

Introduction to the JOCN Special Issue on Open Optical Networks

Original

Introduction to the JOCN Special Issue on Open Optical Networks / Napoli, A.; Curri, V.. - In: JOURNAL OF OPTICAL COMMUNICATIONS AND NETWORKING. - ISSN 1943-0620. - ELETTRONICO. - 12:6(2020), pp. OON1-OON3. [10.1364/JOCN.397704]

Availability:

This version is available at: 11583/2859473 since: 2021-01-04T11:57:02Z

Publisher:

OSA - IEEE

Published

DOI:10.1364/JOCN.397704

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Introduction to the JOCN Special Issue on Open Optical Networks

ANTONIO NAPOLI^{1,*} AND VITTORIO CURRI² 

¹Infinera, Sankt-Martinstr. 76, Munich, Germany

²Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino, Italy

*Corresponding author: anapoli@infinera.com

Received 13 May 2020; published 27 May 2020 (Doc. ID 397704)

We provide an introduction to the paradigm of open optical networks (OON), including its evolution and benefits, followed by a brief summary of the papers in the OON special issue. © 2020 Optical Society of America

<https://doi.org/10.1364/JOCN.397704>

1. OPEN OPTICAL NETWORKS

The open optical network (OON) paradigm is targeted at enabling enhanced interoperability and software control in the optical layer, with the ultimate goals of improved network flexibility, manageability, performance, and economics. It is a broad concept that may encompass hardware, control and management, application programming interfaces (APIs), and data structures. Through the definition of standard interfaces and data-model abstractions, it promotes the softwarization of the network, which is compatible with the growing trend of software-defined networking (SDN). It also enables a mix-and-match best-of-breed strategy with respect to vendor equipment.

OON covers a broad spectrum of implementations, ranging from support for alien wavelengths to fully disaggregated networks (where each network element may be considered as open). Although OON is a relatively new paradigm, it has already had a significant and positive impact on the industry, e.g., by speeding up the development of whiteboxes and pluggable transceivers. This has largely been achieved due to the efforts of recently established industrial initiatives such as OpenConfig, OpenROADM, and the Telecom Infra Project (TIP), in concert with traditional standardization bodies.

To understand the advent of OONs, it is instructive to consider the evolution of optical transmission. Early optical transmission systems were predominantly characterized by intensity modulation with direct detection (IMDD). The strict performance requirements associated with each vendor's transmission system and the lack of adaptability in the network elements led to proprietary single-vendor deployments over fixed network infrastructure (e.g., fiber dispersion maps were implemented with static in-line dispersion-compensation units). In the 2010 timeframe, there was a shift towards coherent optical technologies, which greatly increased the overall capacity of optical line systems. This enabled the explosive

growth of applications such as cloud services, high-speed data-center interconnect (DCI), and distribution of high-definition video. An ancillary benefit was coherent transceivers including advanced digital signal processing (DSP) capabilities, which greatly increased their adaptability and flexibility (such DSP capabilities were absent from IMDD transceivers). Functions such as adaptive channel equalization, rate adaptation, and dynamic dispersion management were now readily supported. The ramifications of this greater flexibility have been profound. It enabled support for dynamic optical networks that can react to rapidly changing traffic demand (a growing trend) and variable network conditions, elastic optical networks that can allocate network resources more efficiently, and low-margin optical networks where much of the excess system margin arising from planning for worst-case scenarios can be removed.

Furthermore, with the ability to dynamically adjust the transmission parameters of each lightpath in order to meet the desired quality of transmission (QoT), simple, yet effective, mathematical models were applicable over a wider range of network conditions. These models are amenable to software implementations that operate on the desired timescales. Each lightpath can be effectively approximated by an additive Gaussian nonlinear channel, where the overall noise is the sum of the amplified spontaneous emission noise introduced by optical amplifiers and the nonlinear interference generated within each fiber span due to the nonlinear crosstalk among the active lightpaths.

Overall, operational advancements and the growing role of software control laid the groundwork for greater interoperability and a move towards open optical networks.

Coinciding with these technical advancements were transformations in the operator/service-provider space. Explosive traffic growth, increased dynamism, and the rise of hyperscale operators changed the competitive environment, especially in the metro DCI sector. This is a complex and competitive ecosystem, which requires fast innovation and development,

a condition that a proprietary-based single-vendor approach cannot fulfill. Consequently, multi-vendor solutions—and the related interoperability between suppliers—were introduced as an alternative to single-vendor closed systems.

Openness was proposed to unleash innovation, reduce costs, and enable adaptive and accelerated development following a pay-as-you-grow approach. OON engenders flexibility in the choice of suppliers, enabling best-of-breed selection for a particular function or scenario. This is gradually leading to the commoditization of both hardware and software. For example, within a fully interoperable ecosystem, an operator can deploy the transponder card from vendor *A*, the reconfigurable optical add/drop multiplexer (ROADM) from vendor *B*, and install the software management from vendor *C*, which, in turn, may utilize APIs from other suppliers. Such flexibility allows operators to enhance existing services and more readily offer new services.

Furthermore, from a vendor perspective, it opens the door to a range of novel applications to manage and utilize the available hardware. Vendors can focus now on specific network elements rather than allocating considerable resources to cover the entire chain of products. This is highly desirable because networks are becoming extremely complex, making it difficult for a single vendor to excel in all aspects. Additionally, several small operators have entered the market and do not have the capability to develop the needed APIs to automate network operation, and thus rely on third party developers.

A significant role in OON realization has been played by component and physical-layer abstraction. Both are needed for open network control, in conjunction with the SDN controller. With respect to component control, there exists data modeling languages/protocols, such as YANG/NETCONF, which support the ability to initialize, configure, and operate network elements. Physical-layer operation has been simplified by the development of approximate transmission models that avoid the time-consuming solution of the complex nonlinear Schrödinger equation that governs propagation over fiber.

The OON paradigm encompasses a range of disaggregation models (partial versus full) and virtualization of network functions. A disaggregated open world mandates an even higher level of scalability for all involved elements and it can work only by providing open APIs. Clearly, OONs require a certain level of standardization for the interfaces managing the hardware and operating the network. Moreover, open APIs may help, for example, in reducing operational inefficiency caused by manual provisioning and restoration, with a goal of full network automation. In this regard, the implementation of the OON paradigm opens up optical network infrastructures to telemetry-based autonomous operation and augmented awareness of the physical layer, enabled by the use of artificial-intelligence functionalities within the control plane.

Several collaborative initiatives have been established by the main players within the industry—operators, vendors, suppliers—to facilitate this process (e.g., OpenConfig, OpenROADM, TIP). These initiatives have paved the way to multi-vendor interoperability demonstrations and helped to establish normative standards for OONs. The work carried out within these initiatives address: open hardware solutions,

data-model abstractions, application interfaces, physical layer software, and implementation of SDN solutions that are suitable for future technologies.

2. SUMMARY OF THE SPECIAL ISSUE PAPERS

The special issue includes eight manuscripts. A brief summary of each work is reported hereafter; we cover the four invited papers first.

“Open and disaggregated optical transport networks for data center interconnects” by C. Xie et al. presents Alibaba’s practices and views regarding its open and disaggregated optical transport network, focusing on metro DCI. The work reports a description of the equipment data models based on OpenConfig YANG models. Additionally, the design and implementation of a home-grown control and management software platform for DCI networks is presented.

“The OpenROADM initiative” by M. Birk et al. provides a detailed overview of the history, goals, design choices, and motivations of the OpenROADM initiative. The current status and the results from interoperability tests are also presented, along with an application involving an OpenROADM-based SDN controller.

As mentioned above, there is a range of disaggregation models encompassed by the OON paradigm. “Open optical communication systems at a hyperscale operator” by M. Newland et al. focuses on two such models: one enables transponder/line-system interoperability and the other supports line-system/line-system interoperability. The discussion includes experiences and lessons learned from Google’s deployed network, including perspectives on optical-link design, software and control, deployment, and operation.

“Modeling and mitigation of fiber nonlinearity in wideband optical signal transmission” by D. Semrau et al. discusses how physical-layer abstraction can enable the development of open and effective network planning tools for interoperable multi-vendor systems. The paper specifically addresses modeling the nonlinear effects observed in wideband optical transmission systems. It presents a closed-form solution that can be used for real-time operation; the paper demonstrates its use in optimizing launch-power spectral profiles.

“Opening up ROADMs: a filterless add/drop module for coherent-detection signals” by J. Kundrát et al. presents an open design for a filterless ROADM add/drop module with a NETCONF northbound interface. Design details and performance results are presented.

“Locally automated restoration in SDN disaggregated networks” by N. Sambo discusses a hybrid centralized/distributed restoration scheme for disaggregated SDN-based networks. The SDN controller uses NETCONF to instruct the network devices how to react if a failure occurs. Simulations show that the scheme can reduce the restoration time with respect to a fully centralized approach.

The other two contributed articles deal with the abstraction of the physical layer and how to exploit this information to operate a network. “GNPy: an open source application for physical layer aware open optical networks” by A. Ferrari et al. investigates the validation of GNPy, which is an application that incorporates a QoT estimator for coherent wavelength

division multiplexed optical networks. Specifically, the validation is performed for mixed-fiber multi-vendor networks over a range of distances, modulation formats, amplifier configurations, and power levels. Results show excellent accuracy in predicting the generalized signal-to-noise ratio (GSNR).

“Using machine learning in an open optical line system controller” by A. D’Amico et al. presents an experimental proof-of-concept on the use of machine learning to predict the GSNR using only the telemetry data from the optical

channel monitors at the end of the line systems, with the optical amplifiers treated as black boxes. Using a deep neural network approach, the authors demonstrate an excellent capability in predicting the GSNR, enabling a significant reduction in the allocated system margin.

Antonio Napoli, Co-Lead Guest Editor
Vittorio Curri, Co-Lead Guest Editor