Real-Time Point-to-Multipoint for Coherent Optical Broadcast and Aggregation – Enabled by Digital Subcarrier Multiplexing

Amir Rashidinejad[1*], Amin Yekani[1], Tobias A. Eriksson[2], Antonio Napoli[3], Robert Maher[4], Aditya Kakkar[1], Vince Dominic[4], Thomas Duthel[3], Mark Missey[4], Parmijit Samra[4], Don Pavinski[4], Peter Evans[4], Warren Sande[4], Mehdi Torbatian[1], Chris Fludger[3], Han Sun[1], Mehrdad Ziari[4], Fady Masoud[1], Azmina Somani[1], Dave Welch[4]

[1]Infinera Canada, [2]Infinera Sweden, [3]Infinera Germany, [4]Infinera USA

[*]arashidinejad@infinera.com

Abstract: We report on the first real-time operation of coherent point-to-multipoint in high-speed fiber-optic communications. The broadcast and aggregation network consists of a 400 Gb/s hub transceiver achieving post-FEC error-free communication with 4×100 Gb/s leaf nodes, 5 – 50 km away. © 2022 The Author(s)

1. Introduction

Optical networks are constantly evolving, with technologies such as the Internet of Things, edge computing, and 5G/6G introducing new requirements at a very fast pace [1]. In addition to the unprecedented demand for increased capacity, service providers are looking for innovative solutions that help reduce the cost of transport, power consumption, latency, and equipment count, together with enhanced scalability and increased service flexibility. This is especially relevant in the fastest growing sectors of telecom, i.e., metro/access and mobile front-haul, where Intensity Modulation Direct-Detection (IM-DD) and coherent compete. Comparatively, coherent provides an order of magnitude more capacity, enables deployment over different fiber types and impairments, supports greater scalability and troubleshooting, and allows for a higher level of service aggregation [2]. Both IM-DD and coherent solutions, however, are mostly implemented in a Point-to-Point (P2P) fashion, that not only requires several aggregation stages compared to an equivalent Point-to-Multipoint (P2MP) counterpart [3, 4], but are also poorly matched to the intrinsically asymmetric traffic patterns of such networks [5].

Digital Subcarrier Multiplexing (DSCM) has been proposed for many applications such as: nonlinear mitigation, dynamic bandwidth allocation, and ROADM filter penalty mitigation [6–9]. In [5, 10], we further proposed to employ DSCM to enable coherent P2MP and recently carried out a simple proof-of-concept experiment [11]. In this work, we successfully demonstrate this concept with the first real-time coherent P2MP transmission/detection over an optical broadcast and aggregation network, where a 400G hub establishes simultaneous communication to 4×100G leaf nodes, located at distances ranging from 5 to 50 km. We show real-time error-free transmission after Forward Error Correction (FEC) decoding in transmission in both uplink and downlink directions.

2. Coherent P2MP Enabled by DSCM

Fig. 1 illustrates the 4×100G ⇔ 400G P2MP experimental setup for the uplink (a) and downlink (b) directions. This scenario resembles a typical metro/access use-case, where a myriad of flows, stemming from different remote sites, are passively aggregated in uplink, or broadcast to in the downlink direction. The experimental setup consists of four 100 Gb/s leaf nodes and a single 400 Gb/s hub transceiver, connected via four SMF-28 spools of 5, 10, 25, and 50 km. Each node is comprised of XR400 prototype transceivers, including real-time DSP/FEC ASIC and integrated Transmitter-Receiver Optical Sub-Assembly (TROSA), housing the PIC and controls circuitry. Each transceiver is capable of generating up to 16 independent Nyquist digital Subcarriers (SCs), each at ∼4 GBd with dual-polarization 16-QAM modulation, resulting in a net datarate of 25 Gb/s per SC. Pictorial representations of the optical spectra for each leaf are shown as insets in Figs. 1(a)–(b), depicted in different colors for each leaf.

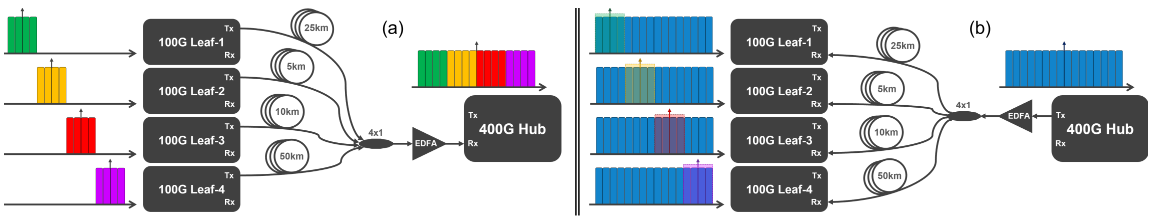

Fig. 1. Setup for 4×100G ⇔ 400G point-to-multipoint: (a) Uplink, (b) Downlink.

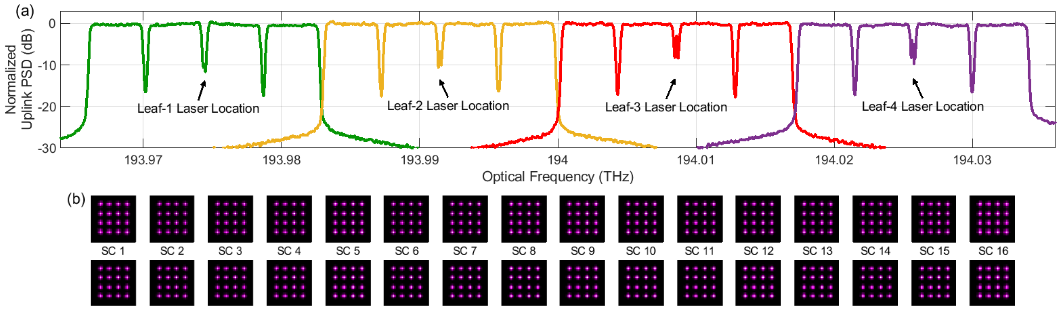

Fig. 2. Real-time measurements for 4×100G ⇔ 400G P2MP uplink: (a) Received optical PSD from each leaf, (b) 16-QAM constellations for each SC at hub (X- and Y-pol separated).

In the uplink direction, each leaf generates 4-SCs that after fiber propagation are passively combined with the flow from other nodes by a 4×1 optical coupler, followed by a single Erbium Doped Fiber Amplifier (EDFA) before arriving at the hub. Conversely, in the downlink direction, the hub generates 16 SCs that after amplification, are broadcast to all leaf nodes via a 4×1 optical splitter. However, each leaf only demodulates and acquires the data within its assigned SCs. Additional AES256 line-side encryption can be implemented to provide a further layer of privacy between leaf nodes. Each leaf's laser is shared between its Transmitter (Tx) and Receiver (Rx), and is automatically adjusted to wavelock to a 4-SC portion of the hub downlink signal [11]. To provide some guardband to laser drift and prevent spectral collision, the gap between SCs is set at ∼300 MHz, <8% of the SC bandwidth [11]. Furthermore, to compensate for path loss differences, the hub instructs all leaf nodes to adjust their transmit power such that a power-balanced Power Spectral Density (PSD) is observed after aggregation.

3. Results and Discussion

Fig. 2 shows the real-time results captured in the uplink/aggregation direction from the 4×100G ⇔ 400G P2MP setup of Fig. 1(a). Each 100 Gb/s leaf is set up to transmit 4 SCs at its laser frequency, at an offset to the hub laser. Fig. 2(a) depicts high-resolution optical spectrum analyzer measurements of these 4 SC-wide PSDs at the hub Rx originating from each individual leaf. The hub laser is located at 194 THz, while the four leaf lasers are wavelocked, based on DSP feedback, to the hub downlink signal at approximately −24 GHz, −8 GHz, +8 GHz, and +24 GHz offset, respectively. Furthermore, as can be observed from the PSD of Fig. 2(a), all SCs arrive at the hub at approximately equal power, confirming that leaf-to-leaf power balancing is performed as intended. The hub Rx has 16 independent clock recovery, carrier estimation, and polarization tracking circuits [5, 11] and can simultaneously demodulate the received data for all 16 SCs, resulting in a net 400 Gb/s post-FEC error-free uplink transmission. Fig. 2(b) illustrates the 16×2 constellation snapshots at the hub Rx, per SC and per X- and Y-polarization, captured over ∼10 minutes.

It should be noted that although the Rx DSP processes each SC separately, the utilized FEC engine uses gain-sharing, whereby symbols within a SC group are mixed, resulting in an averaging of the uplink performance across all four SCs. This can be seen from the hub Rx pre-FEC Q-factor readings of Tab. 1, reported per SC (averaged over both polarizations). As expected, due to FEC gain-sharing, the Q-factor for SCs originating from the same leaf is equal. Furthermore, even though all Q-factor values are very high and significantly far from the FEC limit of 6.4 dB, there is a small spread observed. This can be attributed to the slightly different transceiver SNR, different propagation distances, and ASE tilt experienced by each SC group. However, when comparing the OSNR margin to the FEC threshold for each leaf, this difference drops significantly down to only ∼0.15 dB.

Fig. 3 shows real-time measurements captured in the downlink/broadcast direction from the 4×100G ⇔ 400G P2MP setup of Fig. 1(b). In the downlink, the hub broadcasts all 16 SCs and, as described earlier, each leaf tunes its laser to the center of its designated SC group. Figs. 3(a)–(d) show the received PSD captures, as observed from each leaf's Rx DSP circuit. As expected, due to the frequency offset between each leaf and the hub lasers, the Rx PSDs are asymmetric, and the four designated SCs for each leaf are always to be detected at the four middle SC locations on the receiver side (SC 7–10), as highlighted in Figs. 3(a)–(d). For example, in this demonstration, leaf-1 is assigned the four left-most SCs from the hub downlink spectrum. As a result, it does not receive anything at its locations 1–6. The intended 100 Gb/s 16-QAM downlink traffic for leaf-1 is detected at SC locations 7–10, and although SCs 11–16 are also received by leaf-1, since they are assigned to other leaf nodes, they will not be processed. This can also be observed from the Rx constellation snapshots for leaf-1, depicted in Fig. 3(e).

Table 1. Measured uplink pre-FEC Q-factor per-SC at the hub Rx.

SC #	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Q (dB)	10.2	10.2	10.2	10.2	10.6	10.6	10.6	10.6	10.5	10.5	10.5	10.5	9.8	9.8	9.8	9.8

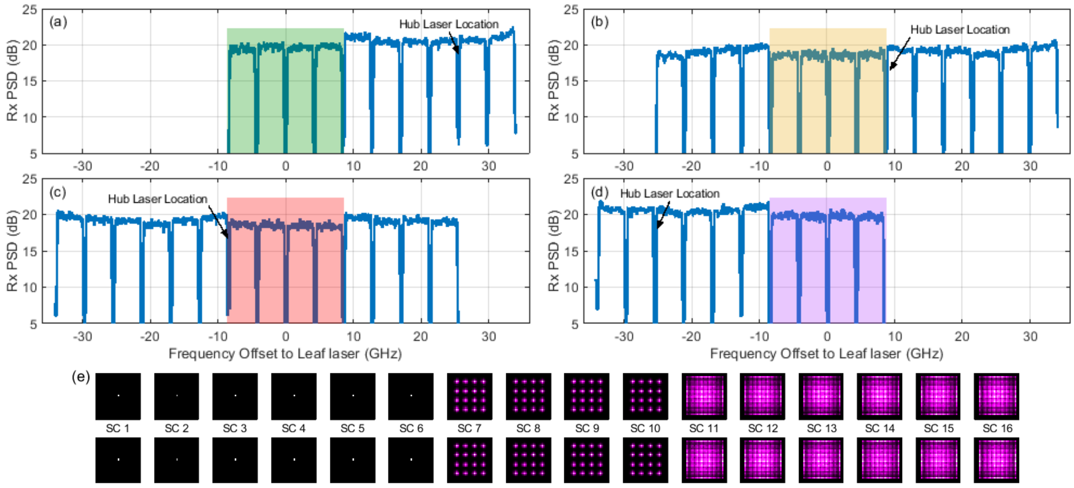

Fig. 3. Real-time measurements for 4×100G ⇔ 400G P2MP downlink: (a)–(d) Rx PSD at each leaf, (e) 16-QAM constellations for each SC at leaf-1 (X- and Y-pol separated).

Finally, in Fig. 4 we report the 4×100 Gb/s downlink performance of the four leaf nodes. For this measurement, we take the gain-shared pre-FEC Q-factor readings from each leaf Rx at 1 second intervals over an $\sim$1 hour soak. As observed, downlink communication to all four leaf nodes is very stable and error-free, with a very large margin to the FEC limit. It should also be noted that leaf nodes 2 and 3, which have the shorter propagation distance, have an approximately overlapping Q reading of $\sim$10.5 dB; whereas leaf nodes 1 and 4 – due to their lower received power level proportional to longer transmission – exhibit, as expected, a lower downlink performance.

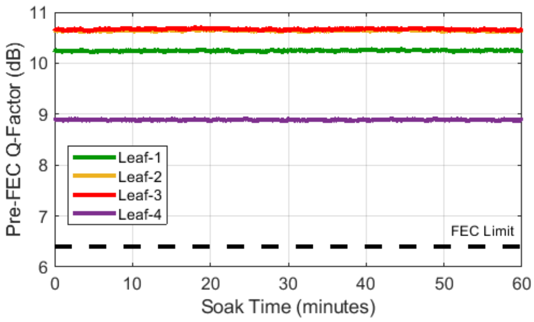

Fig. 4. Downlink pre-FEC Q-factors vs time. Note that the yellow line is partially hidden by the red.

4. Conclusion

Utilizing DSCM, we successfully demonstrate post-FEC error-free coherent P2MP transmission over a high-speed optical broadcast and aggregation setup. Real-time experiments are carried out on prototype coherent transceivers including DSP ASIC, FEC engine, and integrated TROSA.

References

1. "Cisco Annual Internet Report (2018−2023) White Paper." [Online]. Available: https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html
2. J. Cheng *et al.*, "Comparison of coherent and imdd transceivers for intra datacenter optical interconnects," in *Optical Fiber Communication Conference*, 2019, p. W1F.2.
3. J. Bäck *et al.*, "CAPEX savings enabled by point-to-multipoint coherent pluggable optics using digital subcarrier multiplexing in metro aggregation networks," in *European Conference on Optical Communications*, 2020, p. Tu2H.1.
4. A. Napoli *et al.*, "Enabling router bypass and saving cost using point-to-multipoint transceivers for traffic aggregation," in *Optical Fiber Communication Conference*, 2022, p. W3F.5.
5. D. Welch *et al.*, "Point-to-multipoint optical networks using coherent digital subcarriers," *Journal of Lightwave Technology*, vol. 39, no. 16, pp. 5232–5247, 2021.
6. G. Bosco *et al.*, "Analytical results on channel capacity in uncompensated optical links with coherent detection," *Optics Express*, vol. 19, no. 26, pp. B440–B451, 2021.
7. Y. Zhang *et al.*, "Digital subcarrier multiplexing for flexible spectral allocation in optical transport network," *Optics Express*, vol. 19, no. 22, pp. 21 880–21 889, 2011.
8. H. Sun *et al.*, "800G DSP ASIC design using probabilistic shaping and digital sub-carrier multiplexing," *Journal of Lightwave Technology*, vol. 38, no. 17, pp. 4744–4756, 2020.
9. T. Rahman *et al.*, "Digital subcarrier multiplexed hybrid QAM for data-rate flexibility and roadm filtering tolerance," in *Optical Fiber Communication Conference*, 2016, p. Tu3K.5.
10. A. Rashidinejad *et al.*, "Real-time demonstration of 2.4 Tbps (200Gbps/λ) bidirectional coherent DWDM-PON enabled by coherent nyquist subcarriers," in *Optical Fiber Communication Conference*, 2020, p. W2A.30.
11. D. Welch *et al.*, "Digital subcarrier multiplexing: Enabling software-configurable optical networks," *Journal of Lightwave Technology*, pp. 1–18, 2022.