Open XR Concept Introductory White Paper

OXR2022.001.002
January 2022

ABSTRACT:

This white paper introduces the Open XR system concept as a new optical communication network architecture that transforms current point-to-point optical networks into flexible, scalable and cost-effective point-to-multipoint networks. This new architecture, which aligns better with the hub-and-spoke traffic patterns observed in today's metro and access network segments, achieves interoperability across a variety of transceivers operating at different speeds using individually routed, digitally generated subcarriers.

This white paper is an adaptation of "Point-to-Multipoint Optical Networks Using Coherent Digital Subcarriers," by D. Welch, A. Napoli, et al., in Journal of Lightwave Technology, vol. 39, no. 16, pp. 5232-5247, 2021, doi: 10.1109/JLT.2021.

The Open XR Optics Forum
www.openxropticsforum.org

Open XR Optics Forum

The Open XR Optics Forum is the multi-source agreement (MSA) working group for XR optics, the industry's first point-to-multipoint coherent pluggable transceiver technology. The Open XR Optics Forum's mission is to foster collaboration that will advance development of XR optics-enabled products and services, accelerate adoption of coherent point-to-multipoint network architectures, and drive standardization of networking interfaces to ensure ease of multi-vendor interoperability and an open, multi-source solution ecosystem.

For additional information contact:

Open XR Optics Forum
www.openxropticsforum.org

Notice: This white paper has been created to assist the Open XR Optics Forum. The document is offered as a basis for discussion and communication, both within and without the Open XR Optics Forum. The Open XR Optics Forum, reserves the rights to at any time to add, amend, to withdraw statements contained herein. Nothing in this document is in any was binding on the Open XR Optics Forum or any of its members.

Permissions: You are authorized to download, reproduce, and distribute this document. All other rights are reserved. The provision of this document should not be construed as the granting of any right to practice, make, use, or otherwise develop products that are based on the document. Any and all IP rights related to this document and the designs disclosed within, except for the rights expressly mentioned above, are reserved by the respective owners of those IP rights.

Revision history:

Version	Author	Date	Description
V1.0	N.L. Swenson	01/24/2022	Date of initial publication
V1.1	P. Amaya	10/12/2023	Change name from Open XR Forum to Open XR Optics Forum

SOURCE: **Editor**
Norman L. Swenson
Company: Infinera
Address: 6373 San Ignacio Avenue, San Jose, CA 95119
Phone: +1 (408)572-5200
Email. Nswenson@infinera.com

1 INTRODUCTION 5
2 MIGRATING FROM POINT-TO-POINT TO POINT-TO-MULTIPOINT: A SHIFT IN NETWORK PARADIGMS 7
  2.1 Limitations of point-to-point networks 7
  2.2 Benefits and challenges of point-to-multipoint networks 8
3 DESIGN OF POINT-TO-MULTIPOINT NETWORKS 10
  3.1 Point-to-multipoint digital subcarrier coherent transceivers 10
  3.2 Network management and control 14
4 NETWORK APPLICATION SCENARIOS 16
  4.1 Passive optical networks 17
  4.2 Metro aggregation 18
  4.3 Mobile transport: A trip from antenna to DC 20
  4.4 Summary of network applications 21
5 SUMMARY AND CONCLUSIONS 22
6 REFERENCES 23

1 INTRODUCTION

Optical networks have gone through several disruption cycles over the past decades, led by the deployment of breakthrough technologies resulting in improved network architectures. Notable inventions, including the Erbium Doped Fiber Amplifier (EDFA), Photonic Integrated Circuit (PIC), Wavelength Selective Switch (WSS) — which enabled the Reconfigurable Optical Add and Drop Multiplexer (ROADM) — and coherent optical transmission, have all been widely adopted by the optical communication industry. Today's state-of-the-art optical transceivers deliver impressive performance, highly valued in core and submarine applications or any network where fiber is a scarce resource.

Further improvements in transceiver performance are ultimately limited by Shannon's capacity limit. Additional reduction in the cost per bit (CpB) for next generation high-end optical interfaces will have to come in large part from higher symbol rates, rather than by improving the Signal-to-Noise-Ratio (SNR) and Spectral Efficiency (SE), which can be limited by component imperfections. Nonetheless, there is an upper limit on the practical symbol rate that can be achieved, roughly scaling with Moore's law.

While higher data rates will reduce the CpB, it is only worthwhile if the operators can effectively utilize increased transceiver capacity. With the rise of the Internet, mesh voice traffic has been displaced by Hub and Spoke (H&S) data traffic; continued growth is forecast as 5G Radio Access Networks (RANs) and Internet of Things (IoT) enable new use cases [1,2]. Fig. 1 compares qualitatively the traffic patterns — from Point-to-Point (P2P) in (a) to Point-to-Multipoint (P2MP) H&S in (b).

While the total traffic in these networks is high, the traffic requirements at individual sites are relatively low. Typical metro or access network nodes may not be able to fully utilize high-capacity coherent transceivers for several years, leaving the industry in need of a way to continue reducing CpB within the current network architecture [3].

This white paper introduces Open XR based on Digital Subcarrier Multiplexing (DSCM) technology [4-6] to enable P2MP coherent optical networks with streamlined Layer 2/Layer3 (L2/L3) architectures, putting operators on a new trajectory for CpB savings. Open XR features these key characteristics:

- the throughput of a single high-capacity transceiver can be shared by multiple lower-capacity ones,
- digital switching at intermediate nodes is replaced by optical aggregation,
- remote reconfigurability minimizes manual interventions (e.g., for capacity upgrades), and
- smart optical transceivers are managed independently of hosting devices.

The remainder of this white paper is organized as follows. Section 2 discusses evolution of the network architecture from a predominantly point-to-point network to a point-to-multipoint composition. Section 3 discusses the digital subcarrier modem technology and network management that enable a point-to-multipoint network architecture. Section 4 presents Open XR network applications. Section 5 provides a summary and conclusions.

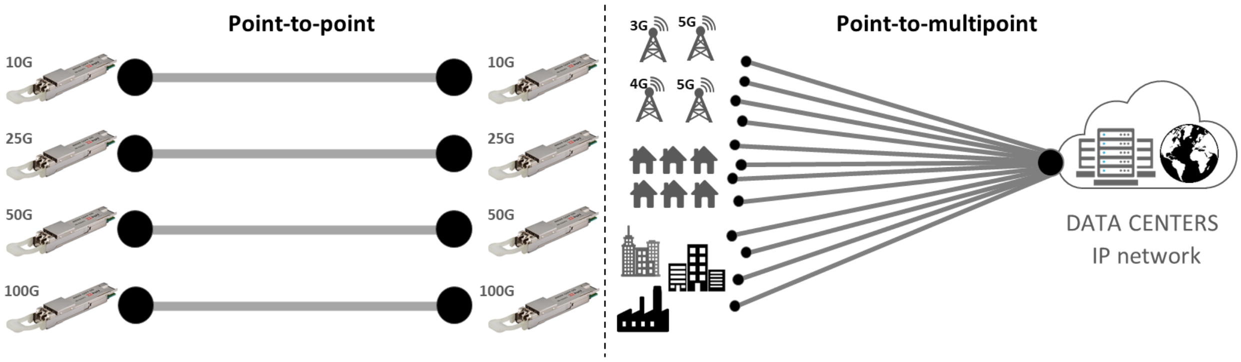

(a) (b)

Fig. 1: (a) Point-to-Point (P2P) versus (b) Point-to-Multipoint (P2MP) networks.

2 MIGRATING FROM POINT-TO-POINT TO POINT-TO-MULTIPOINT: A SHIFT IN NETWORK PARADIGMS

This section compares P2P and P2MP architectures for a given typical metro aggregation network, illustrated in Fig. 2. Sec. 2.1 discusses the relevant limitations arising from the underlying P2P infrastructures, while Sec. 2.2 highlights the benefits of using P2MP.

2.1 Limitations of point-to-point networks

Fig. 2(a) illustrates a simplified common metro/access network architecture, where the transceivers of N endpoints (5G antennas, curb aggregation boxes etc.) communicate with those at the electrical aggregation stage. N low data rate transceivers are needed on each side, i.e., 2N transceivers, plus 2 additional high-speed ones. While optimal for traditional telephony, the P2P architecture becomes sub-optimal with highly asymmetric H&S traffic. The pragmatic solution is to introduce a hierarchy of aggregation devices, typically IP routers, to allow each link in the network to use the optimal rate P2P transceiver; here the aggregation device serves the role of a gearbox, multiplexing traffic from various smaller optical transceivers onto a larger one. There are some drawbacks to this approach.

First and foremost, there is no obvious value to the end customer in regenerating traffic unnecessarily. Fiber spectrum is typically not a problem outside core networks, thus the lowest CpB solution is to minimize the number of regeneration points. Borrowing a term from the Toyota Production System, this is potentially *muda* - waste - a sign of a hidden cost that should be investigated [7].

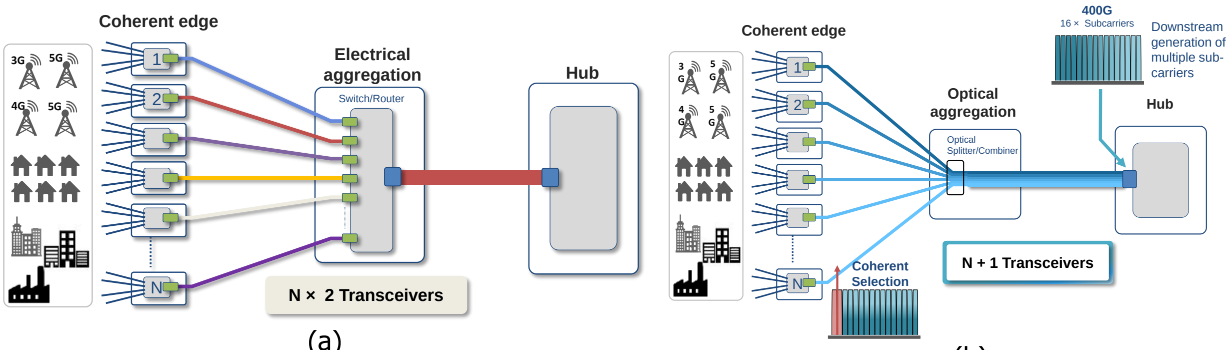

(a) (b)

Fig. 2: Example of intermediate aggregation network of N edge notes for (a) P2P and (b) P2MP. The former requires 2N low-speed transceivers plus 2 high-speed transceivers, while the latter requires only N low-speed transceivers plus 1 high-speed transceiver. In (b), we substitute the aggregation layer (and related transceivers) with a passive N:1 optical coupler thus simplifying Layer 0/Layer 1.

Once the traffic has arrived at the hub site it is terminated in N ports, where each port is sized according to the expected peak traffic demand and matches the edge nodes. If traffic in a particular leaf site increases, a physical site visit may be required to deploy more capacity, burdening Operational Expenditure (OPEX) budgets. Alternatively, the operator can choose to deploy more capacity up front, increasing CAPEX instead. Operators routinely make these decisions based on traffic forecasts and other factors. Despite the variety of traffic patterns, P2P networks have been deployed almost exclusively across the different segments of the optical network[1]. A clear limitation of this architecture is that if an endpoint requires more capacity, all transceivers between the endpoint and the electrical aggregation point must be replaced. Consequently, an operator would opt either for more frequent truck rolls or designing in a safety buffer of additional capacity on day one. Neither solution is economically optimal.

2.2 Benefits and challenges of point-to-multipoint networks

This section discusses network simplification and adaptation to modern data traffic dynamics and how coherent, coupled with DSCM and with advanced DSP, can realize a P2MP architecture.

Fig. 2(b) shows the impact of implementing a subcarrier based Open XR architecture in the same typical metro aggregation network of Fig. 2(a). Here, the electrical aggregation stage of Fig. 2(a) has been replaced by a simple 1:N passive optical combiner. The number of devices and stages needed to aggregate the traffic to be transported to the next hub is greatly reduced.

The Open XR architecture enables routing of digital subcarriers (SC)s independently from an endpoint to the hub. The SCs are locked in frequency in a leader/follower relationship to the high-speed transceiver at the hub, and a muxponder is no longer needed to aggregate the signals from the low-speed transceivers. In Fig. 2(a), even if the electrical aggregation device is replaced by a DWDM mux/demux, the N bookended transceivers simply move from the electrical aggregation device to the hub.

In Fig. 2(b), only N low-speed transceivers and 1 high-speed transceiver are required, 50% less than in Fig. 2(a). In addition, in the P2MP architecture, multiple low-speed transceivers (spokes) are now directly connected to the high-speed transceiver (hub), breaking the book-ended transceiver paradigm. This approach eliminates the need for intermediate traffic aggregation stages while leveraging larger, more efficient switching devices at centralized sites. The related costs for power consumption, footprint, sparing parts, and grooming equipment are consequently reduced. Furthermore, P2MP enables cross-layer Layer 1/Layer 2 (L1/L2) savings by an efficient utilization of the optical transceivers, which can now flexibly provide the exact required amount of capacity.

---

[1] P2MP is employed in access, e.g., Ethernet PON and Gigabit PON. However, the protocols employed there rely on time division multiplexing (TDM).

P2MP can also maximize Layer 3 (L3) efficiency, density, and simplicity by replacing large numbers of low-speed ports with fewer, more efficient, high-speed ones that can be used both for aggregation and as network interfaces. Because capacity can be assigned on demand, seamless network upgrades can be supported, and costly site visits are avoided. For example, an operator may choose to deploy 100Gb/s P2MP pluggable optics in 12 leaf nodes and only turn up a single 25Gb/s digital subcarrier per node on day one. The 12 SCs can be terminated in a single 400Gb/s hub router port, as opposed to 12 100Gb/s pluggable P2P modules. As capacity grows in edge nodes and more SCs are turned up over time, the operator can deploy more 400Gb/s ports in the hub. Compared to using P2P optics, operators can deploy more edge capacity to reduce OPEX associated with site visits, while saving on CAPEX at the hub site.

Maintaining the same definition of SCs over multiple generations of P2MP optics ensures that different generations of pluggables can interoperate, and operators can thus design multi-generational network architectures. Routers and other hosting devices can be independently upgraded. Hub sites can be seamlessly expanded, e.g., to 800G, without requiring upgrade of the leaves, decoupling nodal upgrades from network-wide ones. Network operators can maximize Return on Investment (ROI) and ensure a smooth and cost-effective capacity upgrade of the network.

The P2MP architecture does present some challenges. First, the aggregation devices removed from the network perform more than the aggregation function. Routers are feature-rich devices and operators use them for a variety of functions. Essential features may need to be virtualized or hosted on other types of devices.

A second potential concern is around reliability. While reducing the number of pluggable optics in hub sites will improve network uptime, operators need to ensure that the impact of a failure is contained. While the failure of a P2P hub module connected directly to a single edge device will only bring down the traffic to that one edge device, failure of a P2MP hub optic can bring down the link to multiple nodes; this can be addressed through protection or restoration.

Removing electrical aggregation devices will also improve network uptime but can reduce the operators' ability to reroute traffic around failures. This would be more relevant in a mesh network, and likely matters less for ring, horseshoe and tree topologies in metro and access aggregation networks.

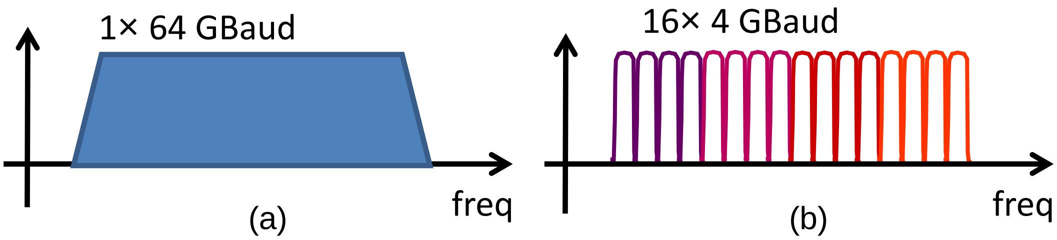

Fig. 3: Pictorial comparison of (a) a single carrier 400G and (b) its equivalent realized with DSCM using 16 subcarriers.

3 DESIGN OF POINT-TO-MULTIPOINT NETWORKS

3.1 Point-to-multipoint digital subcarrier coherent transceivers

Given the range of applications for the proposed technology, it is natural for P2MP transceivers to be productized as digital coherent optical (DCO) pluggable modules. P2MP pluggables present a similar complexity, number of components, and variety of form factors as conventional P2P ones. Conventional pluggables allow communication between same-speed devices (in a P2P format), provided that the modulation format and the coding are identical. By employing DSCM, we can establish a remote communication between all coherent pluggables (hub and all leaves, low- to high-speed devices) using embedded communication channels. Management can be as standard P2P, via host device or as a part of the optical network (versus separate L3 management).

The direct communication between low-speed transceivers and high-speed transceivers is enabled using DSCM subcarriers. This is a digital communication technique that subdivides the transmitted spectrum into digital SCs. From a networking perspective it is a form of sliceable bandwidth. Fig. 3 qualitatively compares the spectra of channels carrying 400Gbps using (a) a single 64GBaud signal, and (b), its equivalent with 16 SCs at 4GBaud using DSCM.

DSCM has been deployed in P2P links for reducing the impact of linear and nonlinear impairments, maximizing reach, and jointly optimizing the linear compensation of Enhanced Equalization Phase Noise (EEPN) [4,8-10]. The cited works demonstrate that a lower symbol rate can reduce the impact of nonlinearities resulting from fiber propagation. DSCM has also been proposed to increase tolerance to bandwidth narrowing by cascaded filters (by adapting the order of the modulation format used in individual SCs [11]) and to reduce power consumption [12].

The SCs are digitally generated at the Transmitter (TX) and do not overlap in frequency. In the case of Nyquist SCs with excess bandwidth β, the subcarriers can be spaced at (1+ β) times the subcarrier symbol rate. This allows for the spectral efficiency illustrated in Fig. 3, where both 400G channels occupy the same bandwidth of (1+β) times 64GHz. The key DSP differences between single wavelength and DSCM are discussed in Sec. 3.2.

Since a DSCM transceiver operates at a symbol rate R_s/N, where R_s is the symbol rate of a single carrier transceiver with the same modulation and bit rate and N is the number of digital SCs, the performance might be limited by the value of the laser linewidth Δf [13]. In this case, the symbol period of the subcarrier is N times that of the single carrier signal. The longer symbol period of the DCSM system requires attention to ensure that laser phase noise does not limit performance.

Digital signal processing for point-to-multipoint transceivers

Fig. 4 compares the DSP at the TX and RX of a typical optical system transmitting (a) a single carrier signal, and (b) a DSCM signal. The general structure of the DSP algorithms for the two transceivers is similar, with a few important differences. The DSP within the DSCM in Fig. 4(b) is the parallelized version of Fig. 4(a), where each stream (SC) operates at a symbol rate R_s/N. In terms of DSP algorithms, this is a relevant feature, as it allows a reduction in hardware complexity [4].

In both TXs, the DSP comprises a mapper into, e.g., a 16-QAM constellation, and a Finite Impulse Response (FIR) filter for the Nyquist pulse shaping with rolloff $\rho \sim 0.05$. This filter also performs digital pre-emphasis to mitigate linear impairments such as analog I/Q skew and rolloff. In the case of DSCM, this structure is instantiated N times in parallel and a digital multiplexer is needed in front of the 4×DAC. Finally, the signal is input to the driver and the dual polarization IQ Mach-Zehnder Modulator (IQ-MZM). In the case of DSCM, we implement an adaptable guard band of ~100 MHz between the SCs.

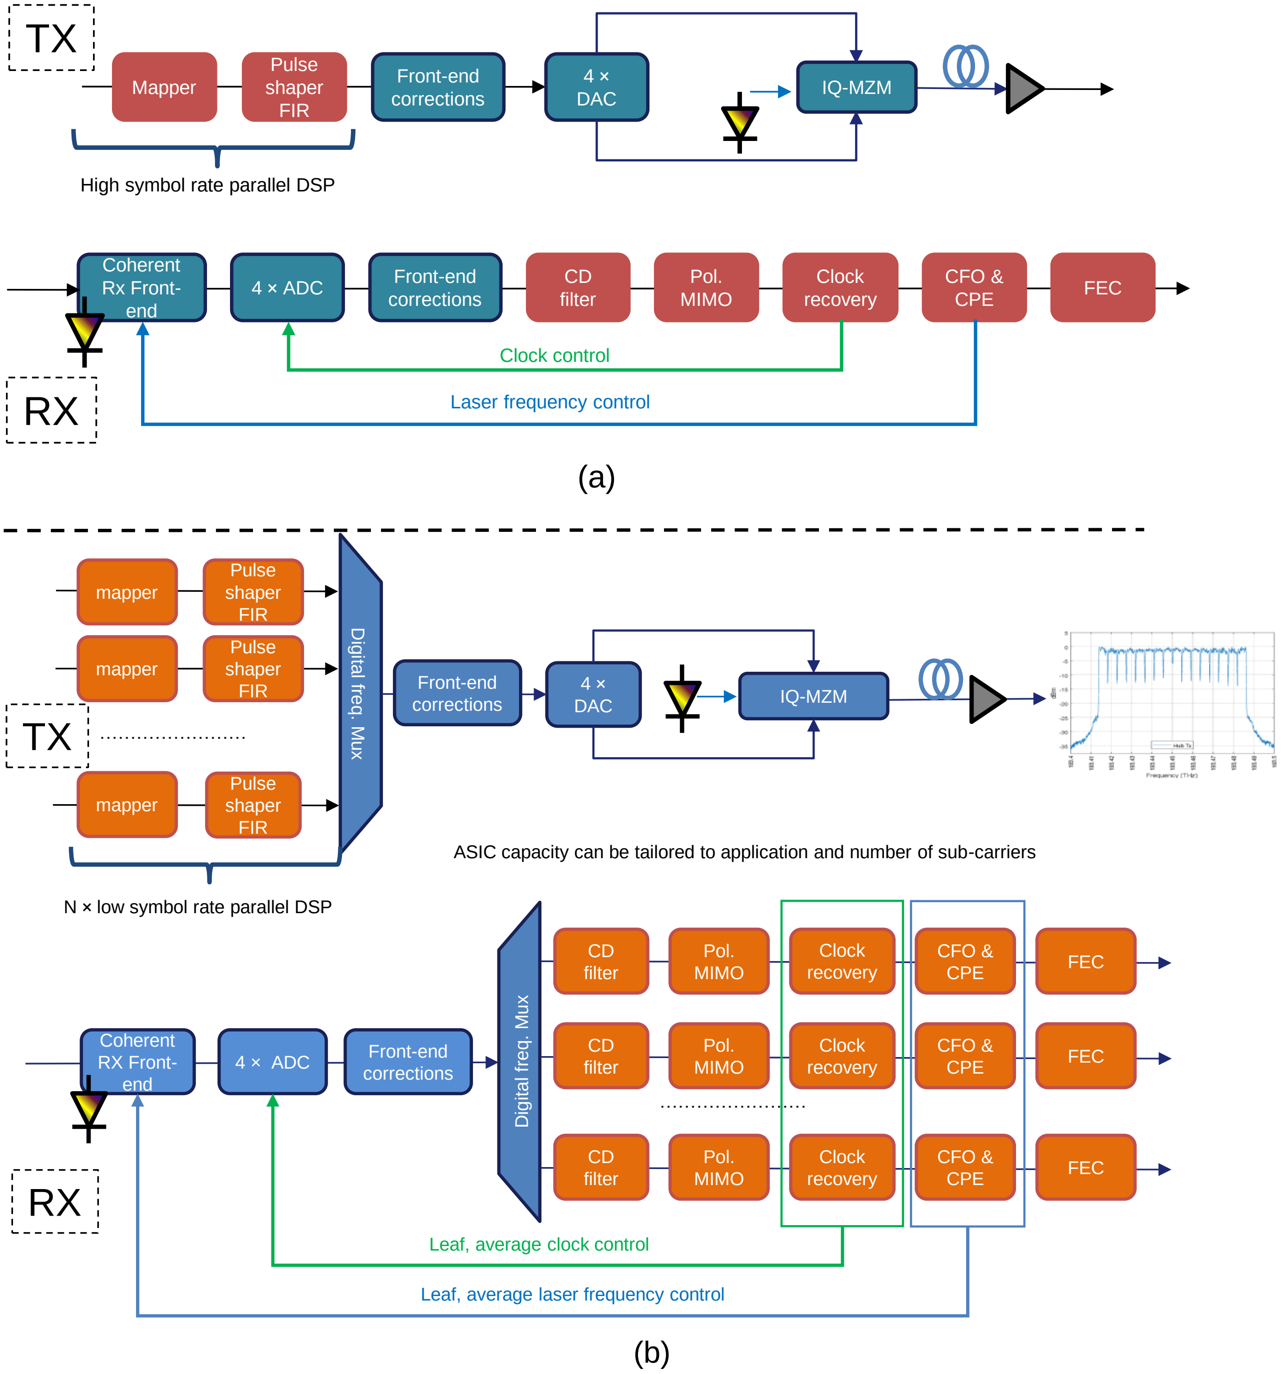

Fig. 4: (a) DSP blocks and system for a single wavelength; (b) the DSP parallelized version for the case of DSCM. CPE: Carrier phase estimation; CFO: Carrier phase offset. The clock recovery comprises a clock phase detector and an interpolator for fast jitter comparison.

After fiber propagation, the signals are detected by the coherent front end and converted into the digital domain by a 4×ADC. Next, the DSP algorithms at the RX mitigate the impairments. Their block structure is essentially the same for single wavelength and DSCM, apart from the parallelization and the feedback signals that are calculated as the average over all SCs.

The first block is the front-end correction, which compensates for effects such as I/Q skew, passband rolloff, power imbalance, and frequency ripple. Although the module design is comparable between single wavelength and DSCM, the penalty introduced by a single effect can be higher for either the former or the latter. For example, as shown in [14] for the case of I/Q skew, the DSP within DSCM generates a mirror image which distorts the SC in the mirror position. In case of DSCM, this additional distortion cannot be compensated for by traditional characterization methods such as those based on the IQ. One possible solution, based on the simultaneous processing of one SC and of its symmetric-frequency counterpart, showed that the introduced mutual interference can be considerably reduced [14].

Next, a digital frequency demultiplexer separates the SCs. The following Chromatic Dispersion (CD) filter compensates for the bulk dispersion accumulated by the channel along the link. Since in the case of DSCM we operate at R_s/N, the complexity of the individual equalizers is reduced with respect to a single carrier RX DSP, as the filter length scales with $1/{R_s}^2$[4].

Next, a Multiple Input Multiple Output (MIMO) equalizer separates the two polarizations and compensates for the residual dispersion and for the Polarization Mode Dispersion (PMD). These two algorithms preserve the same structure between single carrier and SC [15]. As reported in [7], the parallelization needed by the DSCM RX DSP is beneficial in reducing the EEPN, which can be relevant in case of a large amount of accumulated CD.

Once the linear effects have been fully compensated, we can synchronize in time and frequency by estimating the information derived by the clock recovery, the Carrier Phase Estimation (CPE), and the Carrier Frequency Offset (CFO) modules. In both cases, a feedback signal is generated and used in the leaf. At the hub, both clock and laser frequency are free running, locked only by the local timing and laser frequency reference. The first control signal is sent from the clock recovery to the 4×ADC, while the second is fed back from the CPE and CFO to the Local Oscillator (LO) at the coherent front end. These are used to achieve laser frequency and timing synchronization.

A technical challenge in the system described in Fig. 2(b) is the possibility that SCs can collide with each other after we combine them at the passive N:1 optical combiner. To avoid this, we first rely on the leaf module to demodulate successfully at least one SC from the hub, which always transmits its SCs. In a second step, the carrier recovery will provide information to tune the leaf LO laser to lock the LO to the incoming SC, thereby achieving frequency locking between hub and leaf lasers. A similar concept can be applied to achieve timing synchronization. Finally, the FEC and the decision block provide the error-free signal.

It is worth noting that frequency and clock synchronization are among the reasons for selecting DSCM rather than Orthogonal Frequency-division Multiplexing (OFDM) or Discrete Multitone (DMT). In fact, as reported in [16] the computational complexity between OFDM/DMT and DSCM is comparable only if we consider P2P transceivers. As soon as we move to a P2MP scenario, OFDM becomes much more complex because of frequency and phase instability. This is because when the different leaves send their OFDM SCs, the two types of synchronization become extremely complex.

3.2 Network management and control

The introduction of P2MP optical transmission technology in operator networks raises operational network management (NM) questions, particularly regarding the commissioning/management of SCs and the abstraction of the P2MP data flows.

The abstraction challenges in NM for the proposed solution concern the shift in how functions are allocated between modules and products. Functions that were previously housed in a DWDM transponder chassis are now possible within a single pluggable optic.

Moreover, the trend toward disaggregation adds complexity at the NM level when DWDM transceivers are productized in pluggable form factors for use in 3rd party host devices. This issue is not limited to P2MP optical technology, but is industry-wide, and applies to any type of advanced DWDM pluggable optic.

The current state-of-the-art utilizes register-based standardized information models or multi source agreements [17] that hosting device and pluggable vendors must support to achieve interoperability. To introduce a new DWDM transceiver technology, updates to standardized information models need to be agreed upon and implemented within both the pluggables and hosting devices, tying together development and deployment cycles, which in turn slows down technology adoption.

The blending of functions from various layers of the Open System Interconnection (OSI) model presents a challenge with implications for Software Defined Networking (SDN) architectures. A concrete example is the hierarchical architecture using separate controllers for packet and transport layers, with a hierarchical controller above these: this architecture is built on the assumption that the packet and transport controllers do not need to communicate directly or be aware of each other’s configuration state or status. Integrating smart pluggable DWDM optics into IP routers has management challenges, and the same may hold for integration into other network devices (Optical Line Terminal (OLT)s, RAN equipment, network interface cards, etc.).

As smart or advanced pluggables absorb features and capabilities typically present at either the card, chassis or system levels (such as in-band and out-of-band communication channels, a control plane, and topology awareness), new methods are required to enable control without disrupting the existing hierarchy.

In this context, we envision the evolution of advanced pluggable optics to the point where they can be managed through a combination of data plane and/or standardized DCN Ethernet interfaces, e.g., via a Virtual LAN (VLAN) as a separate network entity, combined with several pluggable optic modules interacting with each other to complete coordinated operations directly between modules using both in-band and out-of-band control channels over the fiber link.

These operations include such functions as discovery, automated turn-up, and the ability to manage the lower bandwidth devices (e.g., spoke devices at 25G) via the higher bandwidth optic (e.g., hub at 400G). With these functionalities, NM can establish a connectivity matrix between destinations, enable bandwidth allocation, and manage alarms and network status, allowing the pluggable transceiver to act as an addressable network element within a much larger network configuration.

The functionalities of smart/advanced pluggables are realized both by incorporating processing capability within the pluggable and by implementing communication channels between pluggables; with the associated algorithms, they create a control plane with full topology awareness. All these capabilities are valuable for automated discovery and authentication operations, where an operator needs the modules to establish communication before a datapath has been established between modules. Utilizing the established in-band communication channel available on each SC allows the edge optics to appear as a simple grey optic to their hosting device, while exchanging management information to and via the hub optic. This allows the user to keep a separately managed optical network, independent of the platform within which the optical modules are hosted. We envision that pluggables will absorb even more of the functionalities now implemented on the hosting devices and will require a more flexible and robust communication method than the register-based systems used today.

The preferred method for module-to-module communication depends on the operation, using both in-band and out-of-band control channels. For example, initial locking of a leaf to a hub may rely on out-of-band communication, whereas the monitoring performance of an individual SC may be supported by an in-band communication channel. The preferred method to allow communication from an SDN microservice application directly to the optical pluggables as a separate network entity is over the DCN to the pluggable transceiver through the device hosting the pluggable. The SDN microservice itself is a container-based application with a NETCONF/RESTCONF interface based upon YANG models. It can be run in the cloud, on central servers, or even in containers on the host platform, depending on latency requirements and compute resources available on the hosting device for other applications to run.

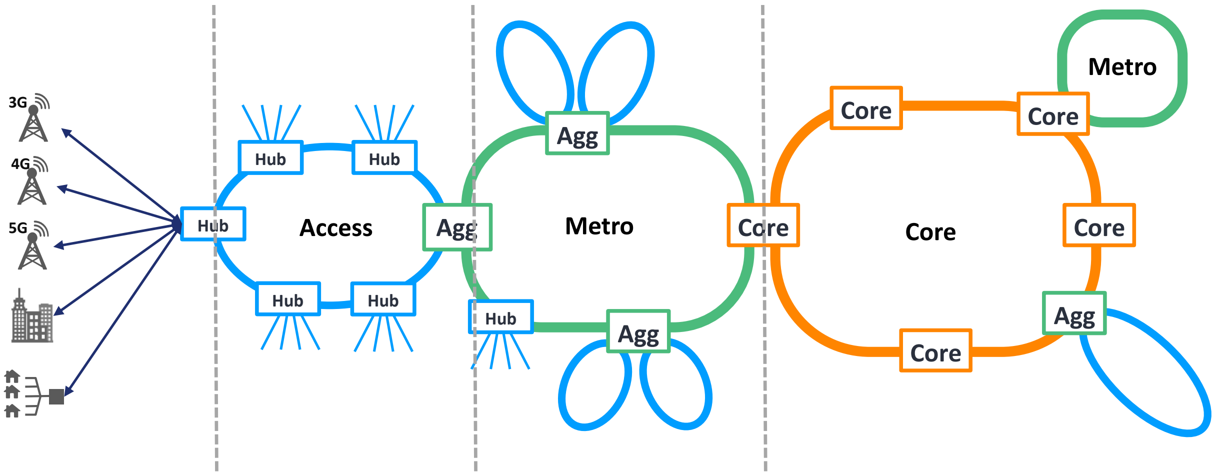

Fig. 5: Generic representation of a telecommunication network from end users to core, passing through access and metro.

The features enabled via this SDN application include per-SC visibility and management, topology awareness, power balancing information, telemetry streaming and Threshold Crossing Alarm (TCA) settings, direct software upgrades, and encryption.

Removing the need for the host to understand and translate all new advanced operations and the required supporting information models greatly reduces the multivendor interoperability complexity for the operator managing a network and allows the innovation cycles each to run at its own pace.

4 NETWORK APPLICATION SCENARIOS

The following discussion on network applications for P2MP is central to the concepts presented in this paper. Fig. 5 illustrates a high-level representation of a telecommunication network from end users to core, passing through access and metro aggregation stages. Based on actual traffic growth, we focus on three of the relevant scenarios displayed in Fig. 5: (i) PON in Sec. 4.1; (ii) metro aggregation in Sec. 4.2 and (iii) mobile fronthaul in Sec. 4.3. Thereafter, we highlight the changes required to re-architect the network in view of P2MP with DSCM. In Sec. 4.1, we also provide a comparison between Time Division Multiplexing (TDM) and Frequency Division Multiplexing (FDM), highlighting the main differences when applied to PON. Sec. 4.4 summarizes the application scenarios.

Core, metro, and access networks are considerably different in terms of geographical size, number of nodes, and magnitude of traffic – both the total and average traffic per node pair. In the sections below, we focus on two different network applications – access and metro aggregation – and then discuss the applicability of P2MP technology to a generic mobile transport use case.

4.1 Passive optical networks

PONs have been one of the enablers of broadband telecommunication systems such as FTTH, and along with the possibility of providing the right bandwidth to the individual user, have significantly contributed to the internet revolution. PONs have been widely deployed using current cost-effective direct-detect (DD) systems [18]; their success is largely due to the adoption of the P2MP approach, an indication of the value of this optical networking architecture. With PON, transmission is realized with TDM, and the transceivers are manufactured in volumes exceeding millions of units per year, as discussed in [19] and its references.

Nevertheless, the P2MP PON architecture does present some limitations. In the case of a data rate upgrade, for example, all devices (hub and leaf modules) need to be replaced because the same communication rate must be utilized. From a technology and power consumption point of view this can become challenging and inefficient, as the end customer needs to deploy the same fast opto/electronics as the hub transceiver, although the traffic requirement may be much lower.

Furthermore, a high data rate with TDM PON might be limited in terms of reach, synchronization, and latency.

For example, with FDM it is relatively simple to increase capacity, but with TDM PON, the requirement for fast transceivers causes an issue with the demands of continuous traffic growth. It is not simple to build TDM PON with coherent technology, and direct detect receivers do not offer the possibility of exploiting more advanced DSP techniques, e.g., to improve the reach. This is an important limitation, and it is exacerbated by the burst mode transmission typical of this architecture. DSP algorithms can be severely impacted during the learning phase. For example, the tracking of polarization must be carried out for each burst.

Moreover, burst transmission induces transients on the electrical current and on the laser output power that affect laser stability. Because of this, it is not suitable for DWDM. As a result, only a few wavelengths will be available.

In addition, great disparity in distances between leaf modules and the hub can make synchronization extremely challenging with a TDM approach.

Finally, latency and RX sensitivity are also critical aspects in PON design. The former is severely affected by the usage of TDM [20], while the latter is not sufficiently high to enable long distances with existing DD.

Because of all these limitations, the standardization of PON advances at a slower pace than other technologies.

To provide the required flexibility and capacity, we propose an FDM-based approach for P2MP as discussed in Sec. 3. Together with coherent and DSP, it supports DWDM, thus enabling a high level of scalability. Issues with traditional PONs are solved when FDMs leverage DSP. Furthermore, in the case of FDM, the costs for end users and hub are different, the upgrades are independent, and they can be performed per individual end user (at a given coarse granularity). Thanks to these characteristics, P2MP based on DSCM and coherent can achieve a capacity exceeding 100G.

Fig. 6 shows how to deploy a P2MP optical network over a brownfield Optical Distribution Network (ODN), with varying hub-to-leaf attenuation. In this example, we show a 400G (16×25G) hub connected to ten leaves at frequency slots of 25G. In this case, either six more leaves can be connected to the distribution network without adding a new hub, or some of the leaves can be upgraded to higher capacities. Each leaf has the maximal throughput of 100G, i.e., 4×25G.

As an example, this type of architecture can initially be adopted for business/science parks with large traffic demands where current PON technologies cannot deliver the required capacity. P2MP optics based on coherent transmission with DSCM offers a roadmap to higher end user capacities than On-off Keying (OOK) implementation, and as traffic demands in access networks grow, the proposed implementation provides additional advantages, as discussed above.

4.2 Metro aggregation

Fig. 7 shows the topology of a typical Content Service Provider (CSP) metro area network link, aggregating H&S traffic from six smaller edge sites into two larger hub sites. The optical line system leverages a filterless architecture, imposing no channel plan requirements on the leveraged optical transceivers. Traffic protection – in dual hub routers – is naturally supported, and equipment in all edge sites can be deployed with full diversity for east- and west-facing traffic.

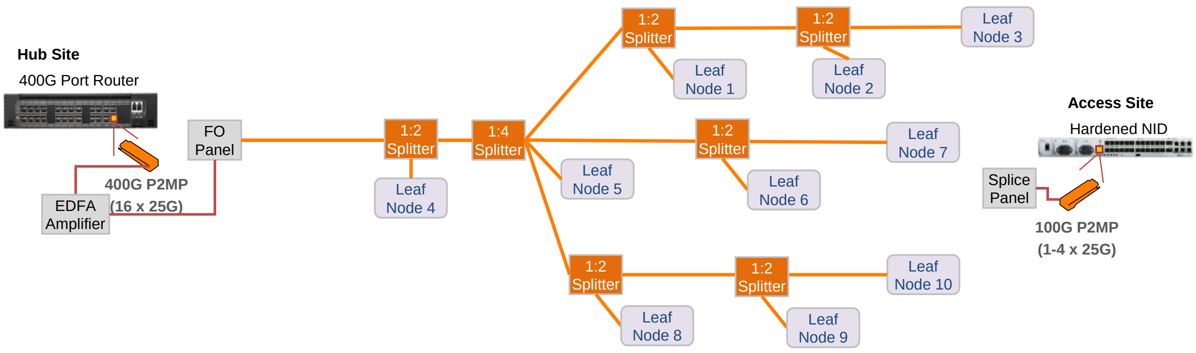

Fig. 6: PON overlay realized with a P2MP transceiver where the 400G can communicate with the 100G leaves.

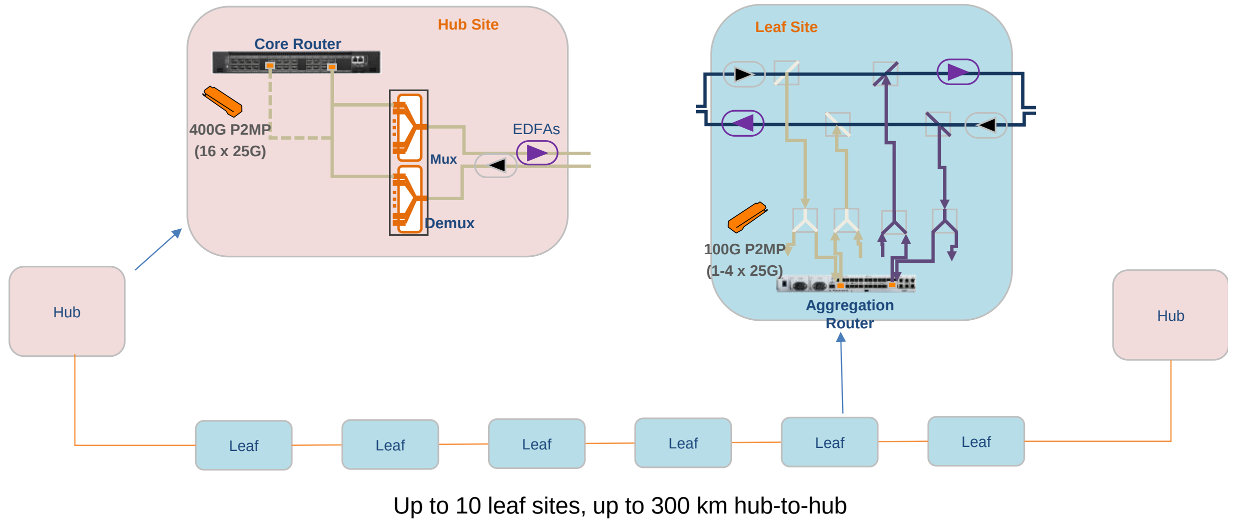

Fig. 7: Application of the P2MP transceiver to a metro aggregation network, with 2×400G hubs and 6×100G leaves.

Traffic volumes are much higher than in access networks as significant aggregation has taken place, and operators today are starting to fill their legacy networks based on 10 Gb/s DWDM transceivers and are in the process of planning and deploying next generation metro networks with higher rate DWDM technology. Using P2MP technology in this section of the network looks very promising as the traffic demands are well aligned with the coherent transceiver technology available today. Edge sites can use 100 Gb/s P2MP optics and fill 25 Gb/s SCs, leaving room for future growth. Hub sites can leverage larger 400 Gb/s P2MP optics to deliver significant savings exceeding 70% [3].

The hub site configuration is a core router with support for 400Gb/s ports, an add/drop mux and an EDFA to guarantee the optimal launch power of the optical signal. At the far end of the link is another hub site, in a different physical location, providing redundancy.

One of the major benefits of this solution is that the hub routers can use larger Ethernet ports, thus reducing L3 port count. For a P2P configuration, there would need to be a smaller hub router port per leaf site, which can result in a significant increase in cost not only at L1, but also at L2/L3. The total capacity required in the hub router can also increase, especially when the hub router ports are poorly utilized and there is a relatively large number of edge sites connected to a hub site.

An illustrative extreme example would be for a hub router to be connected to 16 edge routers, each having a traffic requirement of 25G. With the proposed P2MP 400Gb/s pluggable optics, a single 400Gb/s Ethernet port can serve those 16 nodes, while 100G P2P pluggable optics would require 16×100GbE ports in the hub router. In addition, SC routing can be carried out via software, just as if we built a Colorless-Directionless-

Contentionless (CDC) ROADM architecture in the digital domain within the DSP, leading to savings in both L2 and L3. Capacity can be reallocated as sites grow randomly over time. In the extreme example above, all edge nodes could scale up to 100G of traffic with no need for a truck roll onsite. Only the hub site would need to be visited as traffic grows.

4.3 Mobile transport: A trip from antenna to DC

RANs are deployed by mobile network operators to connect user handsets to the mobile core, and they rely on both wireless and optical technologies for communication. As higher frequency bands are brought into use, cell sizes decrease, and operators must put more focus on the architecture of the optical networks deployed to transport user data from the core nodes out to cell sites. Mobile transport networks are aggregation networks, connecting thousands of cell sites into a core node, making them a straightforward application for P2MP optical technology. In fact, PON is already used in mobile transport applications, especially for cell site backhaul where traffic volumes per link are still modest.

As 5G network architectures are defined, there are several technical, economic, and operational drivers to centralize the baseband functions for several cell sites, relying on optical links to connect distributed units housing the baseband function to radio units at the cell site. Because of the protocols employed, traffic on the fronthaul interface is greater than the underlying user traffic, and the fronthaul traffic has the potential of outsizing mid- and backhaul traffic in RAN transport networks.

This introduces scaling challenges for PON technology, potentially creating a bottleneck to the cell site as Centralized Radio Access Network (CRAN) architectures become prevalent. P2MP optics based on coherent DSCM offers a step function in capacity and is a clear candidate for use in future fronthaul networks.

Fig. 8 shows a set of three 120-degree radio units with a P2MP transceiver directly integrated, sending traffic back to a virtual Distributed Unit (DU) at the hub site, where the P2MP transceiver is hosted in a network interface card. P2MP transceivers can be sized to match the requirements of the hosting devices, removing the need for active aggregation devices in the cell sites and leaf/aggregation switches in the hub sites.

While the fronthaul application can be addressed in the short term with the use of low-cost OOK Coarse Wavelength Division Multiplexing (CWDM) optics, the limited fiber capacity will drive up access fiber costs.

4.4 Summary of network applications

We have shown the application of the P2MP and DSCM technologies in a variety of deployed network topologies, ranging from single and dual fiber to aggregation networks that include PON infrastructure and ROADM networks with protection ring configurations. As a result of the inherent flexibility in the definition and allocation of frequency channels, the Open XR transceiver is compatible with most as-deployed network architectures. To achieve full flexibility and lowest cost implementation, the P2MP transceiver is meant to be deployed in a broadcast configuration as it pertains to the capacity that is aggregated by the individual hub optic. While the broadcast capability results in the full capacity signal being presented at each leaf element, the leaf elements are assigned to one or more particular SCs through the management plane.

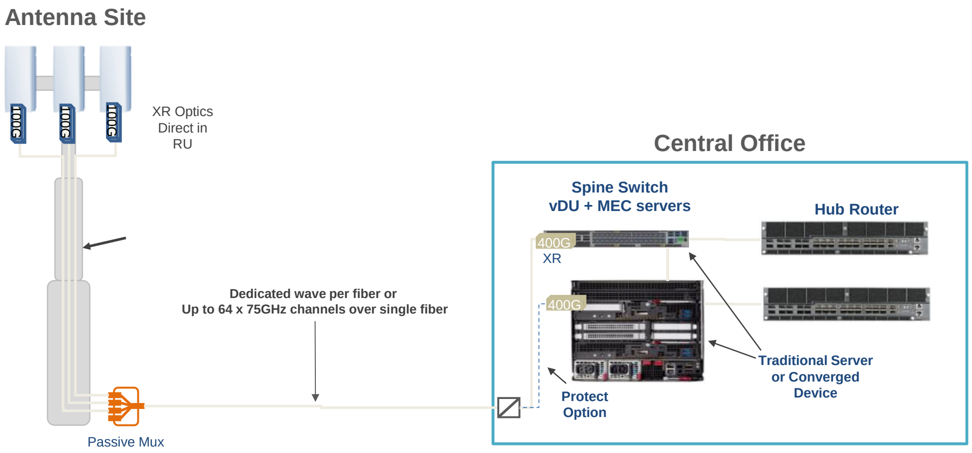

Fig. 8: Wireless fronthaul.

5 SUMMARY AND CONCLUSIONS

This white paper introduces the Open XR architectural concept of point-to-multipoint in optical networks. This paradigm shift in optical networking is a natural follow-on to the introduction of coherent optical technology. This shift is motivated by a comparison with the point-to-point approach, which has so far predominated in optical telecommunications, though not in wireless. One of the strongest motivations for this new approach is the significant evolution of traffic from endpoint-to-endpoint to current hub-and-spoke patterns, at least in the parts of the network experiencing the fastest growth. The necessary technology has been described in terms of hardware, digital signal processing, and network management. Several application scenarios have been analyzed, showing that this approach can be highly beneficial for metro aggregation, access, and fronthaul network segments.

The proposed Open XR solution provides the best possible match between network architecture and traffic patterns. The proposed architecture can coexist with current P2P systems. As 5G and beyond proliferate, high capacity will be needed in access networks, further accelerating the demand for bandwidth. In the context of hub-and-spoke and high-capacity traffic patterns, the Open XR point-to-multipoint technology supported by coherent transceivers using frequency division multiplexing can cost effectively satisfy future requirements of network operators.

6 REFERENCES

[1] Cisco, "Cisco Annual Internet Report (2018–2023)," 2020. doi: 10.1016/s1361-3723(20)30026-9.

[2] M. Schiano, A. Percelsi, and M. Quagliotti, "Flexible Node Architectures for Metro Networks [Invited]," *J. Opt. Commun. Netw.*, vol. 7, no. 12, p. B131, Dec. 2015, doi: 10.1364/JOCN.7.00B131.

[3] J. Back *et al.*, "CAPEX Savings Enabled by Point-to-Multipoint Coherent Pluggable Optics Using Digital Subcarrier Multiplexing in Metro Aggregation Networks," in *2020 European Conference on Optical Communications (ECOC)*, Dec. 2020, no. 1, pp. 3–6, doi: 10.1109/ECOC48923.2020.9333233.

[4] H. Sun *et al.*, "800G DSP ASIC Design Using Probabilistic Shaping and Digital Sub-Carrier Multiplexing," *J. Light. Technol.*, vol. 38, no. 17, pp. 4744–4756, Sep. 2020, doi: 10.1109/JLT.2020.2996188.

[5] D. Krause, A. Awadalla, A. S. Karar, H. Sun, and K. Wu, "Design considerations for a digital subcarrier coherent optical modem," in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, 2017, pp. 1–3.

[6] Y. Zhang, M. O'Sullivan, and R. Hui, "Digital subcarrier multiplexing for flexible spectral allocation in optical transport network," *Opt. Express*, vol. 19, no. 22, p. 21880, Oct. 2011, doi: 10.1364/OE.19.021880.

[7] T. Ohno, *Toyota Production System: Beyond Large-Scale Production*. Productivity Press, 1988.

[8] F. Buchali, W. Idler, K. Schuh, L. Schmalen, and T. A. Eriksson, "Study of electrical subband multiplexing at 54 GHz modulation bandwidth for 16QAM and probabilistically shaped 64QAM," pp. 49–51, 2016.

[9] P. Poggiolini, Y. Jiang, A. Carena, G. Bosco, and F. Forghieri, "Analytical results on system maximum reach increase through symbol rate optimization," in *Optical Fiber Communication Conference*, 2015, vol. 34, no. 8, p. Th3D.6, doi: 10.1364/OFC.2015.Th3D.6.

[10] L. B. Du and A. J. Lowery, "Optimizing the subcarrier granularity of coherent optical communications systems," *Opt. Express*, vol. 19, no. 9, p. 8079, Apr. 2011, doi: 10.1364/OE.19.008079.

[11] 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, doi: 10.1364/OFC.2016.Tu3K.5.

[12] L. Velasco *et al.*, "Autonomous and Energy Efficient Lightpath Operation based on Digital Subcarrier Multiplexing," *IEEE J. Sel. Areas Commun.*, vol. 8716, no. c, pp. 1–1, 2021, doi: 10.1109/JSAC.2021.3064698.

[13] S. Dris, P. Bakopoulos, I. Lazarou, C. Spatharakis, and H. Avramopoulos, "M-QAM carrier phase recovery using the viterbi-viterbi monomial-based and maximum Likelihood Estimators," *2013 Opt. Fiber Commun. Conf. Expo. Natl. Fiber Opt. Eng. Conf. OFC/NFOEC 2013*, pp. 31–33, 2013.

[14] G. Bosco, S. M. Bilal, A. Nespola, P. Poggiolini, and F. Forghieri, "Impact of the Transmitter IQ-Skew in Multi-Subcarrier Coherent Optical Systems," in *Optical Fiber Communication Conference*, 2016, no. 1, p. W4A.5, doi: 10.1364/OFC.2016.W4A.5.

[15] S. J. Savory, "Digital filters for coherent optical receivers," *Opt. Express*, vol. 16, no. 2, p. 804, 2008, doi: 10.1364/OE.16.000804.

[16] B. Spinnler, "Equalizer Design and Complexity for Digital Coherent Receivers," *IEEE J. Sel. Top. Quantum Electron.*, vol. 16, no. 5, pp. 1180–1192, Sep. 2010, doi: 10.1109/JSTQE.2009.2035931.

[17] Optical Internetworking Forum, "Implementation Agreement for Coherent CMIS," 2020.

[18] D. van Veen and V. Houtsma, "Strategies for economical next-generation 50G and 100G passive optical networks [Invited]," *J. Opt. Commun. Netw.*, vol. 12, no. 1, p. A95, Jan. 2020, doi: 10.1364/JOCN.12.000A95.

[19] D. Nesset, "PON Roadmap [Invited]," *J. Opt. Commun. Netw.*, vol. 9, no. 1, p. A71, Jan. 2017, doi: 10.1364/JOCN.9.000A71.

[20] S. Hatta, N. Tanaka, and T. Sakamoto, "Low latency dynamic bandwidth allocation method with high bandwidth efficiency for TDM-PON," *NTT Tech. Rev.*, vol. 15, no. 4, 2017.