Network Programmability – A Reality Check And A Glimpse Into The Future

A Technical Paper Prepared for SCTE•ISBE by

Fady Masoud, M. Eng.
Principal, Product and Technology Marketing
Infinera
555 Legget Drive, Suite 222, Tower B, Ottawa, ON, Canada K2K 2X3
fmasoud@infinera.com

Table of Contents

Title **Page Number**

Table of Contents 2

Introduction 3

Content 3
1. Why Network Programmability? 3
2. The State of Network Programmability 3
   - Programmable Client Interfaces 3
   - Programmable Line Interfaces 4
   - Programmable Optical Carriers/Super-channels 5
   - Programmable Service Activation 6
   - Programmable and Virtualized Infrastructure 7
   - Programmable Telemetry and Network Pulse 7
   - Programmability of Service Restoration 8
   - Programmable Inter-layer Service Setup 8
3. Network Programmability – A Glimpse into the Future 9
   - Programmable and Highly Granular Optical Transmission 9
   - Analytics and Machine Learning-triggered Network Capacity 9
   - Proactive Network and Traffic Protection 9

Conclusion 9

Abbreviations 10

Bibliography & References 10

List of Figures

Title **Page Number**

Figure 1 - Programmable Client Interfaces 4
Figure 2 - Programmable Line Interfaces 5
Figure 3 - Programmable Super-Channels 5
Figure 4 - Programmable Capacity through Sliceable Optics 6
Figure 5 - Dynamically Add, Modify, Move and Retire Optical Capacity with SDC 7
Figure 6 - Programmability of Service Restoration 8

Introduction

The fast-paced migration to the cloud, forthcoming 5G deployments and the proliferation of internet-connected devices (IoT) are fueling the evolution to a new era of intelligent networking that uses advanced analytics and machine learning to build self-optimizing, self-healing and highly autonomous transport networks. Moreover, emerging technologies like Blockchain that enable peer-to-peer distributed ledger for open and secure exchanges over the internet are, more than ever, putting the network at the heart of this resolution. Network programmability is one of the key building blocks for such successful evolution. However, its real-life implementation across the various moving parts of the network (network elements, layers, management platforms, data collection tools, etc.) often comes with technological challenges. This paper focuses on the “state of network programmability” today by highlighting where network programmability has been successful and how it translates into business and operational benefits. It also provides a glimpse into the future, including what the industry is currently working on to extend programmability further into the network and the necessary building blocks to do so.

Content

1. Why Network Programmability?

The proliferation of streamed content over the internet, including the IoT with billions of connected devices, new data exchange and validation models like Blockchain, and accelerated enterprise migration to cloud applications, has dictated a new level of task automation and programmability in order to keep up with the constant and unpredictable demand for bandwidth, streamline operations and eliminate sources of human error. In this new era of hyper-connectivity, where content must be delivered to hundreds of millions of users across the globe with the highest levels of quality, or machine-to-machine communications must be automatically and autonomously initiated without human intervention, network programmability becomes paramount.

The underlying network must be highly programmable and span the various constituent parts of the network to automate tasks, optimize network resources (e.g. optical spectrum, capacity/reach ratio) and make real-time network decisions.

2. The State of Network Programmability

Programmability has been making its mark across various and key parts of network components, layers and functions, from service delivery (client interfaces) to core network traffic management. The following highlight where it has been successfully implemented as well as the value it provides:

- **Programmable Client Interfaces**: The early days of transponders used to be service-and protocol-specific; for example, one transponder would be designed to deliver a predetermined and fixed client service such as OC-192, another type of transponder would be for OC-48 and so on. This often led to the quick exhaustion of a network element’s available service slots and required service providers to maintain a large inventory of circuit packs for each service. Transponders are now programmable, so each port can be set, by software, to operate over a specific client service type or protocol and cover a wide range of bit rates. All different ports of a programmable transponder can operate over any supported service type without restriction, allowing service providers to reduce footprint, maximize return on investment and decrease sparing costs. as depicted in Figure 1.

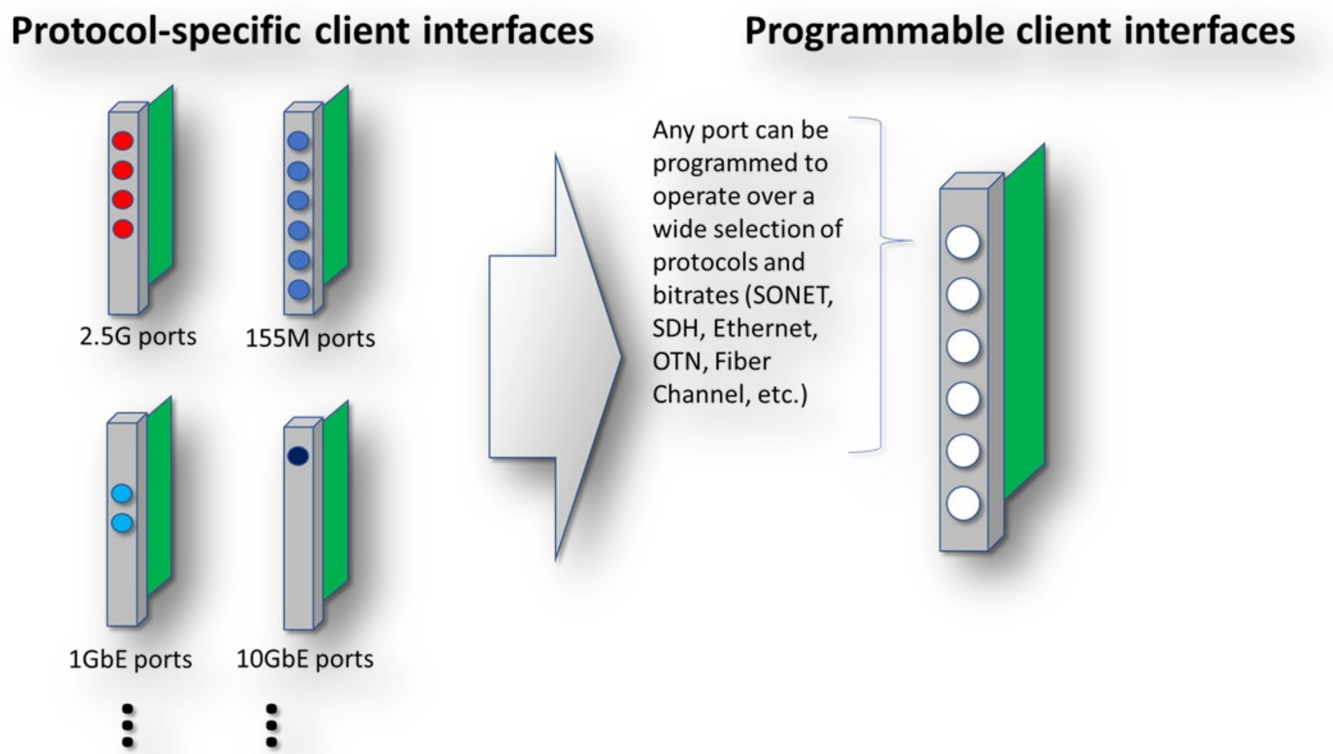

Figure 1 - Programmable Client Interfaces

- **Programmable Line Interfaces**: Programmability has also changed how the network is designed, operated and scaled. Early optical line interfaces operated over a fixed frequency, then later they became tunable across the C-band, where a single line card could be used to transmit over any of the 88 C-band fixed-grid frequencies. The latest innovation in software- and hardware-enabled line interfaces is full programmability: not only can they tune over any of the fixed C-band frequencies, but they also operate over any specific modulation schemes (e.g. QPSK, 8QAM, 16QAM) or a specific baud rate (e.g. 17 GBaud, 22 Gbaud, 33 GBaud), as shown in Figure 2.

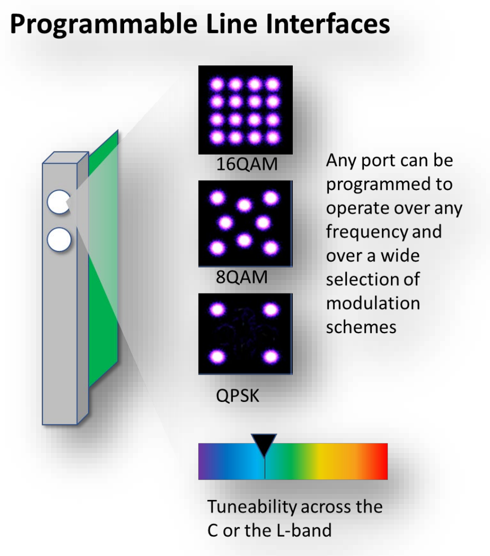

Figure 2 - Programmable Line Interfaces

- **Programmable Optical Carriers/Super-channels**: While DWDM technology disrupted the telecommunication industry by enabling multiple optical carriers to travel in parallel on a fiber to increase fiber capacity, the latest innovation in photonic integration and digital signal processing raised the bar of optical performance and capacity by introducing programmable super-channels. A super-channel includes several optical carriers combined to create a composite line-side signal of the desired capacity that is provisioned in one operational cycle without increasing operational complexity, as depicted in Figure 3. Super-channels are designed to overcome three fundamental challenges: to scale bandwidth without scaling operational procedures, to optimize DWDM capacity/reach and to support the next generation of high-speed services such as 100 GbE and 400 GbE.

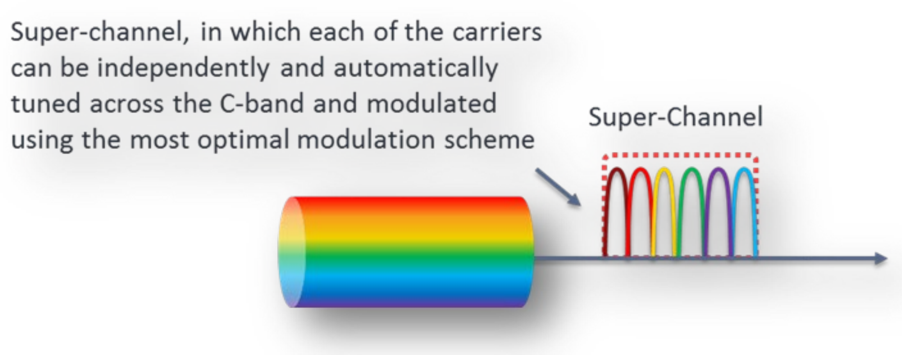

Figure 3 - Programmable Super-Channels

Super-channels are designed to be programmable as they support a capability often referred to as "optical sliceability" by allowing any super-channel to be sliced, so each 100G or N x 100G wavelength can be tuned across the C-band, modulated and then routed in a separate direction from the others to the appropriate destination over any open optical line system, as shown in Figure 4. This ability to "slice and dice" super-channels streamlines operations through programmability of the number and capacity of these bandwidth slices, and it also significantly reduces TCO (power consumption and footprint) while increasing network flexibility.

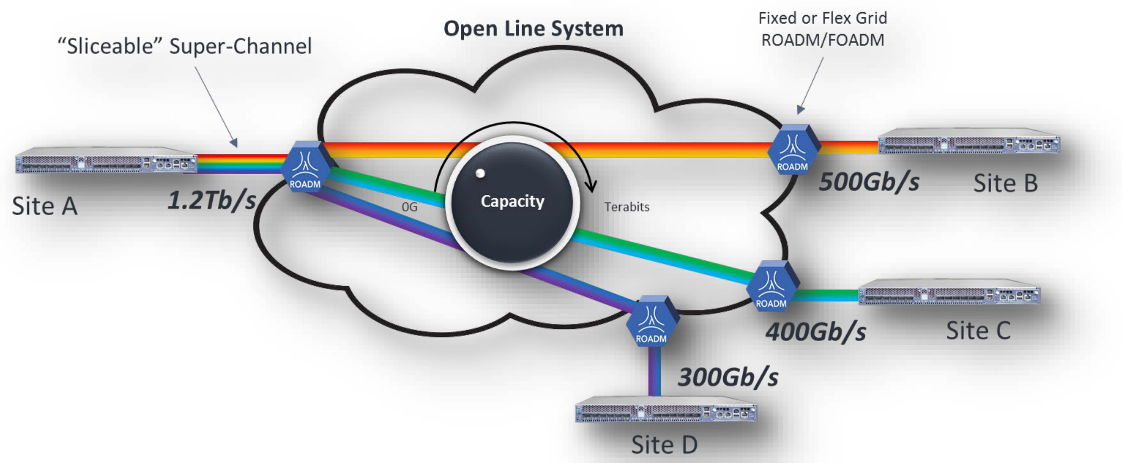

Figure 4 - Programmable Capacity through Sliceable Optics

- **Programmable Service Activation**: Activating new capacity is often complicated, requires truck rolls, is prone to human error and takes months and months to complete. The complexity of this process translates into extended time to revenue, which has a direct and negative impact on service providers' top line. Programmability has also been extended to service activation through the implementation of software defined capacity (SDC) that brings the principles of SDN, which has primarily focused on the Ethernet and packet layers, to the optical transport layer to dynamically add, modify, move and retire optical capacity based on the real-time requirements of upper-layer applications. SDC provides instant software activation of additional capacity, creating a pool of bandwidth that can be dynamically allocated based on traffic demand (Figure 5). Software defined capacity/activation is a true game changer from both business and operational perspectives. It enables a perfect match between the timing of CapEx and service revenue, thus accelerating time to revenue from months to minutes. It also reduces OpEx by streamlining operations and eliminating truck rolls.

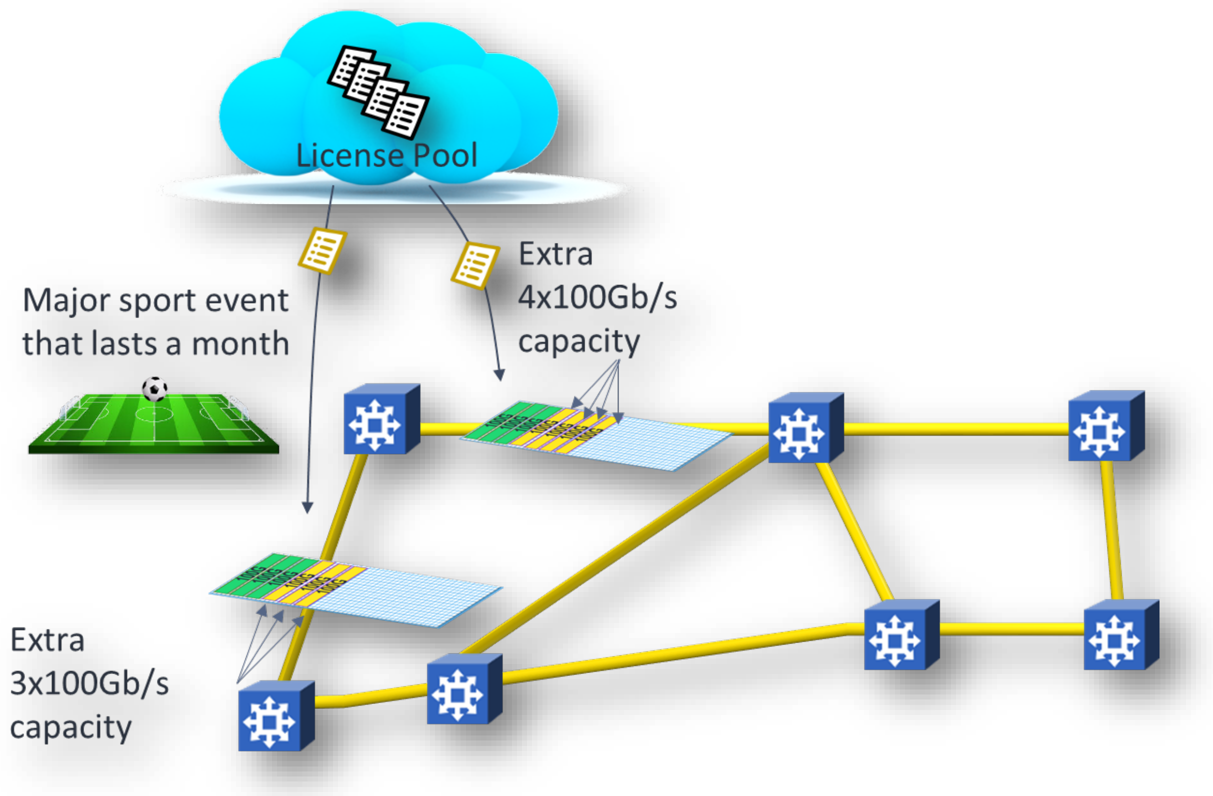

Figure 5 - Dynamically Add, Modify, Move and Retire Optical Capacity with SDC

- **Programmable and Virtualized Infrastructure**: New advancements in software and hardware programmability allow for the creation of virtualized networks at the packet, digital and optical layers and across metro and core domains. Parts of the optical infrastructure can be logically partitioned based on each end customer's (e.g. an enterprise customer's) needs. Premium customers can have dedicated logical partitions of the network with complete visibility and control of their logical network and isolation from other customers, allowing them to customize connectivity and services around their own applications. Similarly, such virtualization capabilities allow network operators to maximize their return on assets and broaden their addressable markets without the additional capital often required to build private physical networks.

- **Programmable Telemetry and Network Pulse**: Cloud and IoT networking rely heavily on streaming telemetry and real-time data analytics to assess network health and proactively avoid failures caused by network degradation. This flow of information between the various parts of the network and the upper-layer software tools and SDN controllers is programmable, leveraging open APIs such as gRPC, RESTCONF, NETCONF/YANG and other northbound interfaces to connect to upper-layer orchestration systems or same-layer intelligent software tools. This programmable flow of telemetry allows the coordination and optimization of resources across the network and the delivery of predictive and prescriptive real-time recommendations and actions for maximum performance.

- **Programmability of Service Restoration**: Network downtime often translates into the loss of millions of dollars as well as irreparable damage to the network operator's reputation and brand image. A network failure can have a disastrous impact on cloud applications where the network plays a vital part in connecting users to the cloud or in machine-to-machine connectivity. Network operators often leverage intelligent software tools such as control planes to operate as the brain of the network, reacting to network changes in real time, without human intervention. These intelligent software tools have evolved to be programmable to increase network availability and protect against failures (e.g. fiber cuts, hardware fa) while maintaining stringent service requirements such as low latency, avoidance of shared risks (e.g. restored traffic going through the same impacted fiber conduit), etc. A programmable control plane can partition a network carrying mission-critical data by allowing specific links, wavelengths, subwavelengths or even nodes to be dedicated to a specified use, with preset thresholds for latency, bandwidth and resiliency. The programmability of service restoration makes networks autonomous, highly reliable and self-healing (Figure 6).

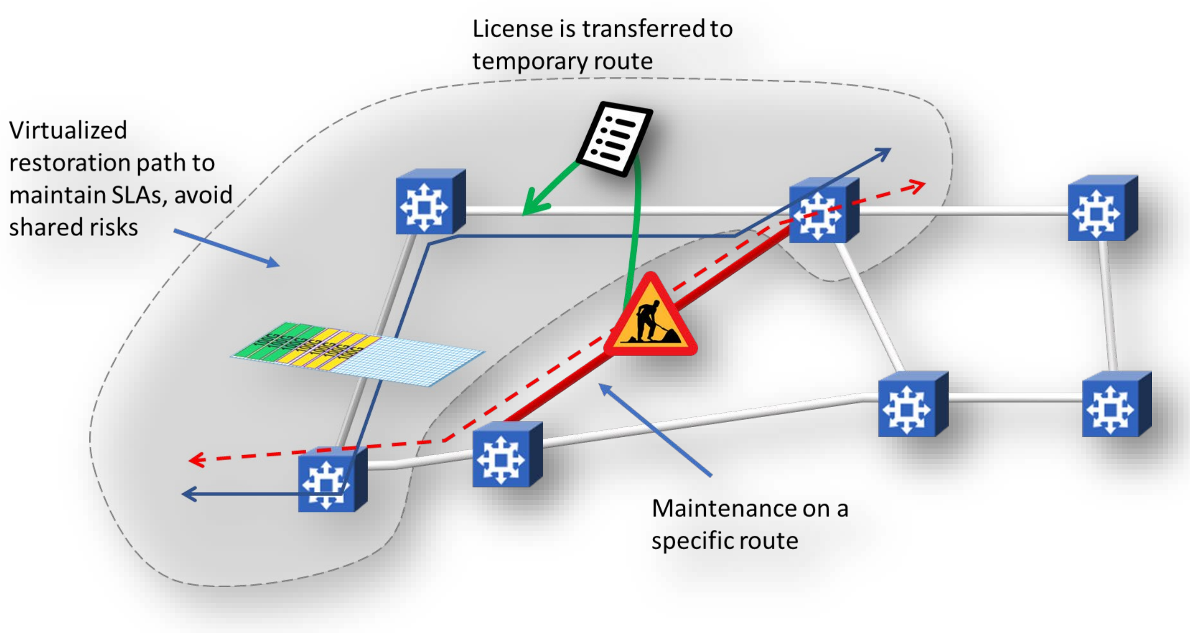

Figure 6 - Programmability of Service Restoration

- **Programmable Inter-layer Service Setup**: Building automation and intelligence across all network layers and various operating tasks is central to underpinning cloud-era networking. SDN controllers and frameworks span network layers to provide programmable and intelligent capabilities to plan, monitor and conduct various network operations, such as real-time capacity planning, bandwidth on demand, network virtualization and many others without human intervention. For example, very sophisticated algorithms and data models can be used to build a microservices-based PCE. The PCE replaces manual offline route and capacity planning processes with highly automated, programmable, real-time service planning and activation over optimal routes across layers, to overcome multiple and often challenging fiber impairments.

3. Network Programmability – A Glimpse into the Future

As discussed earlier, programmability has been implemented across various and key parts of the network, including its components, layers and functions. Nonetheless, it's still poised to further expand and become entrenched in the next generation of connectivity, often referred to as cognitive networking. The following describe how programmability will touch many other aspects of networking, paving the way for highly automated, self-aware, self-organizing and self-optimizing networking infrastructure:

- **Programmable and Highly Granular Optical Transmission**: A practical way to implement higher-order modulation schemes (e.g. 128QAM) and baud rates (e.g. 100 GBaud) in optical transmission systems without a massive compromise in optical performance will certainly emerge. However, advancements in photonic integration and digital signal processing are leading toward a more programmable and highly granular optical WDM line rate through the implementation of hybrid modulations of super-channels and subcarriers. With the ability to program different modulation and dynamically adjust constellation shaping gains on each subcarrier, a variety of spectral efficiencies can be derived [1][2]. This provides higher fiber capacity and a better flexibility for diverse network applications.

- **Analytics and Machine Learning-triggered Network Capacity**: While SDC has already forever changed the way services are planned, provisioned and activated, its next phase will allow real-time network analytics, microservice-based engines and machine-learning algorithms to dynamically increase or decrease network capacity on specific routes based on past trends, spontaneous changes in traffic demand or an anticipated spike in capacity. This data-driven real-time traffic engineering will be fully autonomous, programmable and proactive to identify potential sources of failures before they happen and take the necessary steps to maintain the network at the highest levels of reliability and efficiency.

- **Proactive Network and Traffic Protection**: The use of control plane has significantly increased service and network availability by automatically restoring traffic after a failure over alternative paths while maintaining SLAs. Moreover, Layer 1 or Layer 2 traffic encryption embedded into optical transmission platforms has proven to be an effective way to protect in-flight data from intruders and hacking tools. The evolution of network and traffic protection could be the combination of control plane capabilities, real-time traffic and network topology engineering as well as traffic encryption - all functioning as parts of an intelligent and highly proactive network protection mechanism. Advanced, accurate and fast intrusion detection capabilities could trigger proactive measures to virtually isolate suspicious network areas and automatically encrypt the traffic over specific links. The same proactive capabilities can also be applied to minimize the impact of network maintenance operations by automatically and proactively isolating affected network areas and setting up backup plans during large-scale software upgrades or disruptive network operations.

Conclusion

Network programmability is one of the key building blocks for a successful and fast-paced evolution to the cloud and Blockchain. It's been entrenched across various moving parts of the network, from programmable optical client and line ports to highly sophisticated service creation and activation mechanisms, and it has proven to lead to significant business and operational benefits. Despite such progress, network programmability is still poised for further expansion and evolution by taking advantage of the recent development in analytics and machine-learning tools to elevate the network to a new level of automation, flexibility and efficiency.

Abbreviations

API	application programming interface
CapEx	capital expenditure
DWDM	dense wavelength-division multiplexing
GbE	Gigabit Ethernet
gRPC	gRPC Remote Procedure Call
IoT	Internet of Things
NETCONF	Network Configuration Protocol
OC	optical carrier
OpEx	operational expenditure
OTN	Optical Transport Network
PCE	path computation engine
QAM	quadrature amplitude modulation (8QAM, 16QAM, 64QAM)
QPSK	quadrature phase-shift keying
REST	Representational State Transfer
SDC	software defined capacity
SDN	software-defined networking
SE	spectral efficiency
SLA	service level agreement
TCO	total cost of ownership
YANG	Yet Another Next Generation

Bibliography & References

[1] *Design Considerations for a Digital Subcarrier Coherent Optical Modem*; David Krause, Ahmed Awadalla, Abdullah S. Karar, Han Sun, Kuang-Tsan Wu, OCF 2017

[2] *Subcarriers having different modulation formats*; US patent pending; A. Awadalla, et al.

