

Applications of P4-based Network Programmability in Optical Networks

Filippo Cugini¹, Davide Scano², Alessio Giorgetti³, Andrea Sgambelluri²,
Francesco Paolucci¹, Juan Jose Vegas Olmos⁴, Piero Castoldi²

1: CNIT, Pisa, Italy

2: Scuola Superiore Sant'Anna, Pisa, Italy

3: IEIIT institute, CNR, Pisa, Italy

*4: Nvidia, Ledreborg Allé 130b, Roskilde, 4000, Denmark.
filippo.cugini@cnit.it*

Abstract: This paper presents potentials and challenges of disaggregated metro-edge networking based on open packet-optical nodes encompassing coherent pluggable modules, SONiC open operating system, and P4-based packet switching programmability. © 2021 The Author(s)

1. Introduction

In the last years, compact coherent transmission modules have emerged, moving from traditional 8RU modules to pluggable form factors such as CFP2-DCO and QSFP-DD [1]. Although some limitations still exist with the smallest modules, the trend clearly shows that in the next few years coherent pluggable modules may become the preferred technology particularly for metro-edge scenarios. Indeed, besides space and power savings, such miniaturization process will enable installing coherent modules on off-the-shelf cost-effective packet switches designed for data center operations. This has the potential to open the traditionally closed optical network market to new manufacturers while taking advantage of the high pace of innovation within data centers. Such innovation occurs at both hardware and software levels. At the HW level, besides guaranteeing continuously increasing throughput performance, a new generation of programmable ASIC has been introduced, leveraging on the P4 technology [2]. P4 enables the implementation of novel in-network functions operating at wire speed. Examples includes innovative monitoring and telemetry solutions, cyber-security operations, hardware acceleration of 5G network functions [3]. At the software level, open-source initiatives have pushed disaggregated solutions to production, providing high level of reliability and performance. Among them, the SONiC open operating system (OS), originally designed for intra-data center operations, is gaining consensus in different network scenarios and could represent an attractive OS solution also for next generation packet-optical nodes equipped with coherent pluggable modules [4, 5].

In this work, we present the state of the art, potentials and open issues of open packet-optical nodes supporting coherent pluggable modules, open operating system, and enabling P4 packet switching programmability. Such combined set of technologies is particularly attractive since it may lead to both CAPEX and OPEX savings while guaranteeing more effective traffic engineering solutions. For example, packet-optical nodes may enable the removal of standalone transponders (reducing latency while avoiding O/E/O conversions) and may provide a tight and effective integration between the packet and the optical layers.

2. Packet-Optical node: enabling technologies

Coherent pluggable modules. Digital Coherent Optics (DCO) transceivers are commercially available at rates of up to 400 Gbps in two form factors: CFP2 (C form-factor pluggable type 2) and QSFP-DD (Quad Small Form-factor Pluggable Double Density). The CFP2-DCO transceiver (41.5mm width) provides adequate signal launch power of 0dBm, suitable for metro/regional interconnections traversing multiple Reconfigurable Optical Add Drop Multiplexers (ROADMs). Experiments have shown the capability to cover even 1500km at 400Gb/s in a 75 GHz-spaced DWDM System, using 16 Quadrature Amplitude Modulation (QAM) at 69 Gbaud, probabilistic constellation shaping, and soft-decision forward error correction [6]. In addition, excellent interoperability performance has been already achieved by DCO-CFP2 based on OpenROADM multi source agreement (MSA) [1]. The 400ZR coherent pluggable module based on the smaller QSFP-DD form factor (18.35 mm width) currently provides signal launch power of up to -10dBm and it is mainly designed for data center interconnections, covering single span distances of up to 120km. The larger CFP2 form factor is typically not supported by DC ethernet switches, while QSFP-DD is generally supported even if some electrical power constraints may still induce operational issues. To take full advantage of DC packet switching solutions, a new generation of coherent pluggable modules is needed, able to provide improved optical transmission performance in QSFP-DD form factor.

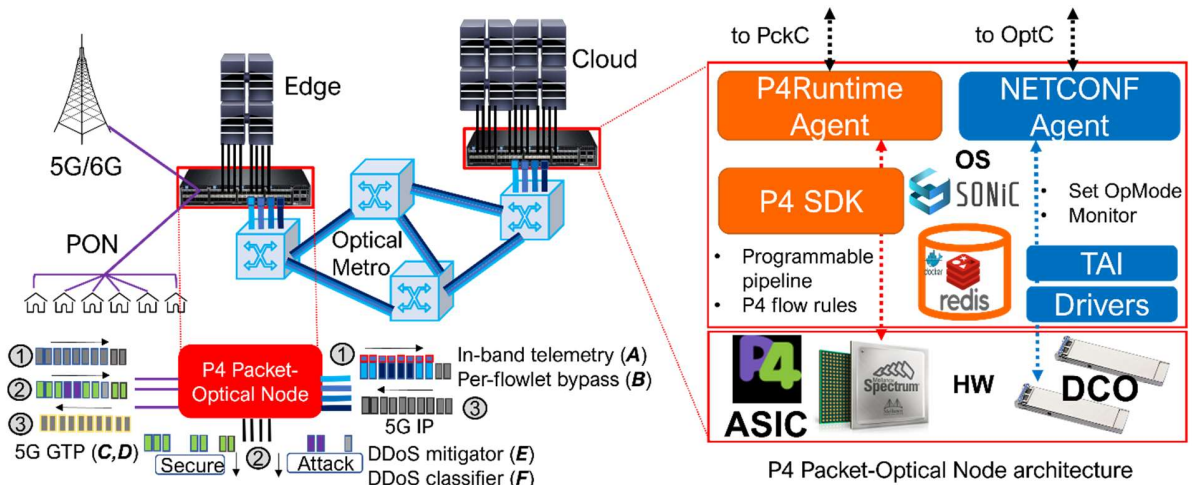


Fig.1. Metro/edge optical network based on open packet-optical nodes supporting advanced P4 applications.

Node operating system supporting coherent pluggable modules. SONiC (Software for Open Networking in the Cloud) is an open-source network operating system already deployed in intra-DC production networks and it is also gaining consensus for devices (including smartNICs) targeting other network segments, including optical networks [5]. It introduces an innovative modular architecture supporting docker-based containerization enabling applications and network tools/protocols to be deployed and operated in secure isolated way. SONiC natively supports the typical network protocols used for DC operations (e.g., Border Gateway Protocol) while it lacks official support to additional relevant protocols for optical networking, such as NETCONF. Overall, as of today, SONiC is not yet a mature solution for production telecom infrastructures, however relevant work by a large industrial community of developers is ongoing and remarkable progresses are introduced at every new release. In addition, at the time of writing this paper, the drivers to support coherent pluggable modules (e.g., 400ZR) have been announced but they are not available yet. The Telecom Infra Project (TIP) is also actively working on a SONiC-based OS (named Goldstone) as well as on transceiver abstraction interfaces (TAI) to support coherent modules [4].

SDN Control of packet-optical nodes. The OpenConfig and OpenROADM initiatives have made a significant effort to standardize YANG models for optical networking [7]. Additional effort is still required to fully and efficiently control coherent pluggable modules. Moreover, open issues are present in the coordination workflow between configurations performed at the packet and optical layer. Traditionally, optical network SDN Controllers (OptC) oversee the configuration of both ROADMs and transponders. When transponders are replaced by pluggable modules equipped within the packet-optical nodes, traditionally controlled only by Packet SDN Controller (PckC), coordinated control need to be provided. Two main solutions are under discussion. In the first solution [5], both OptC and PckC have direct access to the node (see the example in Fig. 1). Coordinated control is achieved by properly providing reading/writing rights to selected portions of the YANG-compliant configuration. However, this kind of coordination is not sufficient since both controllers may take uncorrelated and potentially inadequate actions (e.g., in case of failure). For this reason, the second solution based on hierarchical coordination through a Parent Controller seems the most promising one [8]. In this case, the PckC remains the unique controller for the packet-optical node. The PckC retrieves from the node the supported capabilities of the pluggable module (e.g., supported operational modes), provides this information to the Parent Controller which forwards it to the OptC. OptC is then able to perform end-to-end impairment-aware path computation and compute the suitable transmission parameters. For example, the selected operational mode among those supported is communicated to the Parent which forwards the information to the PckC, in charge of enforcing it into the node. Although this solution enables better coordination, scalability and efficiency issues may be present. Therefore, additional studies and standardized solutions are needed, including finite state machines for effective coordination among SDN Controllers.

3. Applications of P4-based programmability

The limitations that currently affect the three aforementioned enabling technologies are expected to be overcome in relatively short term. This would open the opportunity of adopting within packet-optical boxes innovative solutions

designed for DC, including programmable P4-based ASIC. In the following, a selection of P4 applications is reported that could be of potential interest for integrated packet-optical networking (Fig. 1).

A. Monitoring and Telemetry. The most successful P4 application provides advanced monitoring and telemetry capabilities. For example, In-band Telemetry (INT) enables the introduction of custom packet headers including metadata such as timestamps and the time spent in the traversed queues. This enables accurate monitoring across the whole packet and optical network, potentially leading to improved traffic engineering solutions [9-11].

B. Latency-aware scheduling and forwarding. The possibility to accurately monitor edge-to-cloud latency performance through INT, besides global optimizations, enables dynamic local scheduling and QoS-driven forwarding operations. For example, the local selection among alternative routes across optical and packet resources (e.g., optical bypass) may be dynamically performed even on a per-flowlet basis, instead of adopting the relatively rigid tributary traffic assignment implemented with today's muxponders and transponders [11, 12].

C. 5G function acceleration. Packet-optical nodes may be used to interconnect metro optical resources with 5G X-haul. In this case, P4 has the capability to offload specific 5G functions, such as the User Plane Function (UPF), directly performing GTP protocol encapsulation/decapsulation function and traffic steering [13].

D. Broadband gateway services. In this case packet-optical nodes are used to interconnect the optical metro with the access (e.g., even embedding pluggable OLTs). Open Broadband Network Gateway (BNG) is a TIP initiative leveraging on P4 acceleration for encap/decapsulation and partially provide hierarchical QoS scheduling. Embedding BNG operations in a packet-optical metro node would release the need for dedicated additional BNG hardware [14].

E. Cyber-security. The stateful capabilities of P4 nodes enables the implementation of in-network firewalling solutions operating at wire speed and potentially deployable in all metro nodes. This would constitute a distributed security barrier across the entire network [15].

F. In-network AI. Recent studies are investigating the possibility to run Artificial Intelligence (AI) within P4 ASIC. This would open for disruptively new applications leveraging on in-network AI at wire speed [16, 17].

Conclusions

This paper first introduced the state-of-art and potential of packet-optical nodes based on disaggregated switches designed for data centers. Current open issues in terms of support of coherent pluggable modules, local and SDN Control of the packet-optical nodes were highlighted and discussed. Such issues are expected to be overcome relatively soon, enabling metro-optical networks benefitting from the high pace of innovation of the DC market and technologies, including stateful P4-based ASIC programmability at wire speed. A list of P4 applications suitable for metro packet-optical networks was finally reported, including fine-grade QoS control, monitoring, 5G acceleration, cyber security and in-network AI.

Acknowledgment. This work has been partially supported by the BRAINE project (www.braine-project.eu), funded by ECSEL JU under grant agreement N. 876967, by B5G-OPEN H2020 Project (www.b5g-open.eu) under grant agreement N. 101016663 and by the Department of Excellence in Robotics and Artificial Intelligence of Scuola Superiore Sant'Anna funded by Ministry of University and Research.

References

- [1] E. Pincemin et al., "Interoperable CFP-DCO and CFP2-DCO pluggable optic interfaces for 100G WDM transmission," OFC 2019.
- [2] P. Bosshart et al., "P4: Programming protocol-independent packet processors", ACM SIGCOMM, 2014.
- [3] F. Paolucci, F. Civerchia, A. Sgambelluri, A. Giorgetti, F. Cugini, and P. Castoldi, "P4 Edge Node enabling Stateful Traffic Engineering and Cyber Security," IEEE/OSA Journal of Optical Communications and Networking, vol. 11, no. 1, pp. A84–A95, Jan. 2019.
- [4] V. Lopez et al., "Enabling fully programmable transponder white boxes [Invited]", J. Opt. Commun. and Netw., 2020.
- [5] A. Sgambelluri et al., "Coordinating Pluggable Transceiver Control in SONiC-based Disaggregated Packet-Optical Networks", OFC 2021.
- [6] <https://www.neophotonics.com/press-releases/?newsId=12486>
- [7] A. Sgambelluri, A. Giorgetti, D. Scano, F. Cugini, F. Paolucci, "OpenConfig and OpenROADM Automation of Operational Modes in Disaggregated Optical Networks", IEEE Access 8, 190094-190107, 2020.
- [8] D. Scano et al., "Hierarchical Control of SONiC-based Packet-Optical Nodes encompassing Coherent Pluggable Modules", ECOC 2021.
- [9] J. Vestin, A. Kassler, D. Bhamare, K.J. Grinnemo, J.O. Andersson, G. Pongracz, "Programmable event detection for in-band network telemetry", 2019 IEEE 8th international conference on cloud networking (CloudNet) 2019 Nov 4 (pp. 1-6), IEEE.
- [10] F. Cugini, P. Gunning, F. Paolucci, P. Castoldi, A. Lord, "P4 in-band telemetry (INT) for latency-aware VNF in metro networks", OFC 2019.
- [11] D. Scano, F. Paolucci, K. Kondepu, A. Sgambelluri, L. Valcarengi, F. Cugini, "Extending P4 in-band telemetry to user equipment for latency- and localization-aware autonomous networking with AI forecasting", Journal of Optical Communications and Networking, 2021
- [12] I. Pelle, F. Paolucci, B. Sonkoly, and F. Cugini, "Latency-sensitive edge/cloud serverless dynamic deployment over telemetry-based packet-optical network," IEEE Journal on Selected Areas in Communications, pp. 1-1, 2021.
- [13] F. Paolucci, F. Cugini, P. Castoldi, T. Osinski, "Enhancing 5G SDN/NFV edge with P4 data plane programmability", IEEE Network, 2021.
- [14] R. Kundel et al., "OpenBNG: Central office network functions on programmable data plane hardware", J. of Network Management, 2021.
- [15] F. Musumeci, et al., "Machine-Learning-Enabled DDoS Attacks Detection in P4 Programmable Networks", Journal of Network and Systems Management 30 (1), 1-27, 2021.
- [16] G. Siracusano et al., "Running Neural Networks on the NIC". arXiv preprint arXiv:2009.02353, Sept. 2020.
- [17] F. Paolucci, L. De Marinis, P. Castoldi, F. Cugini, "Demonstration of P4 Neural Network Switch", OFC 2021.