Point-to-Multipoint Optical Networks Using Coherent Digital Subcarriers

Dave Welch, *Fellow, IEEE*, Antonio Napoli, Johan Bäck, Warren Sande, *Member, IEEE*, João Pedro, *Senior Member, IEEE*, Fady Masoud, Chris Fludger, Thomas Duthel, Han Sun, Steven J. Hand, Ting-Kuang Chiang, Aaron Chase, Atul Mathur, Tobias A. Eriksson, Mats Plantare, Magnus Olson, Stefan Voll, and Kuang-Tsan Wu, *Life Senior Member, IEEE*

(Invited Paper)

***Abstract*—A paradigm shift in optical communication networks is proposed, with the introduction of a new ecosystem of devices and components with the capability of transforming current point-to-point optical networks (with their entailed, limiting, electrical aggregation) into flexible, scalable and cost-effective point-to-multipoint networks. In the new architecture, which better aligns with the hub-and-spoke traffic patterns observed in today's metro and access network segments, interoperability across a variety of transceivers operating at different speeds is achieved using individually routed, digitally generated subcarriers. The first comprehensive demonstration of the technical feasibility of the proposed point-to-multipoint architecture based on digital subcarrier multiplexing is presented, along with the remarkable cost savings and simplification of the network it enables.**

***Index Terms*—Point-to-Multipoint, digital subcarrier multiplexing, 5G, metro aggregation, fronthaul, coherent access.**

I. Introduction

OPTICAL networks have gone through several disruption cycles over the past decades, led by the deployment of breakthrough technologies and resulting in improved network architectures. The seemingly unstoppable erosion in Cost per transported Bit (CpB) has allowed network operators to deliver higher capacities to end users with virtually flat Capital Expenditure (CAPEX) budgets. Notable inventions include the Erbium Doped Fiber Amplifiers (EDFA), Photonic Integrated Circuit (PIC), Wavelength Selective Switch (WSS) – which enabled Reconfigurable Optical Add and Drop Multiplexing (ROADM) – and coherent optical transmission, all of which have been widely embraced by the optical communication industry [1]–[6].

Today's state-of-the-art optical transceivers deliver impressive performance, highly valued in core and submarine applications [7] or any network where fiber is a scarce resource [8]. Yet only limited further improvements can be expected in transceiver performance due to Shannon's limit[1]. Further reductions in the CpB for next generation high-end optical interfaces will have to come in large part from higher symbol rates [7], [11]. This maximizes the throughput over the fiber and considerably reduces the amount of cards in the network, thus simplifying its management [12]. This is in opposition to the improvement obtained by increasing the Signal-to-Noise-Ratio (SNR) and Spectral Efficiency (SE), which can be limited by component imperfections [13]. Nonetheless, there is an upper limit on the practical symbol rate that can be achieved, roughly scaling with the Moore's law.

While higher data rates will reduce the CpB, it is only worthwhile if the operators can effectively utilize the 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 Network (RAN)s and Internet of Things (IoT) enable new use cases [14]–[16]. 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, and typical metro or access network nodes may not be able to fully utilize high-capacity coherent transceivers for several years, leaving the

Manuscript received April 11, 2021; revised June 30, 2021; accepted July 7, 2021. Date of publication July 16, 2021; date of current version August 30, 2021. *(Corresponding author: Antonio Napoli.)*

Dave Welch, Warren Sande, Steven J. Hand, Ting-Kuang Chiang, Aaron Chase, and Atul Mathur are with Infinera Corporation, 6373 San Ignacio Ave, San Jose, CA 95119 USA (e-mail: dwelch@infinera.com; wsande@infinera.com; shand@infinera.com; tchiang@infinera.com; achase@infinera.com; amathur@infinera.com).

Antonio Napoli is with Infinera Limited, c/o Fitzgerald and Law LLP, New Penderel House, London WC1V 7HP, U.K. (e-mail: ANapoli@infinera.com).

Johan Bäck, Tobias A. Eriksson, Mats Plantare, and Magnus Olson are with Infinera, Fredsborgsgatan 24, 117 43 Stockholm, Sweden (e-mail: jbaeck@infinera.com; teriksson@infinera.com; mats.plantare@infinera.com; magnus.olson@infinera.com).

João Pedro is with Infinera Unipessoal Lda, Rua da Garagem 1, 2790-078 Carnaxide, Portugal (e-mail: jpedro@infinera.com).

Fady Masoud, Han Sun, and Kuang-Tsan Wu are with Infinera Limited, Ottawa, Ontario K2K 2X3, Canada (e-mail: fMasoud@infinera.com; hsun@infinera.com; kwu@infinera.com).

Chris Fludger and Thomas Duthel are with Infinera GmbH, Nordostpark, 90411 Nuremberg, Germany (e-mail: cfludger@infinera.com; tduthel@infinera.com).

Stefan Voll is with Coriant, part of Infinera Group, 81541 Munich, Germany (e-mail: svoll@infinera.com).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/JLT.2021.3097163.

Digital Object Identifier 10.1109/JLT.2021.3097163

[1]Commercial transceivers are close to Shannon's capacity. When moving to higher-order modulation formats, e.g., 128/256-Quadrature Amplitude Modulation (QAM), severe component limitations arise [9], [10].

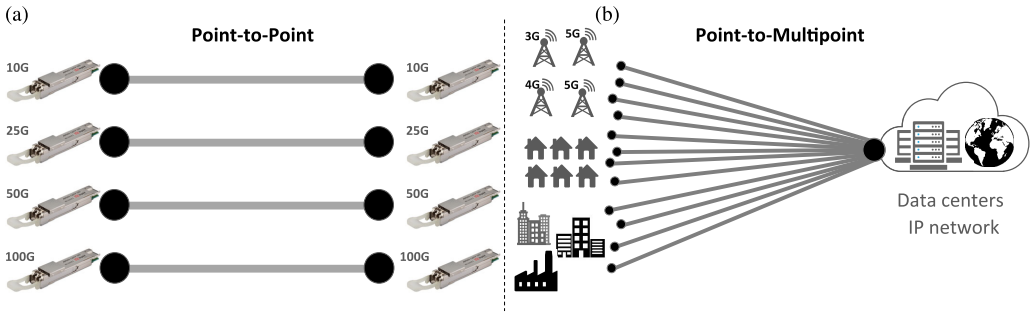

Fig. 1. (a) High-level comparison between Point-to-Point and (b) Point-to-Multipoint traffic patterns.

industry in need of a way to continue reducing CpB within the current network architecture [17].

This paper[2] will outline the use of Digital Subcarrier Multiplexing (DSCM) technology [7], [19], [20] to enable P2MP coherent optical networks with revised Layer 2 / Layer 3 (L_2/L_3) architectures, putting operators on a new trajectory for CpB savings.

We propose a novel network architecture in which:

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

The remainder of the manuscript is structured as follows: Section II provides a comprehensive discussion of the state-of-the-art of available technology; Section III compares P2P and P2MP; Section IV describes the key elements for realizing a P2MP network – DSCM, Digital Signal Processing (DSP), and Network Management (NM); Sections V and VI present relevant application scenarios and the experimental validation of the concept, respectively; Section VII illustrates our vision of next generation optical networks, and Section VIII provides an outlook and conclusions.

II. Technology Background

A. Traffic Evolution

Historically, metro networks were built by telephone operators to create connectivity between local exchanges in metropolitan areas. Most phone calls were local, and required connectivity only to a different local phone exchange in the same metro area. The resulting traffic was highly meshed and P2P in nature, and the few long-distance phone calls were routed to a Central Office (CO) to reach the remote destination [21].

Internet Protocol (IP) data traffic is quite different, as end users are connected to Internet Content Provider (ICP)s hosting their services from a small number of Data Center (DC)s [22]. With the explosion of user data traffic, network patterns have shifted and are now dominated by H&S traffic, regardless of whether the end user is connected to the Internet through a mobile handset, a cable modem, a Fiber-to-the-Home (FTTH) media converter or an Asymmetric Digital Subscriber Line (ADSL) device.

B. Switching Technique Evolution

Transport networks have relied on Time-division Multiplexing (TDM) as the switching technique for several decades. Since it was based on allocating time slots to realize a connection over a given path, TDM was a natural fit to support traffic that was the result of the aggregation of multiple telephone calls, and it was present at both service and transport layers. Early fiber optic networks employed Plesiochronous Digital Hierarchy (PDH). Over time, to overcome key limitations of the former TDM technology, such as complex/expensive multiplexing of lower rate clients and lack of vendor interoperability [23], Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH) gradually replaced PDH. Despite its robustness, SONET/SDH was not natively designed to exploit Dense Wavelength Division Multiplexing (DWDM) as the means to increase greatly the fiber capacity. This limitation was a key driver for the introduction of the Optical Transport Network (OTN) [24], [25].

Internet traffic, given its burst nature, is best supported with a packet-switching technology, such as the IP, since it can take advantage of statistical multiplexing gains. With its rise, IP and IP/Multi-Protocol Label Switching (MPLS) traffic became the dominant protocols traversing telecommunications networks across the globe, while Ethernet emerged from being a low-cost technology for Local Area Network (LAN) to become a replacement for SONET/SDH in Wide Area Network (WAN) [26].

Two main types of infrastructure owners have built today's multilayer core and metro networks. Traditional telecommunication operators have to support a mix of (higher volume) IP/MPLS traffic and (lower volume) legacy services (e.g. leased lines). As a result, their networks usually include IP/MPLS routers and Ethernet switches. The different traffic flows are then multiplexed via OTN switches, which sit on top of a ROADM network. In contrast, ICPs do not have legacy services to support and can simplify the network architecture, preferring to adopt an IP-over-DWDM approach that consists only of routers and ROADMs.

Packet-optical solutions are today becoming prevalent. Developments in pluggable line interfaces that can be deployed

[2]This article is the extension of the invited paper at ECOC 2020 [18].

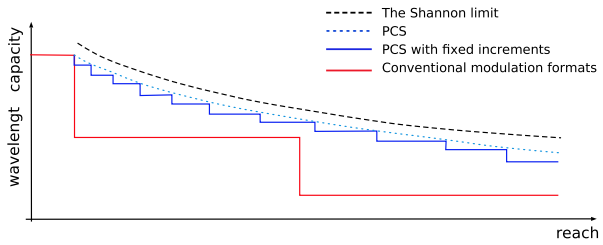

Fig. 2. Comparison of wavelength capacity versus reach for different generations of transceivers.

directly in router ports reinforce the economics of this option in metro networks, by enabling the elimination of dedicated transport boxes and grey interfaces [27]. Fixed access networks still rely on a mix of transmission media, such as copper, coaxial cable and fiber. The rollout of FTTH and Fiber-to-the-Antenna (FTTA) is steadily bringing fiber to users' homes and to antenna towers, respectively. The majority of these deployments use a form of TDM-Passive Optical Network (PON) to carry IP traffic, which, as described above and further discussed in Section V-A, is unlikely to be a scalable and cost-effective option to provide significantly higher data rates.

C. Transceiver Technology Evolution

Until 2010, Direct Detection (DD) systems dominated fiber optics, for any distance and any symbol rate, with a DWDM capacity per fiber as high as ∼1 Tb/s. The explosion of Internet traffic required a rapid introduction of coherent technology and high-order modulation formats, made possible by the fast evolution of Complementary Metal-oxide-semiconductor (CMOS) technology and high-speed Digital-to-Analog Converter (DAC)/Analog-to-Digital Converter (ADC), which enabled advanced DSP [5], [6], [28]–[30] and the transition from Hard Decision (HD)- to Soft Decision (SD)-Forward Error Correction (FEC) [31].

The first generation of coherent transponders achieved 100G–200G at ∼30 GBaud with Polarization Multiplexing (PM)-Quadrature Phase Shift Keying (QPSK)/16-QAM. The next upgrade to 64 GBaud enabled 400G/600G and used 16- and 64-QAM, respectively, and was realized both by compensating for component limitations [32], [33] and by the improvement in silicon integrated circuit technology led by the International Technology Roadmap for Semiconductors (ITRS).

Current commercial coherent systems transmit 800G per wavelength over 1000 km at ∼100 GBaud with Probabilistic Constellation Shaping (PCS) and DSCM [7], [34], [35]. Fig. 2 plots qualitatively the wavelength capacity versus reach for four types of transceivers:

1) a conventional transponder (red solid line) with a limited number of modulation formats;
2) a PCS-based transponder with a finite number of increments (blue solid);
3) an ideal PCS-based transceiver (blue dotted) and
4) a Shannon's ideal transceiver (black dashed).

It is clear from Fig. 2 that we are approaching the Shannon limit in terms of SE and wavelength capacity and that the gains from one generation to the next are diminishing. Similarly, although CMOS and Application-specific Integrated Circuit (ASIC) technology improved from 90 nm [5] to the current 7 nm [7], it also shows signs of a slowdown in the pace of development.

The largest and fastest growing network segments, metro and access, have relatively short link distances and traffic requirements in individual sites are modest. These segments are less well served by today's high-performance coherent transceivers, and the paramount needs are for small footprint, low power consumption and rapid development. The various initiatives to standardize pluggable optics are significantly changing the landscape in the metro segment [27]. Nevertheless, the current industry roadmap for optical transceivers omits what will be presented here as the next key step: the realization of smart pluggables as the basis of next-generation P2MP optical networks.

III. A New Network Paradigm: From P2P to P2MP Optical Networks

This section compares P2P and P2MP architectures in a given typical metro aggregation network, illustrated in Fig. 3. Section III-A discusses the relevant limitations arising from the underlying P2P infrastructures, while Section III-B highlights the potential benefits of using P2MP.

A. Limitations of Point-to-Point Networks

Fig. 3(a) illustrates a common and simplified metro/access network architecture, where the transceivers of N endpoints[3] (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., $N \times 2$ 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 [36].

Once the traffic has arrived at the hub site, it is terminated in N lower-rate 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

[3]In this manuscript, we will interchange the words "leaf/leaves" and "endpoint(s)".

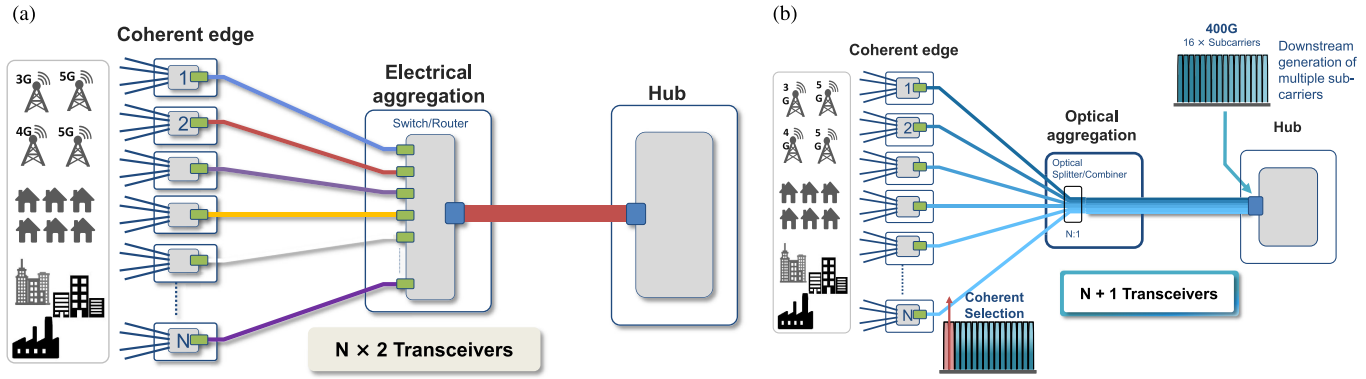

Fig. 3. Example of intermediate aggregation network of N edge nodes for (a) P2P and (b) P2MP. The former requires 2×N low-speed transceivers plus 2 high-speed ones, while the latter requires only N low-speed transceivers plus 1 at high-speed. In (b), we substitute the aggregation layer (and related transceivers) with a passive optical coupler $N:1$ thus simplifying Layer 0 / Layer 1 (L_0 / L_1).

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[4]. A clear limitation of this architecture is that if an endpoint requires more capacity, all transceivers must be replaced. Consequently, an operator would opt either for more frequent truck rolls or having a safety buffer of additional capacity on day one[5]. Neither solution is economically optimal – if an operator wants to align costs to revenues over time, they incur a higher CpB and more frequent site visits.

B. Benefits and Challenges of P2MP 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. 3(b) shows the impact of implementing our proposed P2MP architecture in the same typical metro aggregation network of Fig. 3(a). Here, the electrical aggregation stage of Fig. 3(a) has been replaced by a simple $1:N$ passive optical combiner. The number of devices and stages needed to aggregate and up-speed the traffic to be transported to the next hub is greatly reduced. Note that even if the electrical aggregation device is replaced by a DWDM optical filter, with traditional P2P optics every edge transceiver still needs a corresponding bookended transceiver at the hub. That is not the case with our proposed P2MP optics, because DSCM enables a routing of digital Subcarrier (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. 3(b), only N low-speed transceivers and 1 high-speed one are required, 50% less than in Fig. 3(a). In addition, in the P2MP architecture, multiple low-speed transceivers (spokes) are now directly connected to high-speed ones (hubs), breaking the bookended 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 / 2 (L_1 / L_2) savings by an efficient utilization of the optical transceivers, which can now flexibly provide the exact required amount of capacity.

P2MP can also maximize Layer 3 (L_3) 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. An operator may now choose to deploy 100 Gb/s P2MP pluggable optics in, e.g., 12 leaf nodes, and only turn up a single 25 Gb/s digital subcarrier per node on day one. The 12 SCs can be terminated in a single 400 Gb/s hub router port – a factor of three less than what would be required with P2P optics on day one. As capacity grows in edge nodes and more SCs are turned up over time, the operator can deploy more 400 Gb/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 more generations of P2MP optics ensures that different generations of pluggables can interoperate, and operators can thus define 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 the 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, more than the aggregation function is removed when aggregation devices are taken out of the network. Routers are feature-rich devices and operators use them for a variety of functions.

[4]It is worth mentioning that P2MP is employed in access, e.g., Ethernet PON, Gigabit PON. However, the protocols rely on TDM.

[5]Upgrades of optical systems are costly and prone to human error and service interruptions. The continuous rise in bandwidth demand will increase their frequency, with higher OPEX.

Removing them from networks will take time, as features are 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. Compared to P2P, a P2MP hub optic failure will take down more traffic and affect connectivity to multiple nodes; this can be addressed by using protection or restoration.

A third challenge is that the light paths from a single hub to numerous edge transceivers may include a wide range of losses, necessitating power balancing so that the SCs arrive at the hub with nearly the same power. We envision that the edge P2MP optics will have a wide-range adjustable transmit output power and that a power control loop between the hub and each edge transceiver will be used to equalize the SC powers. In many cases, no additional power balancing elements will be required, but in cases where the path loss difference exceeds the transmit output adjustment range, some external fixed attenuators may be needed.

While P2MP communication is new to coherent optical transmission, it is widely deployed in higher layers. OTN provides multiplexing of payloads for different destinations in a single Optical Transport Unit, and routers provide breakout modes to allow a single higher-rate port to communicate with multiple lower-rate ports over shorter distances. To implement P2MP transmission, the P2MP needs to properly map incoming traffic to the correct digital SC. This could be carried out in accordance with existing standards, such as the breakout modes, or more flexible schemes can be conceived, such as having the P2MP transceiver switch packets based on packet inspection e.g. Virtual LAN (VLAN) tags.

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.

IV. Design of P2MP Networks

A. Point-to-Multipoint Digital Subcarrier Multiplexing 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) pluggables.

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 enable P2MP connectivity with data plane interoperability between coherent pluggables of different capacities – as long as they use the same type of digital SCs. DSCM is a digital communication technique that subdivides the transmitted spectrum into digital SCs [19], [37]. From a networking perspective, it is similar to the Sliceable Bandwidth Variable Transponder (S-BVT) of [38]. Fig. 4 qualitatively compares the spectra of a single 64 GBaud signal carrying a 400G (a), and its equivalent using $16\times$SCs at 4 GBaud with DSCM (b)[6].

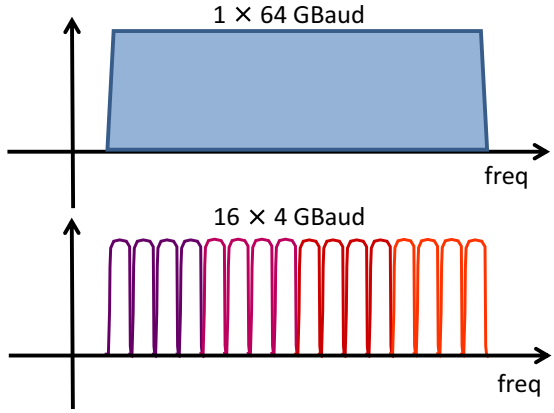

Fig. 4. Pictorial comparison of a single wavelength 400G and its equivalent realized with DSCM with $16\times$SCs.

DSCM has been deployed in P2P links for reducing the impact of linear and nonlinear impairments, maximizing reach, and optimizing the linear compensation jointly to Enhanced Equalization Phase Noise (EEPN) minimization [7], [39]–[41]. The cited works demonstrate that an optimized (lower) symbol rate can reduce the impact of nonlinearities during fiber propagation. DSCM has also been proposed as a way of increasing tolerance against filter cascade by adapting the order of the modulation format used in individual SCs [42] and to reduce power consumption [43].

The SCs are digitally generated at the Transmitter (TX) and do not overlap in the spectrum. Only in the case of ideal Nyquist SCs there is no guard band between them and thus also no overlap in frequency[7]. This allows for the situation illustrated in Fig. 4, where both 400G channels occupy the same 64 GHz bandwidth.

The key differences between single wavelength and DSCM, in terms of DSP, are discussed in Section IV-B. For instance, as a DSCM transceiver operates at a symbol rate R_s/N, where N is the number of digital SCs, the performance might be limited by the value of the laser linewidth Δf [44]. In fact, because of a larger symbol period T_s, the $\Delta f \cdot T_s$ for a DSCM signal with N SCs is N times that of a single wavelength[8]. Hence, a DSCM transceiver is more susceptible to the TX and Receiver (RX) laser variation and characteristics than a single wavelength channel.

B. Digital Signal Processing for Point-to-Multipoint Transceivers

Fig. 5(a-b) compares the DSP at the TX and RX of a typical optical system transmitting a single wavelength carrier (a) and a

[6]Although the number of SCs can be arbitrary, ASIC design is simplified when scaling with a power of two [7].

[7]In reality, a roll-off of the pulse shaper equal to 0 – as in ideal Nyquist SCs – is not possible, and as explained in Section IV-B, a narrow guard band – in our implementation of ∼100 MHz – is required.

[8]A leaf that receives multiple SCs can improve the phase estimation by sharing information between the individual Carrier Phase Estimation (CPE) engines. On the contrary, a hub may have N separate SCs, each with different phase noise contributions. Here, the phase noise information cannot be shared.

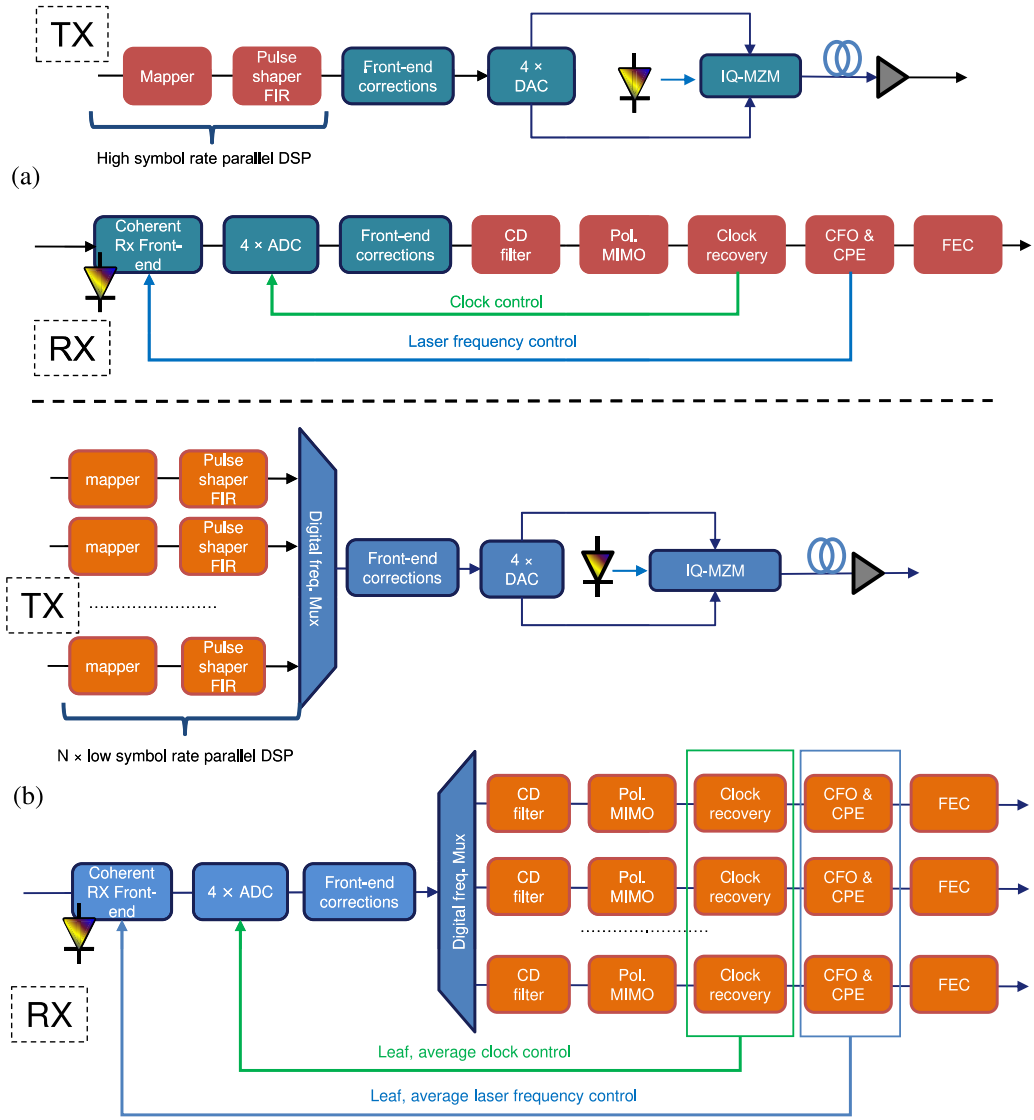

Fig. 5. (a) DSP blocks and system for a single wavelength; (b) the DSP parallelized version for the case of DSCM. CPE: Carrier phase estimation; CFO: Carrier frequency offset. The clock recovery comprises a clock phase detector and an interpolator for fast jitter compensation.

DSCM one (b). 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. 5(b) is the parallelized version of Fig. 5(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 [7].

In both TXs, the DSP comprises a mapper into, e.g., a 16-QAM constellation, and a Finite Impulse Response (FIR) for the Nyquist pulse shaping with roll-off $\rho \sim 0.05$. This filter also performs digital pre-emphasis to mitigate linear impairments such as analog I/Q skew and roll-off. 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.

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, roll-off, 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 [45] for the case of I/Q skew, the DSP within DSCM generates a mirror image which distorts the SC in the mirror position. In the 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 [45].

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 wavelength RX DSP, as the filter length scales with $1/R_s^2$ [7].

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 wavelength and SC [46]. 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 for, we can synchronize in time and frequency by estimating the information derived by the *clock recovery*, and by the CPE together with 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 $4\times$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. 3(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 for time 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 [47] 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.

C. Network Management and Control

The introduction of P2MP optical transmission technology in operator networks raises operational 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 3 rd 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 [48] that hosting device and pluggable vendors must support to achieve interopability.

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 Termination (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 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 include operations such 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 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

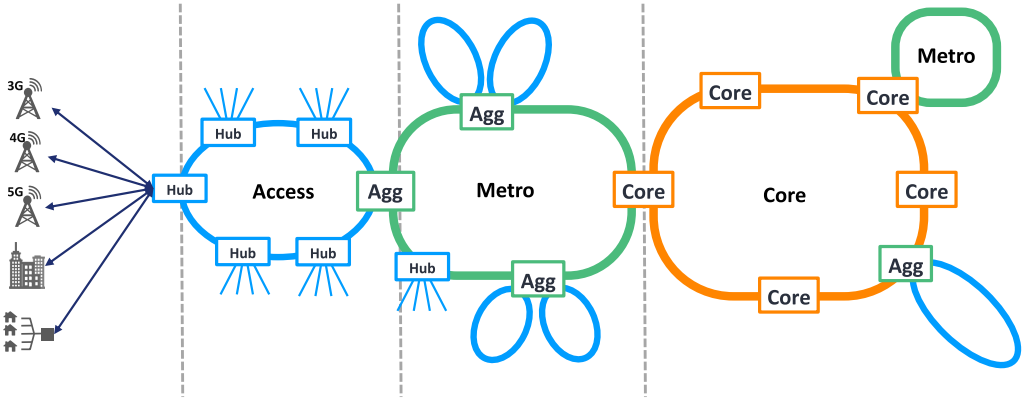

Fig. 6. Generic representation of a telecommunication network from end users to core, passing through access and metro.

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. It is expected that both in-band and out-of-band control channels will be used depending on the operation. 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 plugables 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/Rest 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.

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 to run at their own pace.

V. Network Application Scenarios

The following discussion on network applications for P2MP is central to the argument presented in this paper. Fig. 6 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. 6: (i) PON in Section V-A; (ii) metro aggregation in Section V-B and (iii) mobile fronthaul in Section V-C. Thereafter, we highlight the changes required to re-architect the network in view of P2MP with DSCM. In Section V-A, we also provide a comparison between TDM and Frequency Division Multiplexing (FDM), highlighting the main differences when applied to PON. Section V-D 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.

A. *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[9]. PONs have been widely deployed using current cost-effective DD systems [49]; their success is largely due to the adoption of the P2MP approach, an indication of the value of this optical networking architecture. The transmission is here realized with TDM, and the transceivers are manufactured in volumes exceeding millions of units per year, as discussed in [50] 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.

[9]With TDM PON, end users can receive stream down to Mb/s speed, by applying statistical TDM. This differs from our proposed coherent with DSCM approach, as the laser instability does not allow it to operate at very low symbol rate.

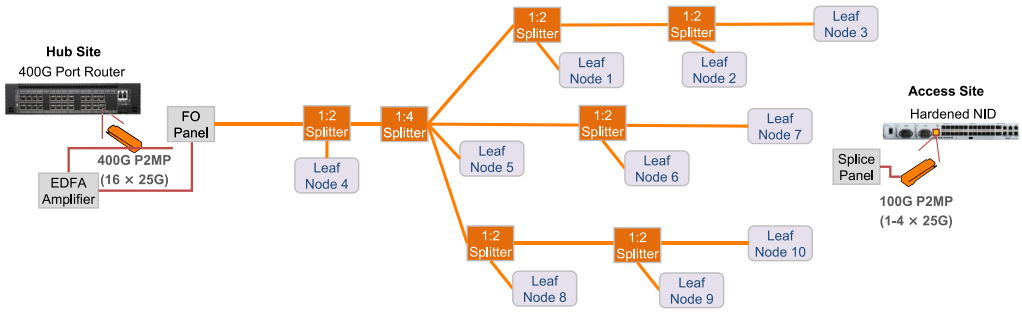

Fig. 7. PON overlay realized with a P2MP transceiver where the 400G can communicate with the 100G leaves.

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 it is not possible to exploit 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.

In addition, as discussed in Section IV-B, if the leaf modules are placed at different geographic locations in the aggregation stage, a TDM approach might suffer synchronization issues during the uplink[10], which cannot be compensated by a proper DSP. Moreover, burst transmission induces transients on the electrical current and on the laser output power and affects the laser stability. Because of this, it is not suitable for DWDM, and as a result, only a few wavelengths will be available.

Finally, latency and RX sensitivity are also critical aspects in PON design. The former is severely affected by the usage of TDM [51], while the latter is not sufficiently high to enable long distances with existing DD.

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

To provide the required flexibility and capacity, we propose an FDM-based approach for P2MP as discussed in Section IV. 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 $\geq$ 100G.

Fig. 7 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.

This type of architecture can initially be adopted for e.g. business 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 does, and as traffic demands in access networks grow, the proposed implementation provides a number of other advantages, as discussed above.

B. Metro Aggregation

Fig. 8 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.

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% [17].

The hub site configuration is a core router with support for 400 Gb/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 L_3 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 L_1, but also at L_2/L_3. The total capacity required in the

[10]PONs become less optimal as required capacities and distances increase and current commercially available solutions are limited in terms of reach to ~25 km between the splitter and OTN.

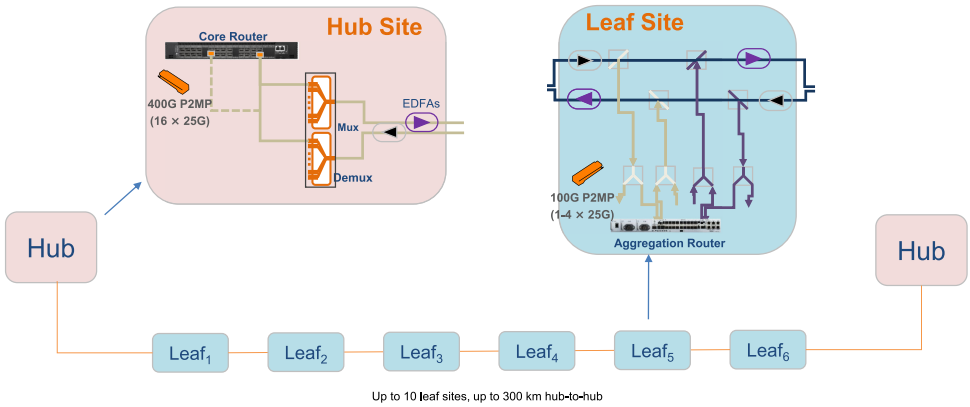

Fig. 8. Application of the P2MP transceiver to a metro aggregation network, with 2×400G hubs and 6×100G leaves.

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 400 Gb/s pluggable optics, a single 400 Gb/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 L_2 and L_3. 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.

C. 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, economical, 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 deployed 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. 9 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.

D. 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, our proposed P2MP 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.

VI. Experimental Validation of the P2MP Concept

This section reports on the experimental verification of the P2MP concept and its main outcomes. The experiment was carried out in collaboration with BT at their labs in the U.K.

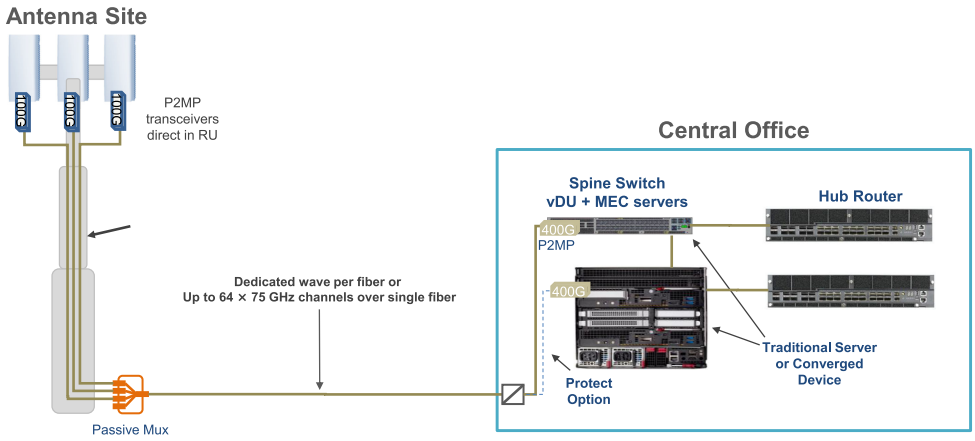

Fig. 9. Wireless fronthaul.

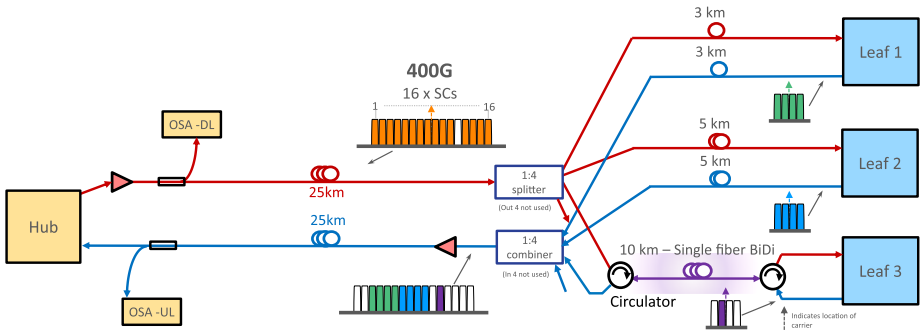

Fig. 10. Experimental test bed used to validate the P2MP concept at BT's lab. The hub can transmit 16 SCs at 25G and the leaves 4 SCs each. Note: leaf_3 sits on a single bi-directional fiber.

and comprises a PON overlay test bed[11]. The aim was to assess the possibility of adapting capacities to instantaneous needs, including the case of asymmetric traffic patterns as visualized in Figs. 10 and 11.

A. *Description of the Test Bed*

Fig. 10 depicts the test bed herein considered. Starting from the left side, the hub is equipped with a 400G P2MP transceiver that operates in multiples of 25G, capable of transmitting up to 16 SCs each with 4 GHz bandwidth employing 16QAM modulation. A portion of the hub's output and input signals is measured by two high resolution Optical Spectrum Analyzer (OSA)s (see Fig. 10), and their spectrum captures are shown in Fig. 11(a–e). The hub is capable of communicating with the remote leaves to set the different configurations that will be discussed in Section VI-B. Next, we place an EDFA[12] to compensate for the loss of the link, which includes 25 km of standard Single Mode Fiber (SMF-28) before reaching the $1:4$ passive optical splitter that connects to the three leaves. These are located at 3, 5, and 10 km, respectively, from the splitter and they are also connected with SMF-28. Each leaf has a 100G transceiver capable of detecting and receiving 4 SCs of 25G each, with 16QAM modulation. Note that each leaf sees the full spectrum of 16 SCs, and the center frequencies of the different leaves are locked to the hub laser with different offsets, depending on which of the 16 SCs they are assigned to receive and transmit on.

In this experiment, the hub transmit power is a constant -9 dBm for the 16 SC signals, or -21 dBm per SC[13]. After EDFA gains and path losses in the experimental setup, the leaf RXs see received powers in the range of -8 dBm to -12 dBm total, or -20 dBm to -24 dBm per SC. The leaf transmit powers are configured such that each SC is received at the hub with close to the same power (within 1 dBm). Leaf transmit powers are between -9 dBm and -12 dBm per SC. After path losses and EDFA gains, the hub receive powers are approximately -17 dBm per SC. In this experiment, the hub receives from 2 to 9 SCs from the leaf TXs, resulting in total received power at the hub from -14 dBm to -7.5 dBm.

For leaf_1 and leaf_2, a dual fiber configuration is used and hence the same SC assignment is defined for both *downlink*

[11] Similar experiments were carried out at different customer labs, as reported in [52], [53]

[12] The use of an EDFA here was only required because of the limitations of the proof-of-concept hardware used for this experiment: a) limited TX power b) limited receiver sensitivity. With production hardware, we will have sufficient link budget to support this configuration without EDFAs

[13] For the last 2 configurations ⓓ and ⓔ shown in Fig. 11, and described below, one hub SC is removed from the transmit spectrum, leaving a 15 SC signal. Total hub TX power is still held at -9 dBm, resulting in a small change in per-SC power.

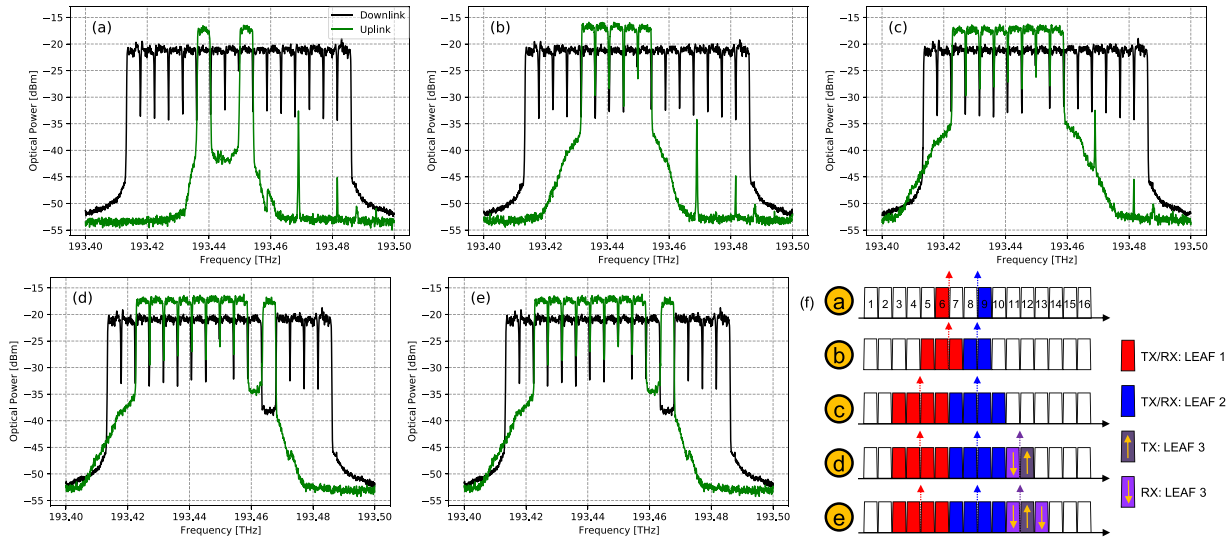

Fig. 11. Experimental validation of: (a) instantiating 25G traffic; (b) instantaneously increasing throughput by lighting up additional SCs with different data rate for different leaves, in this case, 75G on leaf_1 and 50G on leaf_2; (c) frequency tuning of leaves to accommodate additional traffic; (d-e) showing transmission over bidirectional single fiber bidirectional and symmetric (d) and asymmetric (e) traffic in up and downlink. Note that for leaf_3 different SCs are used in up- and down-direction and that the hub turns off the SC used by leaf_3 in the up-link to avoid in-band Rayleigh scattering; (f) reports the sequence of the configured traffic variations.

(hub to leaf) and *uplink* (leaf to hub) to maximize the spectral usage. However, leaf_3 sits at the end of a 10 km single-fiber link with circulators to pass traffic in the two directions. For this case, we can use different SCs in the uplink and downlink direction, if needed, to avoid excess noise coming from Rayleigh back-scattering. The signals from the leaves in the uplink direction are combined by a passive $1:4$ optical combiner at the splitter node location and then propagated back to the hub over 25 km of SMF-28.

B. Experimental Results

Using the test bed described in Fig. 10, different traffic configurations were tested via software running on the hub and the leaf nodes. The hub sends *out-of-band* signals to the leaf nodes with instructions on how to configure themselves for the specific SCs they are assigned. The list of configurations with the active SCs is reported in Fig. 11(f). The operation of the switching of the SCs is a key aspect as reported in Fig. 3 of [18], which shows the power versus time-frequency of a 400G P2MP transceiver, and illustrates the ability to turn on/off the configurations of SCs as a function of the time, i.e., because of traffic demand variations.

In the first configuration ⓐ of Fig. 11(a), the hub assigns SC_6 to leaf_1, and SC_9 to leaf_2, so that the traffic flow can start. Figs. 11(a–e) visualize the two spectra of the two OSAs as shown in Fig. 10: in black we plot the spectrum transmitted from the hub – 16 SCs at 25G – to all leaves; in green that of the SCs being received from the leaves over the combined path. In Fig. 11(a) we are receiving SC_6 from leaf_1 and SC_9 from leaf_2.

The ability to instantaneously provide the requested capacity will be a requirement in next generation optical networks. Our proposed solution achieves this by lighting up more SCs. Consequently, different leaves might have different data rates, as reported in Fig. 11(b), configuration ⓑ. Here, starting from ⓐ, we increase the throughput of $\text{leaf}_{1,2}$. This is realized by activating $\text{SC}_{5,7}$ (+ 50G) at leaf_1, and simultaneously SC_8 (+ 25G) at leaf_2. The traffic has been increased in steps of 25G to 75G at leaf_1 and 50G at leaf_2.

Apart from increasing capacity, another important network management operation is the shift of the leaf center frequency, as reported in configuration ⓒ of Fig. 11(c), where we shift the SCs of leaf_1 by 8 GHz so that the SCs now occupy positions {3,4,5,6}. Leaf_1 moves its laser to the new frequency, and starts transmitting 4 SCs. Leaf_2 then activates two more SCs, {7,10}. We have further hitlessly increased the in-service capacity of leaf_1 (+25G) by adding one SC, and of leaf_2 (+50G) by adding two SCs.

The previous configurations showed dual-fiber traffic flows. Next, in configuration ⓓ of Fig. 11(d), we demonstrate operation over single fiber, to/from leaf_3. We add SC_{11} for the downlink, and SC_{12} for uplink. In order to do that, the hub brings down SC_{12} from its transmit spectrum to free the frequency slot for the uplink coming from leaf_3. This upgrade of the transmission, which only concerns leaf_3 and the hub, does not impact the traffic being sent to and received from the other leaves.

The last configuration ⓔ demonstrates asymmetrical traffic over single fiber between the hub and leaf_3. Here there are two SCs ($\text{SC}_{11,13}$) that provide a total capacity of 50G from the hub to leaf_3 in downlink and one SC, (SC_{12}) providing 25G between leaf_3 and the hub in uplink, thus creating an asymmetric traffic pattern. The spectrum reported in Fig. 11(e) does not change from (d), because the hub was already sending on $\text{SCs}_{11,13}$ and leaf_3 was already transmitting on SC_{12}. The change is that leaf_3 starts to receive the additional SC_{13}. This leads to the configuration where leaf_3 is receiving two non-contiguous SCs ($\text{SC}_{11,13}$) for a total of 50G, and transmitting back one single SC to the hub, i.e., (SC_{12}) for 25G.

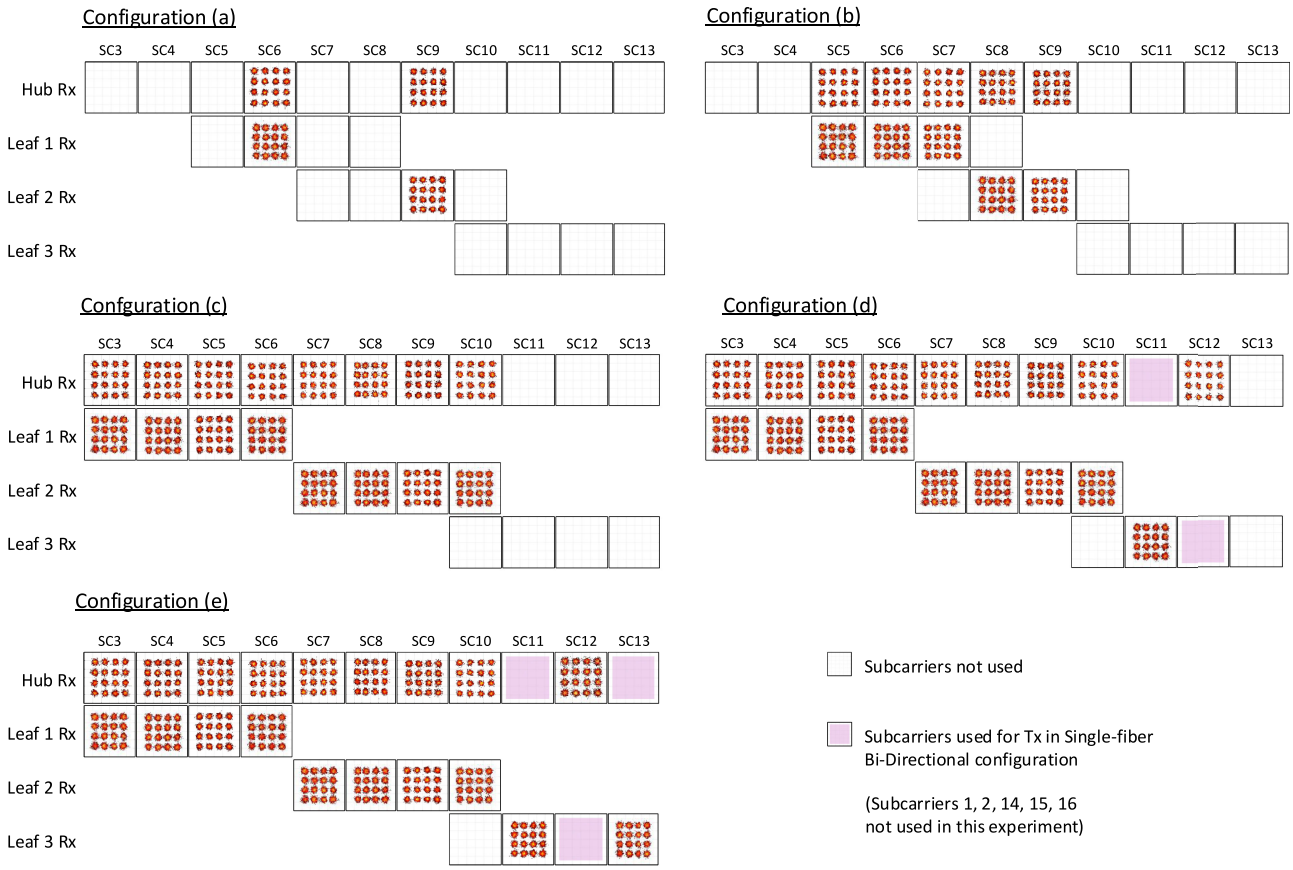

Fig. 12. Experimental signal constellations of the different configurations described in Fig. 11. Note that the hub transmits 16 SCs. However, only the one used in this experimental demonstration are shown in the figure.

To summarize, we started with two SCs, one each for leaf_1 and leaf_2. Then, without any hardware modification, we dynamically increased the capacity of leaf_1 by adding two additional SCs, and the capacity of leaf_2 by adding one SC. We also tuned leaf_1 to free the frequency slot for additional SCs for leaf_1 and leaf_2. Finally, we added leaf_3 to show operation over a single fiber with symmetrical and asymmetrical traffic patterns, and to demonstrate how operation over a fiber pair and a single fiber can co-exist and be connected to the same hub. Finally, Fig. 12 shows the constellation diagrams for each SC based on the configuration illustrated in Fig. 11(f). All SCs were detected with pre-FEC BER that would result in error-free performance after FEC.

VII. Envisioning a Simplified and Modern Telecommunication Infrastructure

In this work, we proposed the use P2MP optics to considerably simplify telecommunication networks. We have shown that DSCM with individual framing enables communication between low- and high-speed transceivers, making separate electrical aggregation redundant, as illustrated in Fig. 3(b). We envision a phased introduction of P2MP optics in operator networks, starting from today's architecture and gradually moving towards the architecture described in Fig. 14 in four steps:

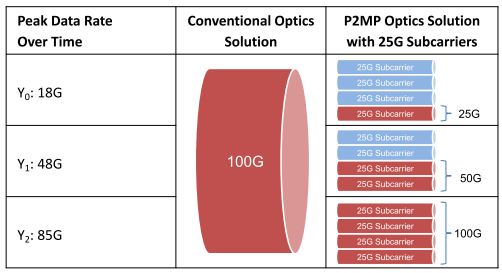

Fig. 13. Visual example of network planning with P2P and P2MP over three years.

(i) The use of DSCM in P2P mode for high capacity links or breakout P2MP mode for lower capacity links. Leveraging flexible software management of individual SCs, capacity can be instantaneously adapted to bandwidth requirements, turning off blocks associated with unused SCs in the DSP and resulting in OPEX savings [43]. Power savings are illustrated by the light blue bars in Fig. 13, in a scenario where SCs are commissioned over a 3-year growth model and compared to P2P optics. The flexibility of DSCM enables operation over legacy line systems with optical filters, greatly simplifying adaptation. The total symbol rate can be configured in steps of 4 GBaud, providing a granularity more than 3×

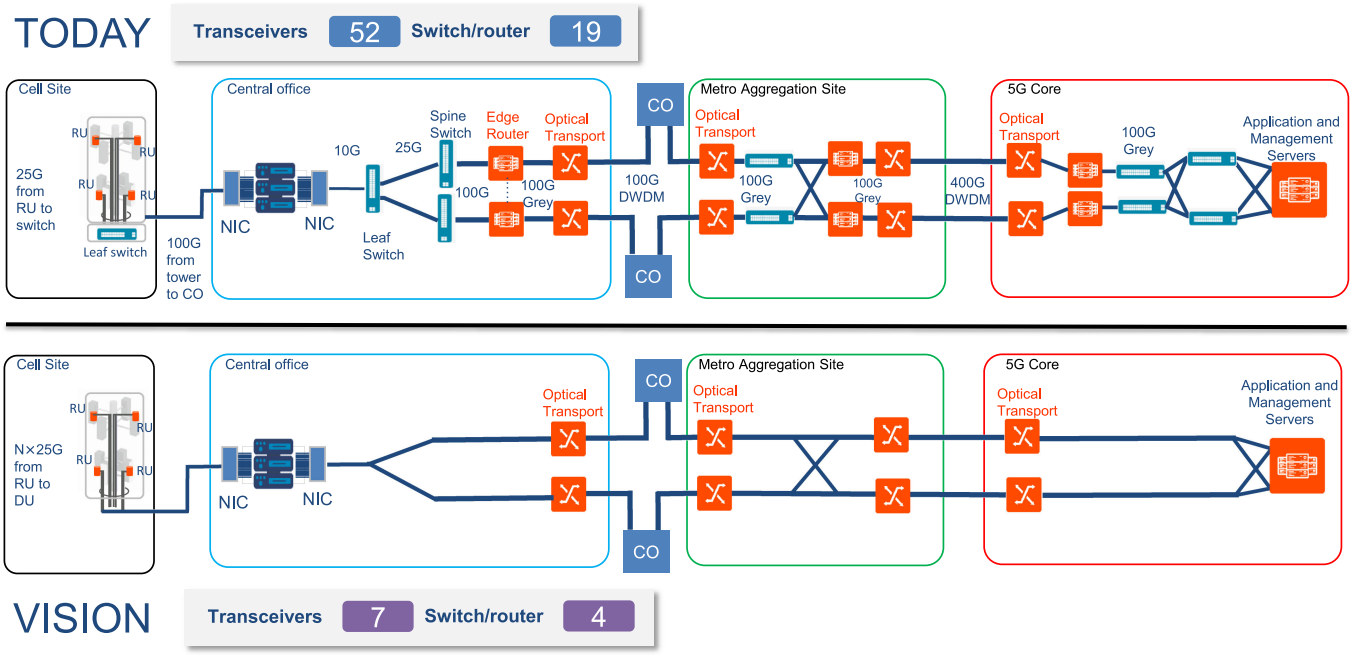

Fig. 14. From today to tomorrow: an example of network simplification.

larger than the most flexible 400G transceivers available on the market in 2021.

(ii) Exploiting the flexible P2MP capability within existing network architectures at higher network layers. L_2/L_3 port speeds can be optimized to total add/drop traffic, providing substantial CAPEX savings by reducing the volume of optical interfaces and leveraging higher data rate transceivers with lower CpB. CAPEX savings as high as 76% have been demonstrated in a 5-year model based on a national footprint of metro networks with realistic traffic requirements, where P2MP pluggable optics deployed directly in core and aggregation routers over a filterless line system were compared against a traditional metro DWDM architecture built with P2P transponders at 100G, 200G and 400G data rates and a ROADM line system [17].

(iii) Optimization of the L_2/L_3 network, using the inherent aggregation of the DSCM to remove intermediate aggregation devices such as leaf switches or aggregation routers, resulting in a flatter network architecture with associated TCO savings. Examples include: a) collapsing metro core and metro access networks and b) simplifying CRAN configurations as shown in Fig. 14. The extent to which L_2/L_3 networks can be collapsed depends on the network geography, as measured by distance from edge to core, and the capacities of available P2MP optical transceivers. In an extreme example, an 800G P2MP transceiver in a core router could aggregate traffic from 25G P2MP transceivers in 32 access sites.

(iv) Integrating P2MP optics directly in other devices, such as Radio Unit (RU)s, Remote PHY Device (RPD)s, and Network Interface Card (NIC)s used in cloud servers, **entirely removing standalone electrical aggregation devices**. While this is technically possible with P2P optics, it may not be economically feasible for aggregation networks with H&S traffic flows. It may also require other networking functions associated with L_2/L_3 to be virtualized or performed by a different device, such as a smart pluggable P2MP optic with packet features integrated in the DSP.

In Fig. 14 we look at the flow of data from an RU in a cell site all the way back to a 5G core located in a DC as it is realized today, and how it will be realized by employing the solution we are proposing. In this example, traffic from the RU to the DU is unprotected, while the traffic from the DU back to the 5G core is protected.

In today's architecture, a fronthaul packet router in the cell site aggregates traffic from multiple RUs for transport back to the DU in the CO. The CO has a leaf and spine switch to aggregate traffic from several DUs to an edge router. From the edge router, a metro access DWDM system with 100 Gb/s wavelengths takes the traffic back to metro aggregation site, where it is further aggregated onto metro core DWDM system with 400G wavelengths. Finally back at the 5G core site, traffic goes through another service router before it is finally handed off to the 5G core itself.

In this journey from the antenna to the core, there are only three places where the bits are processed: in the RU, in the DU and in the 5G core. Any latency-sensitive applications – such as baseband functions or Multi-Access Edge Computing (MEC) applications – must take place in a location close to the end user. Everything else can be located in a central location, where operators can rely on economies of scale to reduce CAPEX and OPEX.

With smart P2MP pluggable optics, we envision a network where the optical transceivers are integrated in the functional

units themselves, in this case the RUs, DUs and servers where the 5G core is hosted. In the cell site, a digital aggregation function in the form of a fronthaul packet router can be eliminated by integrating P2MP directly in the RUs. RUs host P2MP optics directly, and capacity can be provisioned in steps of 25 Gb/s over existing single-fiber infrastructures. As the fronthaul traffic scales due to the use of massive MIMO or the emergence of RUs supporting multiple sector antennas and frequency bands, capacity can be sized accordingly.

In the CO, the backhaul can be simplified by removal of leaf and spine switch architecture connecting the DUs to edge routers, and the edge routers themselves can also be eliminated. Any routing functions required can be integrated directly in the same server as the DU function is hosted in. The P2MP optics are hosted in the network interface cards in the MEC server, and the P2MP functionality can be leveraged to have a single DU connect to many RUs.

The metro aggregation site is bypassed optically, removing all equipment except for the optical line system components, subject to link engineering. At the core site, simplifications are similar to those in the CO, with the core application servers directly hosting P2MP optics and any required routing functions. The P2MP functionality enables a single network interface card in the application server to connect to many DUs.

Optical transponders are removed entirely from the architecture, reducing the dedicated optical mux and potentially WSS functions. All gray optics connecting different boxes also disappear. The result is a significant reduction in the number optical interfaces as described in the table in Fig. 14. The envisioned architecture is greatly simplified, delivering savings across multiple network layers.

VIII. Outlook and Conclusions

This work introduces and demonstrates the architectural concept of point-to-multipoint in optical networks. We believe this will be the next breakthrough in optical communication, following the last great inflection point of the introduction of coherent optical technology.

We motivate this paradigm shift by a comparison with the point-to-point approach, which has so far predominated in optical telecommunications, though not in wireless. We highlight the fact that 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.

We provide an in-depth discussion of the technology needed in terms of hardware, digital signal processing, and network management. We analyze several application scenarios, showing that this solution can be highly beneficial, in particular for metro aggregation, access, and fronthaul network segments.

In Section VI, we report the results of our experimental validation of the concept in a lab experiment carried out jointly with BT. In this experiment, we tested some of the features that point-to-multipoint systems must have. In Section VII we show our vision for future optical telecommunication networks and the phases to realize them.

We believe that the proposed solution shows what the network will be, as it represents an excellent match between network architecture and traffic patterns. We expect its deployment will be facilitated by the fact the proposed architecture can co-exist with current P2P systems. According to market analyses, both 5G and beyond, and the high capacity that will be needed in access networks, will continue to accelerate the demand for bandwidth. In the context of hub-and-spoke and high-capacity traffic patterns, we believe that only a point-to-multipoint technology supported by coherent transceivers employing frequency division multiplexing can satisfy cost-effectively the future requirements of network operators.

Acknowledgment

We would like to thank the BT team for their support during the experimental work, in particular we would like to thank Andrew Lord, Md Asif Iqbal, and Paul Wright. We would like also to thank the two anonymous reviewers for their helpful comments on earlier drafts of the manuscript.

References

[1] E. Desurvire, "Erbium-doped fiber amplifiers," in *Princ. Appl.*, 1992.
[2] D. F. Welch *et al.*, "Large-scale InP photonic integrated circuits: Enabling efficient scaling of optical transport networks," *IEEE J. Sel. Topics Quantum Electron.*, vol. 13, no. 1, pp. 22–31, Jan./Feb. 2007.
[3] S. Gringeri, B. Basch, V. Shukla, R. Egorov, and T. J. Xia, "Flexible architectures for optical transport nodes and networks," *IEEE Commun. Mag.*, vol. 48, no. 7, pp. 40–50, Jul. 2010.
[4] G. Bennett, K.-T. Wu, A. Malik, S. Roy, and A. Awadalla, "A review of high-speed coherent transmission technologies for long-haul DWDM transmission at 100G and beyond," *IEEE Commun. Mag.*, vol. 52, no. 10, pp. 102–110, Oct. 2014.
[5] H. Sun, K.-T. Wu, and K. Roberts, "Real-time measurements of a 40 Gb/s coherent system," *Opt. Exp.*, vol. 16, no. 2, pp. 873–879, 2008.
[6] M. S. Alfiad *et al.*, "111-Gb/s transmission over 1040-km field-deployed fiber with 10G/40G neighbors," *IEEE Photon. Technol. Lett.*, vol. 21, no. 10, pp. 615–617, May 2009.
[7] H. Sun *et al.*, "800G DSP ASIC design using probabilistic shaping and digital sub-carrier multiplexing," *J. Lightw. Technol.*, vol. 38, no. 17, pp. 4744–4756, 2020.
[8] J. Pedro, N. Costa, and S. Sanders, "Scaling regional optical transport networks with pluggable and integrated high-capacity line interfaces," in *Proc. Opt. Fiber Commun. Conf. Opt. Soc. Amer.*, 2021, pp. 1–3.
[9] S. Varughese, D. Lippiatt, S. Tibuleac, and S. E. Ralph, "Frequency dependent ENoB requirements for 400G/600G/800G optical links," *J. Lightw. Technol.*, vol. 38, no. 18, pp. 5008–5016, 2020.
[10] G. Khanna *et al.*, "Single-carrier 400G 64QAM and 128QAM DWDM field trial transmission over metro legacy links," *IEEE Photon. Technol. Lett.*, vol. 29, no. 2, pp. 189–192, Jan. 2017.
[11] W. Heni *et al.*, "Ultra-high-speed 2:1 digital selector and plasmonic modulator IM/DD transmitter operating at 222 GBaud for intra-datacenter applications," *J. Lightw. Technol.*, vol. 38, no. 9, pp. 2734–2739, 2020.
[12] T. Gerard *et al.*, "Relative impact of channel symbol rate on transmission capacity," *J. Opt. Commun. Netw.*, vol. 12, no. 4, pp. B1–B8, 2020.
[13] T. Drenski and J. C. Rasmussen, "ADC & DAC technology trends and steps to overcome current limitations," in *Proc. Opt. Fiber Commun. Conf. Expo.*, 2018, pp. 1–3.
[14] "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
[15] M. Schiano, A. Percelsi, and M. Quagliotti, "Flexible node architectures for metro networks," *J. Opt. Commun. Netw.*, vol. 7, no. 12, pp. B 131–B140, 2015.

[16] "Network Traffic Insights in the Time of COVID-19," [Online]. Available: https://www.nokia.com/blog/network-traffic-insights-in-the-time-of-covid-19-june-4-update/

[17] 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.

[18] D. F. Welch, "Disruption cycles for optical networks: How point to multipoint coherent optics can transform the cost and complexity of the optical network," in *Proc. Eur. Conf. Opt. Commun.*, 2020, pp. 1–3.

[19] D. Krause, A. Awadalla, A. S. Karar, H. Sun, and K.-T. Wu, "Design considerations for a digital subcarrier coherent optical modem," in *Proc. Opt. Fiber Commun. Conf. Opt. Soc. Amer.*, 2017, pp. Th 1D-1.

[20] Y. Zhang, M. O'Sullivan, and R. Hui, "Digital subcarrier multiplexing for flexible spectral allocation in optical transport network," *Opt. Exp.*, vol. 19, no. 22, pp. 21 880–21 889, 2011.

[21] M. Sinclair, "Improved model for European international telephony traffic," *Electron. Lett.*, vol. 30, no. 18, pp. 1468–1470, 1994.

[22] A. Dwivedi and R. Wagner, "Traffic model for USA long-distance optical network," in *Proc. Opt. Fiber Commun. Conf. Opt. Soc. Amer.*, 2000, pp. 1–3.

[23] "G.841: Types and Characteristics of SDH Network Protection Architectures," [Online]. Available: https://www.itu.int/rec/T-REC-G.841-199810-I/en

[24] "G.872: Architecture of the Optical Transport Network," [Online]. Available: https://www.itu.int/rec/T-REC-G.872

[25] D. Cavendish, "Evolution of optical transport technologies: From SONET/SDH to WDM," *IEEE Commun. Mag.*, vol. 38, no. 6, pp. 164–172, Jun. 2000.

[26] N. Bitar, "Transport network evolution: From TDM to packet," in *Proc. Opt. Fiber Commun. Conf. Opt. Soc. Amer.*, 2011, Paper OThR1.

[27] "Open ZR MSA Technical Specification," [Online]. Available: https://openzrplus.org/

[28] Y. Han and G. Li, "Coherent optical communication using polarization multiple-input-multiple-output," *Opt. Exp.*, vol. 13, no. 19, pp. 7527–7534, 2005.

[29] K. Roberts *et al.*, "Performance of dual-polarization QPSK for optical transport systems," *J. Lightw. Technol.*, vol. 27, no. 16, pp. 3546–3559, 2009.

[30] M. Kuschnerov *et al.*, "DSP for coherent single-carrier receivers," *J. Lightw. Technol.*, vol. 27, no. 16, pp. 3614–3622, 2009.

[31] K. Sugihara, K. Ishii, K. Dohi, K. Kubo, T. Sugihara, and W. Matsumoto, "Scalable SD-FEC for efficient next-generation optical networks," in *Proc. 42nd Eur. Conf. Opt. Commun.*, 2016, pp. 1–3.

[32] A. Napoli *et al.*, "Novel DAC digital pre-emphasis algorithm for next-generation flexible optical transponders," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2015, pp. 1–3.

[33] Z. Zhang *et al.*, "Coherent transceiver operating at 61-Gbaud/s," *Opt. Exp.*, vol. 23, no. 15, pp. 18988–18995, Jul 2015.

[34] J. Cho and P. Winzer, "Probabilistic constellation shaping for optical fiber communications," *J. Lightw. Technol.*, vol. 37, no. 6, pp. 1590–1607, 2019.

[35] F. Buchali, F. Steiner, G. Böcherer, L. Schmalen, P. Schulte, and W. Idler, "Rate adaptation and reach increase by probabilistically shaped 64-QAM: An experimental demonstration," *J. Lightw. Technol.*, vol. 34, no. 7, pp. 1599–1609, 2016.

[36] T. Ohno, *Toyota Production System: Beyond Large-Scale Production*. Boca Raton, FL, USA: CRC Press, 1988.

[37] T. A. Eriksson, F. Buchali, W. Idler, L. Schmalen, and G. Charlet, "Electronically subcarrier multiplexed PM-32QAM with optimized FEC overheads," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2017, pp. 1–3.

[38] N. Sambo *et al.*, "Next generation sliceable bandwidth variable transponders," *IEEE Commun. Mag.*, vol. 53, no. 2, pp. 163–171, Feb. 2015.

[39] F. Buchali *et al.*, "Study of electrical subband multiplexing at 54 GHz modulation bandwidth for 16QAM and probabilistically shaped 64QAM," in *Proc. 42nd Eur. Conf. Opt. Commun.*, 2016, pp. 1–3.

[40] P. Poggiolini, Y. Jiang, A. Carena, G. Bosco, and F. Forghieri, "Analytical results on system maximum reach increase through symbol rate optimization," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2015, pp. 1–3.

[41] L. B. Du and A. J. Lowery, "Optimizing the subcarrier granularity of coherent optical communications systems," *Opt. Exp.*, vol. 19, no. 9, pp. 8079–8084, 2011.

[42] T. Rahman *et al.*, "Digital subcarrier multiplexed hybrid QAM for data-rate flexibility and ROADM filtering tolerance," in *Proc. Opt. Fiber Commun. Conf.*, 2016, Paper Tu3K.5.

[43] L. Velasco *et al.*, "Autonomous and energy efficient lightpath operation based on digital subcarrier multiplexing," *IEEE J. Sel. Areas Commun.*, early access, 2021, doi: 10.1109/JSAC.2021.3064698.

[44] 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," in *Proc. Opt. Fiber Commun. Conf. Expo. Nat. Fiber Optic Eng. Conf.*, 2013, pp. 1–3.

[45] 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 *Proc. Opt. Fiber Commun. Conf. Opt. Soc. Amer.*, 2016, pp. W 4A-5.

[46] S. J. Savory, "Digital filters for coherent optical receivers," *Opt. Exp.*, vol. 16, no. 2, pp. 804–817, 2008.

[47] B. Spinnler, "Equalizer design and complexity for digital coherent receivers," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 5, pp. 1180–1192, Sep./Oct. 2010.

[48] "Implementation Agreement for Coherent CMIS," [Online]. Available: https://www.oiforum.com/wpcontent/uploads/OIF-C-CMIS-01.0.pdf

[49] D. van Veen and V. Houtsma, "Strategies for economical next-generation 50G and 100G passive optical networks," *J. Opt. Commun. Netw.*, vol. 12, no. 1, pp. A 95–A103, Jan. 2020.

[50] D. Nesset, "PON roadmap," *J. Opt. Commun. Netw.*, vol. 9, no. 1, pp. A 71–A76, 2017.

[51] 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, pp. 1–7, 2017.

[52] "Infinera and American Tower Complete Live Network Demonstration of First Point-to-Multipoint Coherent Optical Transmission in Latin America," [Online]. Available: https://www.infinera.com/wp-content/uploads/pr20210315-Infinera-and-American-Tower-Complete-Live-Network-Demonstration-First-in-Latin-America.pdf

[53] "Virgin Media Trials Infinera's Cutting Edge Multi-Gigabit Network Technology," [Online]. Available: https://www.infinera.com/wp-content/uploads/pr20210303-Virgin-Media-Trials-Infineras-Cutting-edge-Multi-gigabit-Network-Technology.pdf