VIRTUAL EXPERIENCE
OCTOBER 11-14

2021 Fall
Technical Forum
SCTE • NCTA • CABLELABS

Delivering Access Beyond 10G

Coherent Subcarrier Aggregation as Backhaul for Next-Generation R-OLT, RMD, and Wireless

A Technical Paper prepared for SCTE by

Colin Howlett
Chief Technology Officer
Vecima
771 Vanalman Ave, Victoria, BC, Canada V8Z 3B8
+1 (250) 881-6235
colin.howlett@vecima.com

Aaron Chase, Infinera
achase@infinera.com

Antonio Napoli, Infinera
anapoli@infinera.com

Kevin A. Noll, Vecima
kevin.noll@vecima.com

Jay Rolls, Pacband
jrolls@pacband.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

Over the last decade, the demand for network capacity has been continuously growing. This is particularly true in the last mile of the access and metro networks. This has been fueled by an unstoppable sequence of web-based applications such as video streaming, cloud services, mobile transport, and machine-to-machine communication. Recently, the traffic demand further increased driven by the ongoing pandemic though the impact has been, so far, manageable.

This growth was enabled by the wide deployment of broadband connectivity. This is an evolving landscape, which is now being addressed by the realization of newer technologies such as 10 gigabit per second (10G) capable Passive Optical Network (PON), DOCSIS 3.1, and fifth generation (5G) wireless, each offering access to subscribers at gigabit speeds. At the same time, the consolidation of hub sites and plant extensions are driving deployment of Distributed Access Architecture (DAA). As a result, the Converged Interconnect Network (CIN), the glue that connects these diverse access networks to the operator's headends, is being asked to provide more capacity to an increasing number of devices that are spread over an ever-widening geographic area.

Modern telecommunication networks transport Internet Protocol (IP) traffic utilizing primarily hub & spoke topologies. Here low-speed transceivers (spokes at the end user's location) are connected to high-speed transceivers (hubs). The optimal network topology for this would be point-to-multipoint (P2MP) and, usually, the physical plant is built as a P2MP architecture. However, the traffic is usually transported via point-to-point (P2P) architectures, such as those typically deployed for traditional telephony services, with the exception of PONs, where P2MP is realized using time-division multiplexing (TDM).

These network segments are the closest to the end users, and because of the large number of devices, cost and power consumption play a crucial role. At the moment, the preferred transmission method in access and metro networks is Intensity-Modulated Direct-Detection (IM-DD) systems, which employ modulation formats such as non-return-to-zero (NRZ) or pulse-amplitude-modulation (PAM)-4.

This technical solution has significantly contributed to the boom of the Internet, but initial doubts about its long-term sustainability – in terms of capacity – are arising as next-generation optical networks are being called upon to support even larger amounts of data. In fact, IM-DD is limited in spectral efficiency compared to the advanced modulation formats employed in coherent systems. Another strong limitation of the technology used in the access network is the utilization of TDM architectures, which significantly reduce the maximum throughput of the network and does not allow the usage of advanced digital signal processing (DSP) algorithms. These boundaries, and the underlying P2P network architecture, have resulted in a growing awareness in the industry that P2P 10G IM-DD systems cannot meet the requirements for backhaul while simultaneously making efficient use of network infrastructure, e.g., in the case of fiber scarcity.

In this context, coherent optics, as has happened in other segments, might come to the rescue where capacity and reach become the core concerns, but they still fall short where high device counts or simply geographic distribution are the key issues. A first important step towards a low-cost coherent marketplace is represented by the Optical Internetworking Forum (OIF) implementation agreement of 400ZR [OIFIA2020]. Coherent offers a wide array of advantages ranging from the ability to employ advanced DSP technologies – that can compensate for fiber propagation effects such as accumulated dispersion – to the enabling of wavelength division multiplexing (WDM).

This evolving storyline requires new solutions to address the given requirements. This paper provides operators with an evaluation of possible approaches to this issue ranging from higher speed IM-DD systems to coherent optics with a special focus on the emerging concept of coherent P2MP transport using digital subcarrier multiplexing (DSCM) [Sun2020] for aggregation [Welch2021].

Fiber access links have traditionally been arranged in a P2P fashion, with paired electronics on each end of the fiber. If electronics are changed on one end, they must be changed on both ends to remain interoperable. Consequently, to combine low-speed interfaces to higher-speed ones, an intermediary aggregation device is required. In [Welch2021], a paradigm shift has been proposed that targets a universal approach to connectivity in telecommunication systems. This exploits DSCM to simplify the network architecture by removing the electrical aggregation layer and the bookended transceivers and replacing them with a simple passive optical combiner. Moreover, DSCM allows the realization of a truly P2MP network architecture [Welch2021].

This is relevant because most current-day last-mile architectures are broadcast in nature – meaning that transmitted downstream signals are “heard” by all endpoints. This is true for PON, DOCSIS, and Mobile networks. Wired networks (PON/DOCSIS) might be single or dual fiber. Whether one considers ring topologies, star topologies, passive tree architectures, PON overlay architectures, or enterprise/wavelength extensions, the deployment approaches are all similar. “Last mile” networks are predominantly P2MP in nature, but still served by P2P optics. Such designs tend to drive high optical port counts, while the solution based on DSCM might lead to significant savings [Bäck2020].

At the same time, that network capacity continues to grow, marketed network speeds are also growing, with 1 gigabit per second (Gbps) service now commonly available, and plenty of industry discussion around emerging multi-Gbps speeds on the horizon. We see this from a few perspectives, including residential broadband rates, demand from business customers, wireless capabilities now made possible via 5G, as well as an accelerating build rate of fiber networks. In addition, the cable industry has put forward its strong 10G mantra over the past few years. The discussion will address critical operator needs such as capacity, scalability, deployment lifespan, network reach, CIN aggregation architecture impacts, the convergence of DAA/wireless/business services, and operational simplification. A road map toward a new network paradigm will be also provided.

The remainder of this paper is organized as follows. Section 2 describes the state-of-the-art in terms of traffic, networks, integration with cable access technologies, and the reasons to move beyond 10G. Section 3 describes existing optical transmission technologies and how those can move beyond 10G. Section 4 presents a proposal that aims to cope with existing traffic patterns and high-capacity networks. Here, the concept of DSCM is introduced and we explain how this helps to realize fully capable point-to-multipoint optical networks. Section 5 provides a high-level qualitative comparison of technology options in moving beyond 10G before wrapping up.

2. Evolving Traffic and Deployment Requirements

In this section, we describe the evolving traffic and deployment requirements to serve various cable operator access solutions for HFC, PON, and wireless.

2.1. Network segments and traffic patterns

Figure 1 illustrates a generic tiered telecommunication infrastructure typical of systems deployed by cable operators today, spanning all the way from subscribers back to a wide area core network, with all links except for most subscribers, consisting of digital optical transport.

The access segment includes equipment such as Converged Cable Access Platform (CCAP) in DOCSIS hybrid-fiber coax (HFC), optical line terminal (OLT) in PON, and centralized unit (CU), distributed unit

(DU), radio unit (RU) in 5G wireless which provide last mile connections to users. Connections are generally tens of kilometers and limited to 80-120 km maximum.

The metro segment connects multiple access segments together across a larger, generally metropolitan area with links <200 km.

The core segment provides multi-city transport between different metropolitan areas covering hundreds or thousands of kms.

In this paper, we focus primarily on the access segment, but the solution described is also applicable to the metro segment.

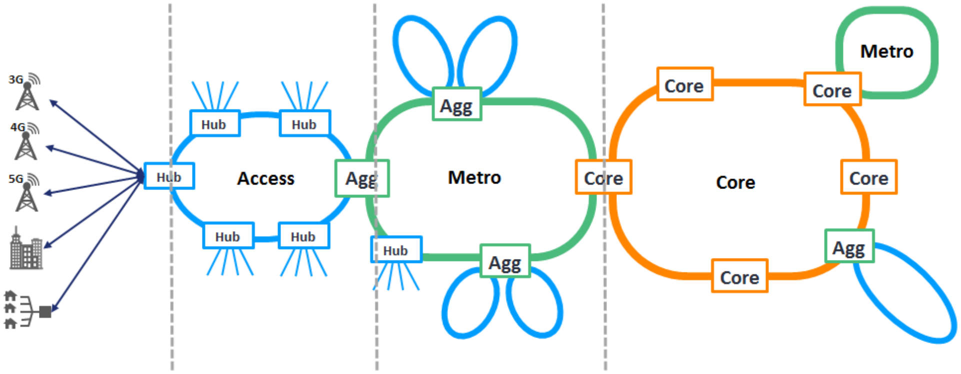

Figure 1 - Generic Cable Operator Network Architecture

A significant portion of the traffic within a cable operator's network is used for providing access from a large number of subscribers to a smaller number of service access points delivering services such as operator-delivered video content, streaming video providers, or general Internet services. This traffic is inherently organized in a hub and spoke pattern. Relatively low-speed (compared to optical transport) subscriber last mile interfaces connect to the Internet and other services via a high-speed hub, and traffic patterns resemble a point-to-multipoint (P2MP) network.

Today though, optical networks are still designed utilizing a point-to-point (P2P) approach, with the exception of access via passive optical network (PON). Figure 2 illustrates the differences between these traffic patterns.

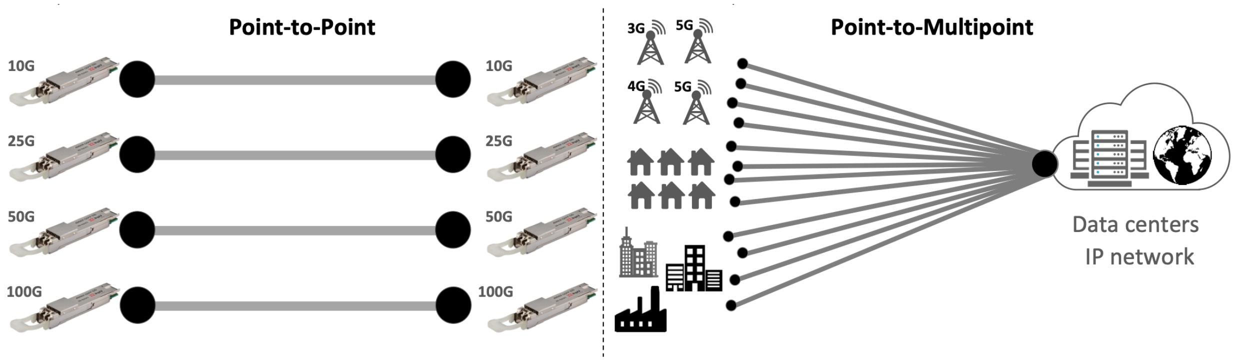

Figure 2 - Point-Point and Point-to-Multipoint Network Traffic Patterns

2.2. Evolving capacity needs for cable access segments

Over the last decade, the demand for capacity in the last mile access network has been growing continuously. At the same time, the access topology has also changed, with the "last mile" links becoming shorter, and active electronics being pushed deeper. Whereas fiber, coaxial, and twisted pair copper north bound endpoints previously terminated in facilities (headends, central offices, huts, etc…), it has become far more common for these termination points to be placed somewhere in the outside plant infrastructure.

Examples include:

- Cable operators' hybrid fiber coax networks have typically been fed by a combination of video QAMs (signal modulators) and Converged Cable Access Platform (CCAP) systems for broadband. These were typically rack mounted in an environmentally controlled facility. With the advent of DAA using Remote physical layer (PHY) or Remote media access control + physical layer (MACPHY) devices, those electronics take on a smaller more modular format that gets pushed out of the rack and into the fiber node installed in the neighborhood that is being served.

- Since the early 2000's, TelCo's have deployed remote digital subscriber line access multiplexers (DSLAMs) to deliver broadband over twisted pair copper.

- Mobile network operators (MNOs) have experienced a fast evolution of cellular radios and the radio access network (RAN) with electronics moving in multiple directions within the network topology. Increased use of virtualization and evolution of the location at which processing happens within the RAN has created increased need for a variety of optical transport solutions ultimately feeding radios located deep in the network.

- And for operators deploying fiber to businesses and residences, PON fiber terminations now frequently occur in remote OLTs that are pedestal or strand mounted, or within MDUs and office buildings.

The result is that the number of locations requiring fiber terminations is growing exponentially, while requiring higher and higher speed interfaces at the same time. Such network demands require a well thought out approach to connectivity.

Figure 3 below shows common scenarios for cable operators as they consider the required capacity in the access segment of their network.

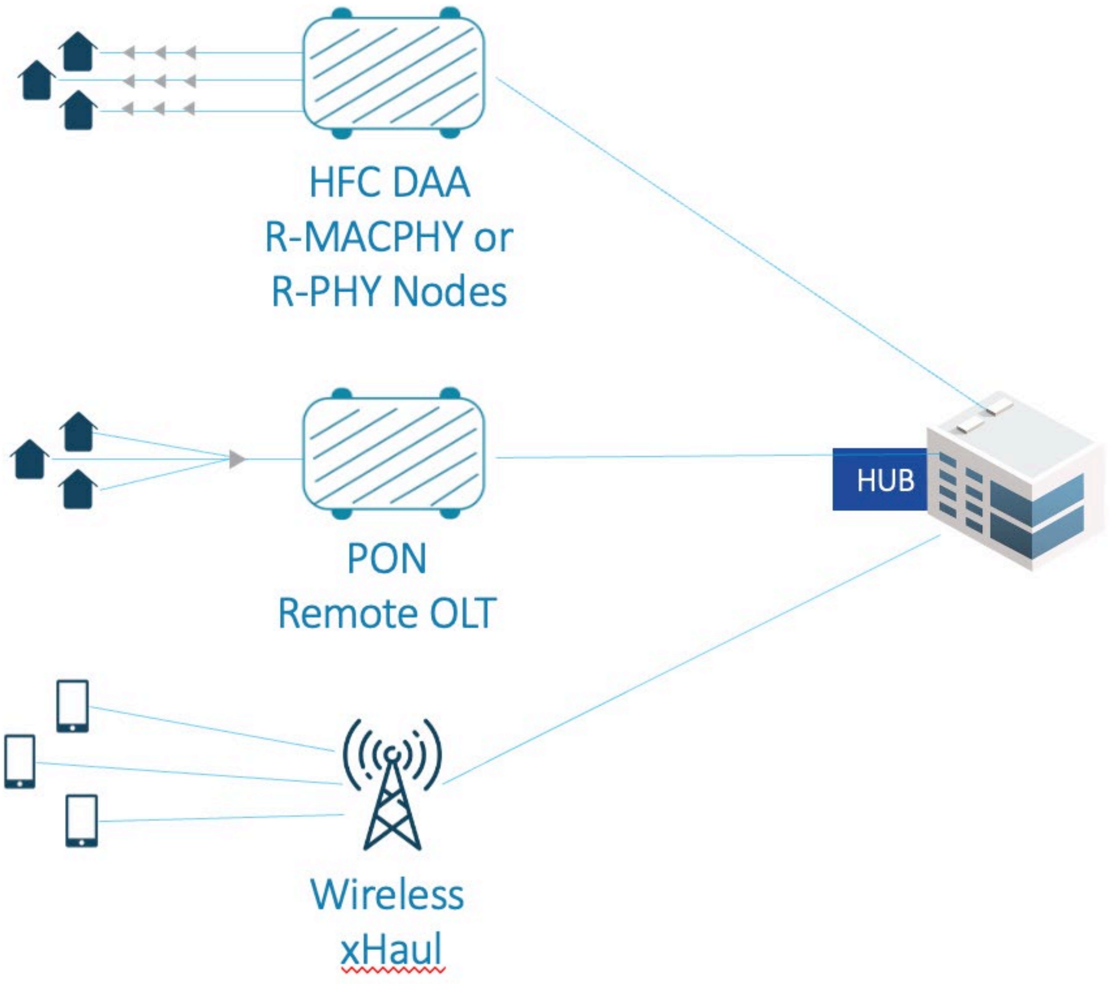

Figure 3 - Access Segment Transport Applications

2.2.1. HFC DAA

Many operators are evolving their DOCSIS and QAM video delivery infrastructure to DAA to:

- Save space and power in hub facilities as the number of HFC service groups expands with ongoing capacity growth
- Improve signal fidelity and resulting capacity by moving RF processing to the edge
- Converge HFC access backhaul with other forms of access backhaul by removing analog optics and replacing them with IP/Ethernet

DAA deployments today typically consist of 1 downstream (DS) service group with a mix of DOCSIS and QAM video. The total throughput required to feed a single DS service group with 1218 MHz of spectrum can theoretically exceed 11 Gbps but operators have generally coalesced around using 10G small form factor plus (SFP+) with 10GBASE-LR (long reach), ER (extended reach), or ZR/dense wavelength division multiplexing (DWDM) optics based on IM-DD as a pragmatic, cost-effective solution.

The expense of updating the nodes once deployed leads operators to build outside plant (OSP) and the converged interconnect network (CIN) feeding DAA nodes for long term (5-10 year) capacity. Capacity needs are soon expected to exceed 10 Gbps in DAA nodes due to:

- Increasing desire to deploy segment-able nodes which provide an option to support 2 (or more) DS service groups in a single DAA node at 10 Gbps each (20 Gbps for 2 DS SG)
- DOCSIS 4.0 (D4.0) frequency division duplex (FDD) extension to 1794 MHz can support DS capacity in an individual service group beyond 13 Gbps; this allows D4.0 to exceed the capacity of 10 Gbps PON and offer true 10 Gbps service which may be an important marketing difference in years to come

As a result, 25G links to the DAA nodes is becoming a desirable target for the combination of D4.0 and multiple-downstream service groups in next-generation DAA nodes.

HFC DAA nodes are also generally power constrained, with increasing pressure to fit more and more processing or other features in a housing (in North America) which is thermally limited to 160-180 Watts in harsh outdoor environments. In many DAA node designs, pluggable uplink Ethernet optics are generally the constraint for maximum operating temperature. Larger and more thermally friendly pluggable form factors such as C form-factor pluggable half-size (CFP2) are favored in this environment over small form factor pluggable (SFP) or quad small form factor pluggable (QSFP) formats.

The upcoming standardization of a generic node housing with the SCTE Generic Access Platform (GAP) standard will drive this further by allowing other modules for edge compute or Ethernet switching in the same node. Power efficiency and power consumption are critical in DAA applications.

2.2.2. Remote PON

With advancements in IP video delivery and the ability for operators to deploy video service without traditional QAMs, cable operators are increasingly deploying more fiber to the home (FTTH) in the network. This is especially true in greenfield builds such as housing subdivisions, and also in rural areas fueled by government broadband funding. Common technologies for deployments today are 10G-EPON or XGS-PON.

Due to the nature of the cable network with long distance links to hubs and the desire for high split ratios of 64:1 or 128:1, many operators cannot easily support hub-based OLT. As a result, the emerging dominant deployment model is the Remote OLT (R-OLT), containing 4 PON segments within a single outside plant housing powered from the coax cable.

In terms of capacity, each PON segment is generally limited to ~9 Gbps of maximum throughput due to the nature of the PON system overhead. An R-OLT is typically fed using the same 10G SFP+ that would be used for HFC DAA including 10GBASE-LR/ER/ZR optics as a pragmatic solution. A one-to-one arrangement of SFP+ to PON segment is common, but some operators optimize cost using R-OLT platforms that support aggregation of the uplink interfaces to multiple PON segments.

R-OLT implementations could already use a cost-effective solution beyond 10G to support aggregation over a single interface for fiber/wavelength savings. Year over year capacity growth and evolution in the R-OLT is also expected to drive further need for capacity beyond 10G due to:

- Support for 25G-PON (25GS-PON or 25G-EPON)
- Potential increases in PON segments beyond 4 per R-OLT
- In some countries, cabinet-based deployments with up to 16 PON segments in a single chassis are also common

R-OLT are not as thermally or space constrained as HFC DAA nodes due to the lack of RF amplifiers, but pluggable uplink optics are still a significant design consideration.

2.2.3. Wireless xHaul

Many cable operators operate wireless networks directly or sell services to provide transport for wireless operators. With convergence of the access segment of the network to serve HFC and PON, it's natural to also consider wireless xHaul within that same network.

Figure 4 illustrates the split options in various 5G deployment architectures which differ in where and how functions are positioned between elements containing some or all RAN components consisting of CU, DU, and RU. Cable operators may deploy a mix of these architectures depending on the siting and capacity needs within a particular market.

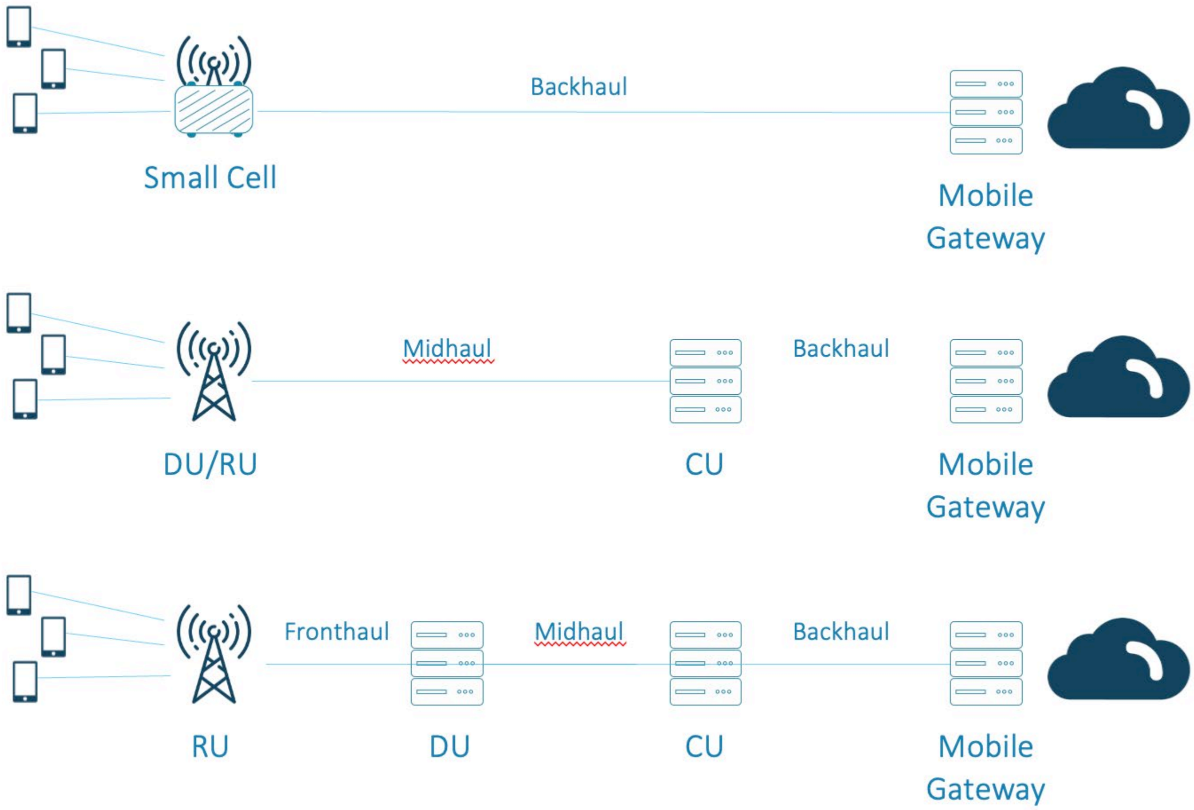

Figure 4 - 5G xHaul Transport Splits

Architectures feeding backhaul and common mid-haul options (F1 split, option 2) through the access segment will require overall capacity approximately equal to user bandwidth. Except in the case of massive multiple-input, multiple-output (MIMO) mm-wave bands, this capacity will generally not consume > 10 Gbps.

The increased use of virtualized RAN, with significant benefits in using elasticity to match capacity to where users move during the day, leads to growth in a fully disaggregated RAN with significant fronthaul usage possible through the access segment. While it is relatively common today to use dark fiber and gray optics with the need for fiber construction, an access segment which is naturally increasing capacity could allow for convergence and increased reuse of the existing fiber to boost wireless capacity.

Fronthaul bandwidth requirements, as indicated in [Cisco5G] can easily exceed 10G or even 25G as cell site capacity and the use of massive MIMO expands. An access segment architecture which supports endpoint capacities up to 100G could be well suited to serve these wireless needs.

Table 1 - 5G xHaul Transport Capacity Requirements

Band	Bandwidth	MIMO Layers	Fronthaul Data Rate for 3 Sectors No Compression	Midhaul/Backhaul Data Rate for 3 Sectors
850 MHz	10 MHz	4T/4R	7.35 Gbps	330 Mbps
1.8 GHz	20 MHz	4T/4R	14.7 Gbps	660 Mbps
3.5 GHz	100 MHz	64T/64R, 8 layers	57.87 Gbps	13.5 Gbps
28 GHz	400 MHz	64T/64R, 4 layers	82.32 Gbps	19.5 Gbps

2.3. Pushing beyond 10G

In the access space, service providers are faced with a challenging question – once an operator needs to upgrade or support capacities beyond 10G, should they upgrade to 25G or to 100G? If an operator upgrades to 25G, that increased capacity, as is often the case, will have a limited lifetime because there will be a need to upgrade soon again, assuming today’s 30% yearly traffic growth persists. However, deploying 100G would likely be cost prohibitive today.

A quandary operators face in the “last mile” access network is that the number of network touch points far eclipses those of all other parts of the network. This leads to a situation where the operational expenses encountered for equipment upgrades drive costs that are often vastly more than for the equipment itself. Operators are well served to think this through and work to deftly address these ever-growing needs, going beyond just a repetitive process of constantly adding capacity in an incremental fashion. Minimally, such changes should have a lifetime of at least five years, and more ideally upwards of ten years.

Many choices have a ripple effect beyond just the end point solution. Commonly end points require an aggregation router at a hub/headend location. Aggregation devices with SFP/SFP+ interfaces usually have two or more 100G interfaces and those are used to communicate with two or more spine switches. When a network operator is considering bandwidths above 10G, each leaf switch is only aggregating four to seven 25G optics, or in other words, the cost of that leaf switch is applied to only six 25G end points on average. The hub facility at the edge of the network may also be experiencing space and power constraints so density of connections is increasingly important, especially when considering operator intention to add new functions like mobile edge compute (MEC). If an operator chooses to upgrade to 100G optics, then they might only be using 15%-20% of the capacity, so again, an aggregation device is needed higher up to drive more efficient port utilization on routers.

In HFC networks, it is becoming increasingly observed that fiber termination points at the fiber node demarcation points (where fiber transitions to coax occurs), are becoming an increasingly strategic point in the network that can be leveraged for a multitude of services. Fiber is such an important resource, as it represents an asset that was constructed at great expense. Furthermore, it is also a medium which is incredibly flexible and expandable at modest cost. These fiber end points represent the "new network edge" for MSOs that can be used to serve not only coax-fed residential and SMB customers, but also enterprise customers, cell backhaul services, as well as an operator's own RAN.

3. Existing optical transmission technologies

Direct detection (DD) was one of the first solutions to detect optical signals at the receiver. It is a technique where only the amplitude information is preserved after the photodiode optical signals are converted into electrical currents. The first commercial optical systems were realized with this technique, and DD dominated all market segments until approximately 2010. Up to that time, optical systems could transport 10 Gbps per channel over a total of 80 wavelength division multiplexed (WDM) channels on a single fiber.

Since 2010, there have been rapid advancements in coherent optical transmission technology to extract more and more capacity from a single channel. As a result, transmission systems with 800 Gbps or higher capacity are now commercially available. Coherent optics, in a similar way to the classical concept of coherent radio transmission, uses polarization along with amplitude and phase modulation to realize higher order quadrature modulations and provide higher capacity from a specific line rate or channel bandwidth.

3.1. 10G

10 Gbps per channel is widely used today. 10 Gbps Ethernet, first standardized as IEEE 802.3ae in 2002, is used widely in enterprise networks and access network backhaul especially for DAA in DOCSIS and fixed wireless networks. 10 Gbps was standardized as an option for FTTH in 2009 as IEEE 802.3av (10G-EPON) and in 2010 as ITU-T G.987 and is now the fastest growing access network technology in the world. 10 Gbps isn't just for fiber, either – DOCSIS support for DS speed approaching or exceeding 10 Gbps is becoming possible with DOCSIS 4.0 technology.

3.2. 25G (and beyond)

IM-DD has continued to evolve. 40 Gbps was the next step, with IEEE 802.3bg in 2011 creating the standard for single-lane 40 Gbps Ethernet (4 × 10 Gbps was standardized in 2010), but it is limited to a distance of 2km and is all but defunct in the market.

25 Gbps, as a by-product of 100 Gbps (4 × 25 Gbps) was the next step beyond 10 Gbps that was attractive to the market. IEEE 802.3by and IEEE 802.3cc standardized 25 Gbps in a single lane, and 32G Fiber channel (really 28 Gbps) was introduced at the same time.

25 Gbps has taken on a life of its own, independent of 100 Gbps. Economically 25 Gbps fits a sweet spot in network designs – fulfilling a need for more bandwidth but at a much lower cost than a jump to 100 Gbps. 25 Gbps per channel is now being applied to passive optical networking in the 25/50G EPON standard, IEEE 802.3ca. The limits of today's IM-DD technology began to be exposed in development of 50G PON at the ITU-T where 50Gbps over a single channel in the downstream was adopted for the ITU-T G.9804 series. Nonetheless, specification of 50 Gbps in the upstream was postponed in favor of 10 Gbps and 25 Gbps in the upstream due to the difficulties of burst reception at 50 Gbps.

Achieving 100 Gbps on a single channel with IM-DD was no small feat. In 2021, IEEE 802.3cu was able to achieve a standardized 100 Gbps on a single channel at a distance of 10 km. There are proprietary versions of single-lane 100Gbps that are able to achieve up to 40 km, but greater distances are relying on advanced modulations.

4. Future coherent transmission technologies

A race is on to establish the next-generation technology that will replace current P2P 10G optics based on IM-DD to support increasing capacity requirements for access and metro to serve DOCSIS, PON, and 5G wireless services. Among the most relevant challenges, one could single out bandwidth demands, compensation of the physical layer impairments, and continuous and dynamic network upgrades. Next, we will describe coherent systems in the form of single-carrier systems as well as DSCM.

4.1. Single-carrier coherent systems

Coherent detection has two key advantages over IM-DD. The first is that coherent detection preserves the phase information, and DSP algorithms can be applied within the transceiver to decode information embedded in the phase. Secondly, the local oscillator allows excellent channel selection. The first property was not fully exploitable until complementary metal–oxide–semiconductor (CMOS) and digital-to-analog (DAC) and analog-to-digital converter (ADC) technologies became mature. The second became less important after the introduction of the first commercial optical amplifiers. Because of this, coherent was not used during the first optical revolution that started with WDM systems and IM-DD. After the dot-com era and explosion of IP data traffic, the throughput of IM-DD systems became insufficient to cope with the exponential bandwidth growth.

This led to the rapid development of 100G thanks also to the implementation of the DSP algorithms and advanced forward error correction (FEC) code within the first generation of application-specific integrated circuit (ASIC) [Sun2008, Roberts2009]. These achievements led to the spreading of coherent technology from subsea to core and regional networks, where it replaced the IM-DD. Now, coherent transponders can achieve data rates beyond 1 Tb/s [Sun2020, Buchali2016], employing high-order modulation formats enabling a multitude of modes by varying the symbol rate, FEC overhead, and probabilistic constellation shaping (PCS). Although coherent has been so far confined to core and submarine networks, thanks to recent initiatives such as 400ZR, it is set to enter new markets such as data center interconnects.

IEEE 802.3ct has released a specification for 100 Gbps over a single channel for 80 km using dual polarization-differential quadrature phase shift keying (DP-DQPSK) (not direct detection). CableLabs® released a coherent optics specification in 2019 that describes the architecture and requirements for a 100 Gbps and 200 Gbps capable of up transmitting data up to 120km. As of today, though, only a handful of vendors are offering products that meet the 100Gbps requirements and fewer are offering a 200 Gbps option. The CableLabs coherent optics specification requires only a single fiber for transmission which is a big advantage over most other options that require 2 or more fibers. One potential disadvantage of coherent optical modules is that they incorporate a DSP into the housing which requires more power and a larger format due to the component size and required heat dissipation.

Type	Launch Power (dBm)	Loss Budget (dB)	Distance (km)	Speed (Gbps)	Fiber Count
10GBASE-ZR	0	25	80	10	2

Type	Launch Power (dBm)	Loss Budget (dB)	Distance (km)	Speed (Gbps)	Fiber Count
25GBASE-ER	-1	20	40	25	2
25GBASE-LR	-5	18	10	25	2
10GBASE-ER	-1	14	40	10	2
10GBASE-LR	-11	8	10	10	2
50GBASE-LR	-4.5	12	10	50	2
50GBASE-ER	-0.5	15	40	50	2
100GBASE-FR1	-2	4	2	100	2
100GBASE-LR1	-1	8	10	100	2
CableLabs P2PCO 100G	-6	21	80	100	1
CableLabs P2PCO 200G	-7.5	19	80	200	1
10GBASE-PR	2	30	20	10	1

Table 2 - Comparison of Single Channel Optical Systems

Coherent is also gaining momentum in the PON space with CableLabs recent kickoff of Coherent PON which is targeting 100Gbps PON using coherent optics as the underlying physical layer.

4.2. Digital subcarrier multiplexing based coherent systems

An alternative and more flexible solution to standard coherent is digital subcarrier multiplexing (DSCM), which creates subcarriers (SC) in the digital domain using only one laser. From a networking perspective, it is similar to the sliceable bandwidth variable transceiver (S-BVT) introduced by [Sambo2015]. The key difference between the single carrier of Section 4.1 and DSCM is that now the DSP operates at a symbol rate (R_s) of R_s/N, where N is the number of digital SCs. Clearly, in this case. the transceiver performance might be limited by the value of the laser linewidth Δf [Dris2013]. In fact, given the larger symbol period T_s, the $\Delta f \cdot T_s$ product for a DSCM signal with N SCs is N times larger than that of a single wavelength [Welch2021].

DSCM is used in P2P links [Sun2020], and it provides different benefits. First, it enables, by parallelizing with respect to the individual SCs, simplifcation of the DSP algorithms. For example, the compensation of accumulated dispersion. Second, it significantly reduces the impact of equalization enhanced phase noise (EEPN), which is relevant in the case of long-haul transmission. Third, by optimizing the modulation format of the SCs, it helps to increase the tolerance against filter cascade [Rahman2016]. Last, thanks to its intrinsic flexibility, it enables transmission to be adapted to the current traffic, thus reducing the power consumption [Velasco2021].

A pictorial representation of the spectrum of a 400G 16QAM DSCM signal – with 16 SCs at 4 gigabaud (GBd) each, and its equivalent single-carrier are depicted in Figure 5, where both channels occupy the same 64 GHz bandwidth. Only in case of ideal Nyquist shaping (i.e., roll-off = 0) there is no need for a guard band as there is no overlap. Roll-off = 0 is not possible, and therefore a guard band is needed [Welch2021].

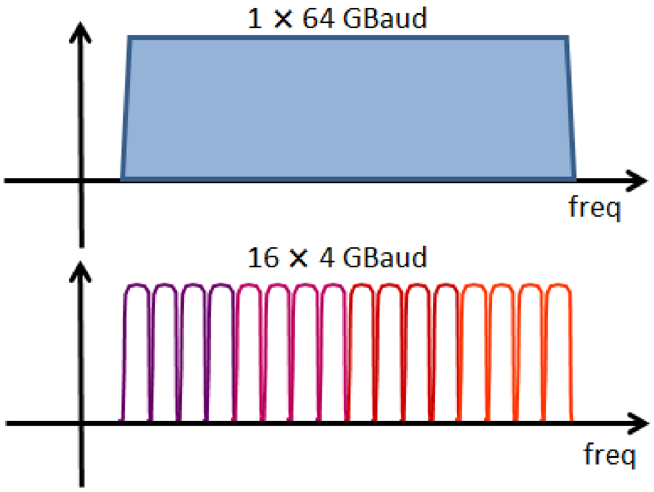

Figure 5 - Single Carrier and DSCM Spectrum

4.3. Enabling point-to-multipoint optical aggregation networks with DSCM

Beside the benefits described in Section 4.2, DSCM and coherent can enable the re-architecting and simplification of future optical networks. Figure 6 illustrates a simplified metro/access aggregation network realized with P2P transceivers (a) and with P2MP ones enabled by DSCM (b).

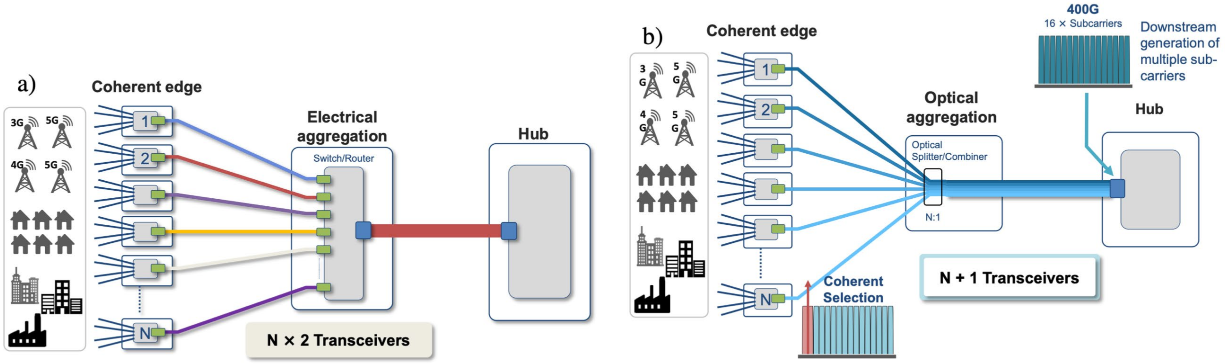

Figure 6 - Point-to-Point vs. Point-to-Multipoint Aggregation

In Figure 6-a, the N end users are equipped with N low-speed transceivers connected with their N pairs at the electrical aggregation stage. Next, a high-speed transceiver communicates to the hub, for a total of 2N (low-speed) + 2 (high-speed) transceivers. The architecture shown in Figure 6-a can be realized by employing IP routers for aggregation and multiplexing of the low-speed into the high-speed transceivers, where each link uses the optimal rate P2P transceivers.

This approach is sub-optimal for hub and spoke traffic and is a major drawback of this solution. As optical spectrum in fiber is not scarce outside core networks, to achieve the lowest cost solution, we minimize the number of regeneration points. This approach hides a waste as there is no value to the end-users in regenerating traffic unnecessarily. In fact, this approach is constrained by decisions that are based on the peak traffic value. Consequently, operators will decide (1) to minimize the number of physical site visits, thus deploying higher-capacity transceivers – i.e., higher CAPEX or (2) to optimize the capacity of the transceivers; thus, increasing the number of truck-rolls, i.e., higher OPEX. Neither solution is economically optimal. With this architecture if an endpoint requires more capacity, the two transceivers in that link must be replaced. In the situation where the electrical aggregation devices utilize common port capacities, as is typical in most commercial high density devices, the entire aggregation device and all attached transceivers may need to be replaced. A typical example of this could include a DAA CIN built to support 10G with high density 10G switches for electrical aggregation; a need to change a link to the next step higher (25G) cannot be easily accommodated in the existing aggregation device.

Figure 6-b shows the same aggregation network realized with DSCM. The electrical aggregation stage is now replaced by a simple 1:N passive optical combiner, which greatly reduces the number of devices and stages needed to aggregate and up-speed the traffic to be transported to the next hub. In Figure 6-b, only N low-speed transceivers and 1 high-speed are required, i.e., 50% less than in Figure 6-a. This P2MP architecture connects multiple low-speed transceivers (spokes) to high-speed ones (hubs), breaking the bookended transceiver paradigm. These benefits directly translate into lower power consumption, smaller footprint, less sparing of parts and grooming equipment.

Furthermore, P2MP with DSCM enables cross-layer savings related to an efficient utilization of the optical transceivers, which can flexibly provide the required amount of capacity. For example, an operator deploys 100G (4×25G) P2MP pluggable optics in 12 leaf nodes. On day 1, only 1 SC per node is active, and all 12 SCs are terminated in a single 400G hub router port, which can receive up to 16 SC at 25G each. This is equivalent to approximately one-half (13 vs. 24) of the devices compared to what would be required with P2P optics on day 1. If the capacity grows in the edge nodes, then more SCs will be needed over time, and they can be activated via software. This operation is possible if the definition of SCs is maintained over more generations of P2MP optics, so that different versions of pluggable can interoperate, and operators can thus define multi-generational network architectures. By enabling this, routers and hubs can be independently upgraded, thus decoupling nodal upgrades from network-wide ones. Network operators can maximize the return on investment and ensure a smooth and cost-effective capacity upgrade of the network.

Nevertheless, there are challenges for P2MP architectures. First, it requires time to remove routers from the network, as these elements are rich in features, and operators use them for several network functions. Second, having fewer transceivers leads to network simplification, but if one device fails, the amount of traffic that is affected will be larger. A third challenge is the transmission of SCs over different light paths. If different SCs are propagated over different links, at the hub, they might suffer power imbalance. The pluggable will have an adjustable range to transmit and receive, and this will be monitored via a power control loop between the hub and each edge transceiver. Nevertheless, if the path loss difference exceeds the transmit output adjustment range, an external attenuator might be required.

As bandwidths increase and the number of devices in the geographic area expands, there may be a desire for protection or redundancy. The access network is typically unprotected over the last mile, but redundant optics can be used at the hub location to protect against module failure; where there is an active/standby architecture deployed. If diverse fiber paths are available, then current protection mechanisms can be used to protect against fiber cuts. The cost of redundant equipment at the access edge of the network is typically prohibitive but for those applications that can justify the incremental expense that is also an option.

Overall, coherent with DSCM is a combination of mature technologies and it has been successful in high-end products, but when it comes to mass production, the required components, e.g., the laser, might have higher costs relative to traditional IM-DD or single carrier coherent solutions. On the other hand, coherent with DSCM significantly reduces the number of devices, eliminates the electrical aggregation stage and reduces the number of truck-rolls. It has been shown that, if the coherent pluggable price is below 50% of the existing pluggable, the solution presented in [Welch2021] brings cost benefits, when we consider the entire network [Marino2021].

4.3.1. Time-division versus frequency-division multiplexing

Time division multiplexing (TDM) and frequency division multiplexing (FDM) serve different scenarios with their advantages and disadvantages. In access, TDM – and its variants, e.g., statistical TDM – is used in PON networks before aggregation, where the multiplexing is carried out in time. This is a complex operation in the case of high data rates.

FDM – with DSCM – has been deployed in core networks [Sun2020] and enables Tb/s channels after aggregation via high-speed coherent transponders. In transport, the FDM-based channels are seen as one carrier by the remaining elements of the network. Hereafter, we compare TDM versus FDM, based on network application scenarios.

In PONs thanks to the application of statistical TDM, end users can receive streams down to Mb/s speed. This differs from the proposed coherent solution with DSCM approach, as the laser instability does not allow it to operate at symbol rate well below 1 GBd. From a transmission point of view, PONs today use the highly cost-effective IM-DD, whose transceivers are manufactured in volumes exceeding millions of units per year [Nesset2017].

On the other hand, in the case of data rate upgrade in a PON network based solely on TDM, all devices (hub + leaf modules) need to be replaced because the same peak communication rate must be supported. With FDM, the technology used and power consumption applicable to lower data rate locations can be optimized for that lower data rate. Transmission and reception at the peak rate is not needed. Maintaining high peak rates in all locations may increase costs and integration challenges for those optoelectronics in applications that do not need the full channel rate. .

Next, PON TDM transceivers cannot exploit the capability of advanced DSP techniques. DSP algorithms have difficulties in the learning phase required for burst mode transmission that is used in today's PON. For this reason, data rate upgrade might be limited for what concerns reach and synchronization. The last is particularly critical as the leaf modules will be placed at different locations before the aggregation stage. Because of these limitations, PON standards advance at a slower pace compared to other technologies.

In opposition to PON based solely on TDM, FDM simplifies data rate upgrades and enhances the network flexibility. By using DSP and continuous mode transmission/reception in both directions, it can extend the reach, and it supports DWDM. An FDM network enables different costs for users versus hub, 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 of 400G. P2MP optics based on coherent transmission with DSCM offers a roadmap to higher end user capacities than IM-DD implementations can achieve, and as traffic demands in access networks grow, the proposed implementation provides a number of other advantages, as discussed above.

5. Beyond 10G Solution Comparison

A flexible access segment optical network topology should provide:

- Cost effective technology
- Scalability as well as ease of upgrading
- Sufficient distance reach
- Integration into outside plant node platforms
- Service convergence to allow for a unified access network segment
- Efficient use of hub space and power
- Long deployment lifespan
- Operational simplicity

Table 2 provides the authors' view of a summary comparison of the available technologies with these criteria in mind . A "+" score indicates a technology which excel in a particular area while "–" indicates the poor ability to support the desired criteria. For single-carrier coherent, some criteria are greatly impacted by the use of electrical aggregation (EA) and the "with EA" variation covers the case where electrical aggregation in the outside plant is added.

Table 3 - Beyond 10G Solution Comparison

Criteria	IM-DD 25G	Single-Carrier Coherent 100G/200G+	DSCM 25G-400G
Aggregate Module Capex	+	-- - with EA	-
Scalability	-	Neutral + with EA	++
Reach	-	+	+
Node Integration	+	Neutral	Neutral
Service Convergence	Neutral	+	+
Hub Space/Power	Neutral	Neutral	+
HFC DAA Suitability	+	- Neutral with EA	Neutral
PON R-OLT Suitability	Neutral	Neutral	+
Wireless xHaul Suitability	Neutral	Neutral	+
Deployment Lifespan	-	+	+
Operational Simplification	Neutral	-	+

Aggregate Module Capex

- 25G IM-DD solutions provide very reasonable module costs while coherent solutions are expected to be significantly higher than IM-DD for some time

- DSCM allows for a significant reduction in the number of transceivers due to the P2MP operation and optical splitting/combining

Scalability

- The selective use of subcarriers in DSCM allows for a single module to support a range of total bandwidths while the capacity is fixed in IM-DD; this could allow for specific variants of the modules which only serve small numbers of subcarriers

- Single carrier coherent can scale in capacity through modulation, but does not generally support scaling below 100G per link without the use of electrical aggregation

Reach

- IM-DD solutions in the 25G and above range (and especially in DWDM configurations) are not readily available beyond 40km while coherent solutions can provide reach out to 120 km in point-point operation. P2MP operation with DSCM can be achieved beyond 40 km at rates up to 100G with split ratios of 32:1 or less

Node Integration

- IM-DD 25G solutions using SFP28 are nearly identical in space and thermal footprint compared to 10G SFP+ used for DAA and R-OLT nodes today

- The extra power in the optics and DSPs in coherent solutions drive increased power beyond 7W and require larger form factors for node integration such as CFP2 to support effective passive cooling

Service Convergence

- Moving past 10G generally allows service convergence but the maximum limit may prevent higher rate wireless fronthaul or business services from being offered over the same access segment optical network

- Coherent solutions allow for greater data rates to more easily aggregate multiple services and support higher xhaul capacity needs

Hub Space/Power

- The additional power consumed by single-carrier coherent solutions is offset by the reduction in total switch ports if aggregation of multiple services is possible externally (outdoor coherent termination device for example)

- DSCM provides significant benefits for hub space and power due to the P2MP operation which allows for aggregation of multiple end points into a single high-rate hub module along with simple optical passive combining/splitting

HFC DAA Suitability

- Single-carrier coherent, because it cannot scale down without separate electrical aggregation, is not well suited to the 20-25G requirements of segment-able D3.1 and D4.0 DAA nodes

- DSCM can scale down to a single subcarrier to suit 25G needs but does come with tradeoffs on power and space in HFC DAA nodes already constrained in those areas

PON R-OLT Suitability

- 25G IM-DD supports 2:1 aggregation relative to 10G PON technologies (peak capacity design targets)
- DSCM scales well through the range of possible uplinks in a typical 4 port R-OLT (40-100 Gbps), even considering some oversubscription while single-carrier coherent overserves capacity in today's 10G PON environments

Wireless xHaul Suitability

- DSCM supports lower scale than single-carrier and the dynamic P2MP operation is very well suited to the elasticity and number of radio endpoints that are served

Deployment Lifespan

- IM-DD will eventually bump into limits as 25G is no longer sufficient for transport to last mile access devices
- Coherent solutions provide significant capacity to cover near and long-term requirements

Operational Simplification

- IM-DD 25G operations are just like DAA and R-OLT today
- Single carrier coherent will generally require a coherent termination device to fit the wide range of devices driving up transceiver count and adding another layer of electrical aggregation devices into the network
- DSCM use of P2MP operation ideally suits the needs of the network and allows passive optical splitting/combining instead of electrical aggregation

6. Conclusions

This article presents a comprehensive discussion on why access segment optical transport Beyond 10G is required while reviewing the solutions available using IM-DD and coherent transmission technologies. A new concept of coherent digital subcarrier multiplexing (DSCM) provides benefits in scalability using individual subcarriers while supporting P2MP operation to allow for transceiver reduction and replacement of electrical aggregation from single-carrier coherent with passive optical splitting and combining.

The near future is sure to be very active in this area with interoperability and standardization happening on multiple levels including the newly created OpenXR Forum and just launched CableLabs 100G Coherent PON projects.

Abbreviations

5G	Fifth Generation
10G	10 gigabits per second
25G	25 gigabits per second
ADC	Analog-to-Digital Converter
ASIC	Application-Specific Integrated Circuit
CCAP	Converged Cable Access Platform
CFP2	C Form-Factor Pluggable Half-Size
CIN	Converged Interconnect Network
CMOS	Complementary Metal-Oxide Semiconductor
CU	Centralized Unit
D4.0	DOCSIS Version 4.0
DAA	Distributed Access Architecture
DAC	Digital-to-Analog Converter
DP-DQPSK	Dual Polarization Differential Quadrature Phase Shift Keying
DSCM	Digital Subcarrier Multiplexing
DSLAM	Digital Subscriber Line Access Multiplexer
DSP	Digital Signal Processing
DU	Distributed Unit
DWDM	Dense Wavelength Division Multiplexing
EA	Electrical Aggregation
EEPN	Equalization Enhanced Phase Noise
FDM	Frequency Division Multiplexing
FEC	Forward Error Correction
FTTH	Fiber To The Home
GAP	Generic Access Platform
GBd	Gigabaud
Gbps	Gigabits per second
HFC	Hybrid-Fiber Coax
IM-DD	Intensity-Modulated Direct-Detection
IP	Internet Protocol
MACPHY	Media Access Control + Physical Layer
MEC	Mobile Edge Compute
MIMO	Multiple-Input, Multiple-Output
NRZ	Non-Return-to-Zero
OIF	Optical Internetworking Forum
OLT	Optical Line Terminal
P2MP	Point-to-multipoint
P2P	Point-to-point
PAM	Pulse-Amplitude-Modulation
PCS	Probabilistic Constellation Shaping
PON	Passive Optical Network
QSFP	Quad Small Form Factor Pluggable
R_s	Symbol Rate
R-OLT	Remote OLT
RAN	Radio Access Network
RU	Radio Unit

S-BVT	Sliceable Bandwidth Variable Transceiver
SC	Subcarrier
SFP+	Small Form Factor Plus
T_s	Symbol Period
TDM	Time-Division Multiplexing
WDM	Wavelength Division Multiplexing

Bibliography & References

[OIFIA2020] OIF-400ZR-01.0: *Implementation Agreement 400ZR*; Optical Internetworking Forum (OIF); https://www.oiforum.com/wp-content/uploads/OIF-400ZR-01.0_reduced2.pdf

[Sun2020] H. Sun *et al*., "800G DSP ASIC Design Using Probabilistic Shaping and Digital Sub-Carrier Multiplexing," in *Journal of Lightwave Technology*, vol. 38, no. 17, pp. 4744-4756, 1 Sept.1, 2020, doi: 10.1109/JLT.2020.2996188.

[Nesset2017] D. Nesset, “PON roadmap,” *J. Opt. Commun. Netw.*, vol. 9, no. 1, pp. A71–A76, Jan 2017.

[Welch2021] D. F. Welch *et al*., "Point-to-Multipoint Optical Networks Using Coherent Digital Subcarriers," in *Journal of Lightwave Technology*, doi: 10.1109/JLT.2021.3097163.

[Bäck2020] J. Bäck *et al*., "CAPEX Savings Enabled by Point-to-Multipoint Coherent Pluggable Optics Using Digital Subcarrier Multiplexing in Metro Aggregation Networks," *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1-4, doi: 10.1109/ECOC48923.2020.9333233.

[Sun2008] H. Sun *et al*., “Real-time measurements of a 40 Gb/s coherent system,” *Optics Express*, vol. 16, no. 2, pp. 873–879, 2008.

[Roberts2009] K. Roberts *et al*., “Performance of dual-polarization QPSK for optical transport systems,” *Journal of lightwave technology*, vol. 27, no. 16, pp. 3546–3559, 2009.

[Buchali2016] F.Buchali *et al*., “Rate adaptation and reach increase by probabilistically shaped 64- QAM: An experimental demonstration,” *Journal of Lightwave Technol- ogy*, vol. 34, no. 7, pp. 1599–1609, 2016.

[Sambo2015] N. Sambo *et al*., “Next generation sliceable bandwidth variable transponders,” *IEEE Communications Magazine*, vol. 53, no. 2, pp. 163–171, 2015.

[Dris2013] S. Dris *et al*., “M-QAM carrier phase recovery using the Viterbi-Viterbi monomial-based and maximum likelihood estimators,” in *2013 OFC*. IEEE, 2013, pp. 1–3.

[Rahman2016] T. Rahman *et al*., “Digital subcarrier multiplexed hybrid QAM for data-rate flexibility and roadm filtering tolerance,” in *Optical Fiber Communication Conference*, 2016, pp. Tu3K–5.

[Velasco2021] L. Velasco *et al*., “Autonomous and energy efficient lightpath operation based on digital subcarrier multiplexing,” *IEEE Journal on Selected Areas in Communications*, pp. 1–1, 2021

[Marino2021] P. P. Marino *et al*., "Point-to-Multipoint Coherent Optics for Re-thinking the Optical Transport: Case Study in 5G Optical Metro Networks," *2021 International Conference on Optical Network Design and Modeling (ONDM)*, 2021, pp. 1-4, doi: 10.23919/ONDM51796.2021.9492393

[Cisco5G] Bringing it All Together in 5G, S. Ajmeri, BRKSPG-3477, https://www.ciscolive.com/c/dam/r/ciscolive/apjc/docs/2020/pdf/BRKSPG-3477.pdf, 2020