder. The Xponder multiplexes lower-line rates from the router onto a high-speed, high baud-rate wavelength connected to a ROADM. The lightpath passes-through interim nodes, and when regeneration is required, it is done exclusively with Xponders (no router required). Since the Xponder can be a muxponder, it can also drop lower-rate (in multiples of 100Gb/s) traffic at regen sites. The Xponder supports multiple line-rates from 100Gb/s to 800Gb/s by changing the baud-rate and modulation. The reach-table (Fig. 1d) shows the performance of ZR, ZR+ and Xponder at different line-rates and for different number of interim ROADMs. The actual reach-table utilized has a finer granularity, but space restrictions limit the number of configurations shown.

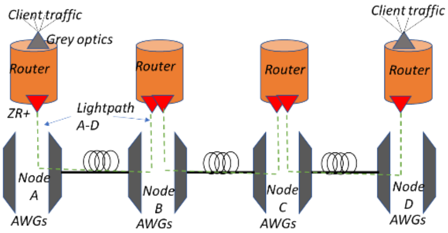

Fig. 1a. Hop-by-Hop (HBH) Optical Network.

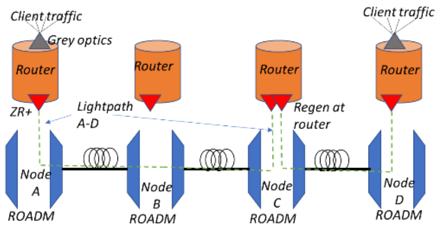

Fig. 1b. ZR+ ROADM.

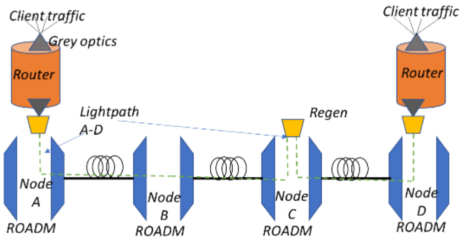

Fig. 1c. Xponder-based solution.

	ZR 400Gb/s	ZR+ 400Gb/s	Xponder @400Gb/s	Xponder @600Gb/s	Xponder @800Gb/s
Reach in km for 0 interim ROADMs	120	960	6000	2620	720
Reach in km for 4 ROADMs	0	880	5680	2320	400
Reach in km for 8 ROADMs	0	720	5360	1920	0
Reach in km for 12 ROADMs	0	0	5040	1600	0
Reach in km for 16 ROADMs	0	0	4720	1200	0
Reach in km for 20 ROADMs	0	0	4400	0	0

Fig. 1d. Example reach table for ZR+ and Xponders.

Network name	Nodes	Links	Max link (km)	Min link (km)	Avg Span (km)	Nodal degree	Average hop	Average lightpath length (km)
MN1	12	18	650	120	60	1.5	2.56	660
MN2	14	20	842	180	70	1.42	2.89	820
MN3	16	25	1210	210	70	1.56	3.12	910
CN1	34	96	4590	242	80	2.82	3.32	1824
CN2	80	128	5820	245	80	1.6	3.6	2356
CN3	100	140	6400	281	80	1.4	3.98	2480

Table 1: Specifics of the networks.

Network name	100Gb/s	200Gb/s	400Gb/s	OTU-4	OTU-2	Year 1 (Tb/s)	Year 3 (Tb/s)	Year 5 (Tb/s)
MN1	24-96	12--48	24-96	2--8	10	10.38	36.7	58.5
MN2	14-56	28	28-70	28	56-112	17.78	32.18	43.12
MN3	48-96	32-48	64-128	32-64	0	40	58.4	76.8
CN1	68-340	17-170	68-306	102	136	48.96	132.96	201.96
CN2	160-480	80-240	240-480	160	80	144.8	224.8	304.8
CN3	200-400	200-400	300-900	500-700	200	232	392	552

Table 2: Stochastic traffic models.

3. Traffic Provisioning and Optimization Model

Table 1 shows the key properties of the networks that we consider, whereas Table 2 presents a 5-year traffic model. We assume 5 types of traffic – 100GbE, 200GbE, 400GbE, OTU-2 and OTU-4. 400GbEc dominates the rest of the traffic (P-router-to-P-router, inter-data-center). Metro network cases (MN1, MN2, MN3) traffic increases from 10-to-58Tb/s (MN1), 17-to-43Tb/s (MN2) and 40-to-76Tb/s (MN3). For core networks (CN1, CN2, CN3) the traffic increases from 48-to-201Tb/s (CN1), from 144-to-304Tb/s (CN2) and from 232 to 552Tb/s (CN3). Traffic increases approximately 2x every other year. The goal is to provision traffic on high-capacity optical channels (400Gb/s and beyond). In the case of HBH and ZR+ROADM, the maximum capacity per optical channel is 400Gb/s, while Xponders are for this study assumed to support 400, 600, 800Gb/s optical channels, though other interim rates (like 700Gb/s) are also possible. We provision traffic using a constrained optimization model. The objective of the mixed integer linear program (MILP) is to minimize the number of regenerators, as these are: (a) the primary variable components directly impacting network cost; and (b) placement of regenerators have a strong impact on all the lightpaths passing through a fiber. We observe that minimizing regenerators leads hardware count minimization. To this end, the constrained optimization model takes as inputs architecture and provisioning constraints (Section 2), wavelength continuity for ZR+ROADM and Xponders cases, wavelength and fiber degree constraints, and reach-table constraints. The MILP takes an additional input: possible regenerator sites (nodes) and determines their optimal locations. To make the solution tractable in time (given the huge combinational possibilities in site selection), we restrict the number of nodes that the MILP considers to 6 in each iteration, with a max of 3 iterations, thereby resulting in speedup of the MILP to less than 5 minutes with a 99.92% accuracy. The objective function of minimizing regenerators is also subject to a provisioning constraint of provisioning every traffic request. For the Xponders case, we induce an additional line-rate constraint, whereby, for a traffic request that is in multiples of 400Gb/s, the LP *attempts* to provision this onto 600Gb/s or 800Gb/s containers if the reach tables' constraint is satisfied, else on 400Gb/s lightpaths. Spectrum conversion across a lightpath through a regenerator is allowed to improve fiber utilization. OTN traffic in the case of HBH and ZR+ROADMs is mapped directly onto exclusive OTN-specific ZR+ modules, while in the case of Xponders, OTN traffic at OTU-4 is directly mapped onto an Xponder. The Xponder can subsume up to 16x100Gb/s clients mapped to 2-line-side optical channels, in any combination of OTU-4 and 100GbE.

4. Results

We built a Python-based simulation model that calls the MILP optimization module that runs on an Intel core i7 9^{th} gen processor. Results are plotted in Fig. 2-4. Each figure has 6 load values, corresponding to a different network (see Table 1). From Table 1, it can be expected that Xponders' superior reach+higher line-rate leads to the least number of regens. In fact, only in one network (CN3), does Xponders need regens, while HBH, and ZR+ROADMs do require sizable number of regens/mandatory regens, due to ZR/ZR+ reach limitations. Fig. 2 plots the number of transceivers required for the three architectures, whereas Fig. 3 presents the number of 400Gb/s router ports used. In the metro cases, Xponders needs 78% fewer transceivers than HBH and 54% less than ZR+ ROADM averaged over 5 years, while in core networks, Xponders needs 79% less transceivers than HBH, and 67% than ZR+ over 5 years. In case of router ports used, for metro networks, Xponders needs 66% less than HBH, and 55% less than ZR+ROADM, while for core networks Xponders need 63% less than HBH and 44% less than ZR+ROADM. The performance improvement in core can be attributed to the better reach and spectral performance of Xponders, which reduces or even avoid the need of regenerators, but it may seem counter-intuitive that in metro networks, Xponders continues to perform well. The reason is the ability to exploit transmission at 800Gb/s and 600Gb/s line-rates, whenever feasible. Also, since a router cannot mux IP and OTN traffic onto the same ZR+ interface, a number of lightpaths running OTN are effectively at sub-400Gb/s. This aspect and the spectral efficiency of Xponders leads to reducing not only the number of transceivers and router ports, but also a better fiber (wavelength) utilization as shown in Fig. 4 for both metro and core networks.

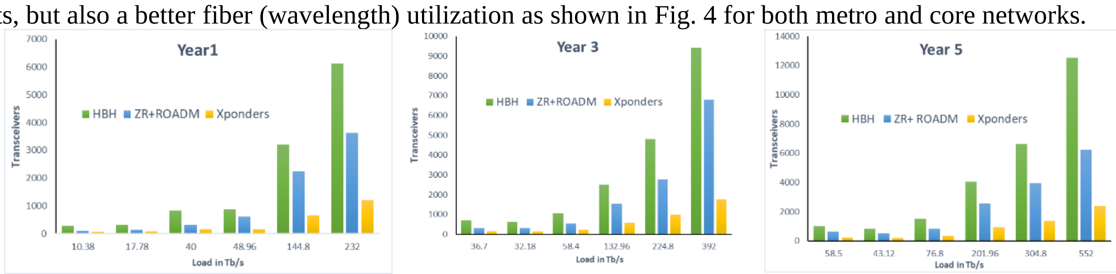

Fig. 2. Transceiver count over a five-year period.

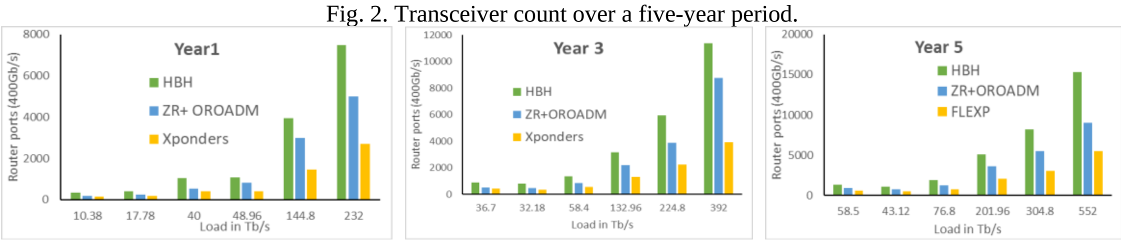

Fig. 3. Router port count over a five-year period.

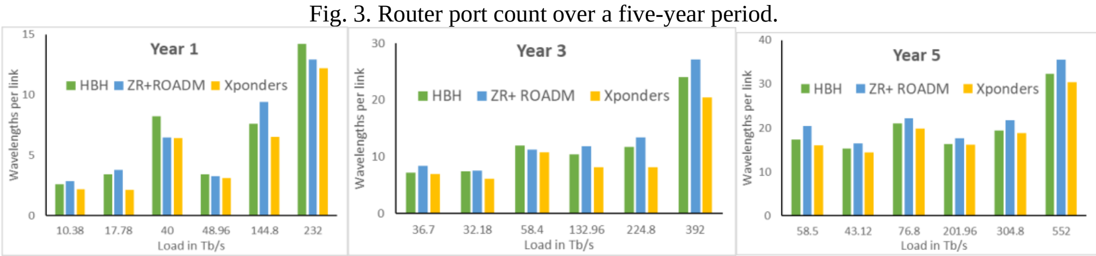

Fig. 4. Wavelengths used per-link over a five-year period.

5. Conclusion

We have compared hop-by-hop (HBH), ZR+ ROADMs and Xponders as plausible IP-over-WDM architectures for mixed traffic over three metro and three core networks. The value of using Xponders is clearly visible from an equipment, spectral efficiency and fiber utilization perspective in both metro and core networks.

References

[1] P. Wright, R. Davey, and A. Lord, "Cost model comparison of ZR/ZR+ modules against traditional WDM transponders for 400G IP/WDM core networks," 2020 European Conference on Optical Communications (ECOC), (IEEE, 2020), pp. 1–4.

[2] S. Melle, T. Zami, O. Bertran-Pardo, and B. Lavigne, "Comparing IP-optical architectures & WDM transport technologies in metro, regional and long-haul networks," in Optical Fiber Communication Conference, (Optical Society of America, 2021), pp. M3E–2.

[3] T. Zami, B. Lavigne, and O. B. Pardo, "Benchmarking of opaque versus transparent core WDM networks featuring 400ZR+ QSFP-DD or CFP2 interfaces," 2020 European Conference on Optical Communications (ECOC), (IEEE, 2020), pp. 1–4.

[4] J. Pedro, N. Costa, and S. Sanders, "Scaling regional optical transport networks with pluggable and integrated high-capacity line interfaces," OFC 2021M3E.1.