

# Coordinating Pluggable Transceiver Control in SONiC-based Disaggregated Packet-Optical Networks

Andrea Sgambelluri<sup>1</sup>, Davide Scano<sup>1</sup>, Alessio Giorgetti<sup>2</sup>, Francesco Paolucci<sup>4</sup>, Emilio Riccardi<sup>3</sup>, Roberto Morro<sup>3</sup>, Piero Castoldi<sup>1</sup>, Filippo Cugini<sup>4</sup>

*1: Scuola Superiore Sant'Anna, Pisa, Italy*

*2: IEIT-CNR, Pisa, Italy*

*3: TIM, Turin, Italy*

*4: CNIT, Pisa, Italy*

*andrea.sgambelluri@santannapisa.it*

**Abstract:** Effective control of pluggable transceivers in SONiC-based packet-optical nodes is demonstrated. A workflow for multi-layer recovery upon soft failure detection is validated, showing no traffic disruption and fast node-driven coordination between packet and optical operations.

© 2021 The Author(s)

## 1. Introduction

Disaggregated optical networks have attracted remarkable interest due to potential savings in CapEx as well as for their fully standardized open interfaces for Software Defined Networking (SDN) [1-3]. OpenConfig and OpenROADM represent the most relevant initiatives for standardizing disaggregated solutions. The former is more focused on the device model definition and is designed for partial disaggregation where transponders work in pairs while the optical transport network is managed and controlled by a single vendor. The latter is designed for full disaggregation where both transponders and reconfigurable optical add drop multiplexers (ROADMs) are controlled in a vendor neutral way. Both define YANG models for NETCONF-based SDN agent implementations [4]. Considerable work on disaggregation is also ongoing in the telecom infra-project (TIP) and in the Open Disaggregated Transport Network (ODTN) initiatives [5]. In this context, most of the scientific work on disaggregation has focused on transmission modules as standalone network elements, like transponders and muxponders [6-8]. However, the recent advances in transmission technology have driven the introduction of coherent pluggable transceivers that can be equipped within packet switching devices. For example, Digital Coherent Optics (DCO) transceivers are commercially available at rates of 400 Gbps with configurable transmission parameters in different form factors, such as CFP2 and the smaller QSFP-DD 400ZR.

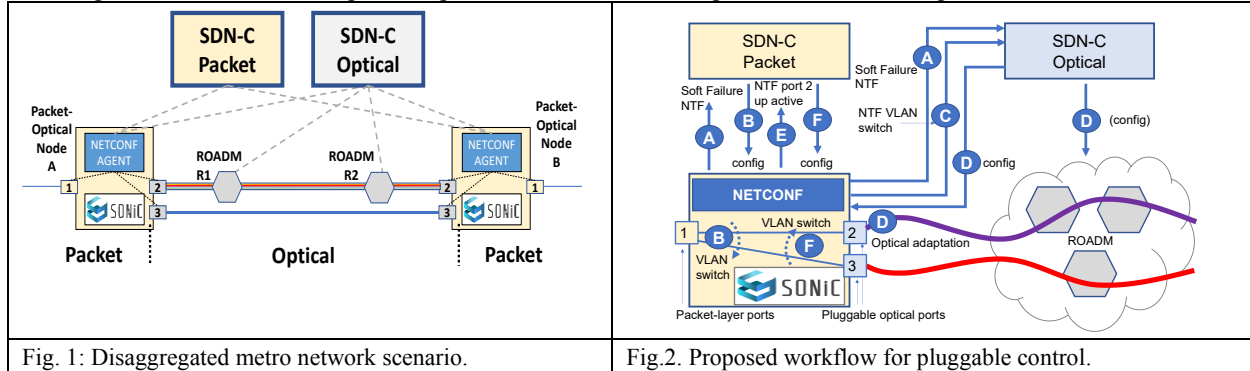
Replacement of standalone transponders with pluggables modules in the packet devices directly connected to the optical network drives significant benefits in terms of CapEx, power consumption and occupied space in central offices. Furthermore, it enables a tight integration between packet and optical networks, which is of special interest as transport is dominated by Ethernet and IP traffic. For example, a single packet switch can provide both intra-data center (DC) traffic aggregation and, thanks to coherent pluggables, effective DC-to-DC interconnection. However, controlling packet-optical solutions requires a complete operating system that is much more complex than traditional NETCONF/YANG software agents employed in standalone transponders [4,5]. SONiC (Software for Open Networking in the Cloud) is an open-source network operating system already deployed in production intra-DC networks and it is also considered a strong candidate to control packet-optical nodes although some operational extensions are needed to fill the existing architectural gaps. For example, SONiC does not natively support NETCONF and it does not encompass the needed software components to operate on coherent pluggable transceivers. Work is ongoing within TIP to provide Transponder Abstraction Interface (TAI), enabling abstracted control of pluggable transceivers [9]. Another gap to be filled is the coordination between packet and optical parameters on the same node, which are often provided by two different SDN controllers, one in charge of packet resources and one in charge of optical transport. So far, this aspect is yet undiscussed in the scientific literature.

This work designs and implements a novel comprehensive workflow enabling coordinated control by SDN packet and optical controllers concurrently operating on a packet-optical node equipped with coherent pluggable modules and using SONiC enhanced with NETCONF/YANG components.

## 2. Reference disaggregated scenario and problem description

The reference metro network disaggregated scenario is illustrated in Fig.1. Two types of nodes are present: ROADMs providing optical switching and packet-optical nodes providing packet switching. Packet-optical nodes are equipped with pluggable transceivers. In large metro networks, a single controller with visibility on both packet and optical resources is hardly implementable due to scalability issues. Two controllers are then typically considered: an Optical SDN Controller (OptC) in charge of the the disaggregated optical transport network (e.g., adopting ODTN solutions for the support of OpenROADM and OpenConfig models [4]) and a Packet Controller (PckC) supporting Layer 2-7 configurations. Traditionally, each SDN controller has full visibility on all components and software

modules of every controlled network element. For example, OptC, besides configuring the ROADMs, is responsible for the configuration of the line interfaces of transponders. However, in the considered scenario, transponders are replaced by pluggable modules equipped within the packet-optical nodes, traditionally controlled only by PckC. That is, two different controllers need to concurrently operate on packet-optical nodes. Thus, a proper workflow needs to be defined to enable the SDN agent of the packet-optical node to coordinate the operations imposed by each controller. Indeed, without proper coordination, complex multi-layer operations, such as recovery upon soft failure, would lead to management conflicts on the packet-optical nodes as well as to potential traffic disruptions.



### 3. Proposed solution for coordinated control of pluggable modules

The proposed workflow used to coordinate PckC and OptC operations is reported in Fig.2 (steps A-F). The workflow exploits the NETCONF-based SDN agent deployed in the packet-optical node. Both controllers are connected to the agent. In particular, the agent exposes the whole YANG model including packet and optical descriptors. To avoid non-standard and complex peer/hierarchical operations, the two controllers do not communicate each other to coordinate their actions. Instead, they leverage on the proposed workflow to avoid conflicts and guarantee segregation of control. Ownership segregation has been implemented exploiting the NCACM solution as detailed in RFC 8341. In particular, the OptC is provided with writing rights on the optical parameters and read-only rights (including enabling notifications upon subscriptions) on packet parameters. Similarly, PckC is provided with writing rights on packet parameters and read-only rights on optical parameters.

It is assumed that two lightpaths are configured on the network. Upon soft failure detection affecting the lightpath provisioned through the optical pluggable module (e.g., port 2 in Fig.1), a NETCONF notification is sent to both controllers (step A). This triggers PckC to initiate the recovery workflow, while OptC becomes aware of the soft fault but, to avoid concurrent operations potentially leading to traffic disruption, it does not enforce optical reconfigurations yet. In step B, PckC enforces new forwarding rules (e.g., new routes/VLANs associations) to an alternative pluggable module (from port 2 to port 3 in the figure) exploiting the protection lightpath. At the end of step B, step C triggers a further NETCONF notification to OptC, indicating that no tributary traffic is forwarded by the pluggable 2. This triggers step D: OptC at this stage can enforce the optical transmission adaptation of the pluggable modules and the potential reconfigurations of the transit ROADMs along the lightpath affected by soft failure. Once the adaptation procedure is concluded, the lightpath of pluggable 2 returns available, and a NETCONF notification is sent, notifying the end of the optical recovery (step E). This allows PckC to revert the routes/VLANs associations to the original state, successfully reconfiguring tributary traffic through the optical transceiver of port 2 (step F).

### 4. Experimental validation and results

The proposed solution has been implemented in a network testbed reproducing the scenario of Fig.1. The testbed includes two ROADMs, two packet-optical nodes equipped with pluggable transceivers, and two ONOS-based SDN Controller [5]. Within the packet-optical node, OpenConfig model is assumed for optical resources while a simplified YANG model for VLAN switching is considered for packet forwarding. Packet-optical node architecture is depicted in Fig. 3, it consists of a Mellanox SN2010 Ethernet switch running SONiC operating system over ONIE. On top of SONiC, a specifically designed docker container runs the ConfD-based NETCONF agent communicating with the controllers. The container retrieves node status information by directly accessing the SONiC Redis database and enforces node configurations using custom-built REST APIs integrated within SONiC. Node ports 1 and 3 are equipped with 10Gb/s SFP+ pluggable transceivers, monitored by SONiC pmon container; VLAN settings are applied through the SONiC swss and syncd containers. Port 2 is attached to an external 100 Gb/s coherent system configured as being a pluggable module, i.e., its driver is accessed via REST by the docker only, with no direct connection to the

SDN controller as it would be for a standalone transponder. As also illustrated in Fig.1, Port 1 acts as tributary interfaces, port 2 handles the working 100Gb/s coherent communication across the optical transport network. Port 3 provides the alternative path for protection purposes. Under working condition, tributary traffic of port 1 is forwarded to port 2 only. Then, soft failure is generated by using a Variable Optical Attenuator (VOA).

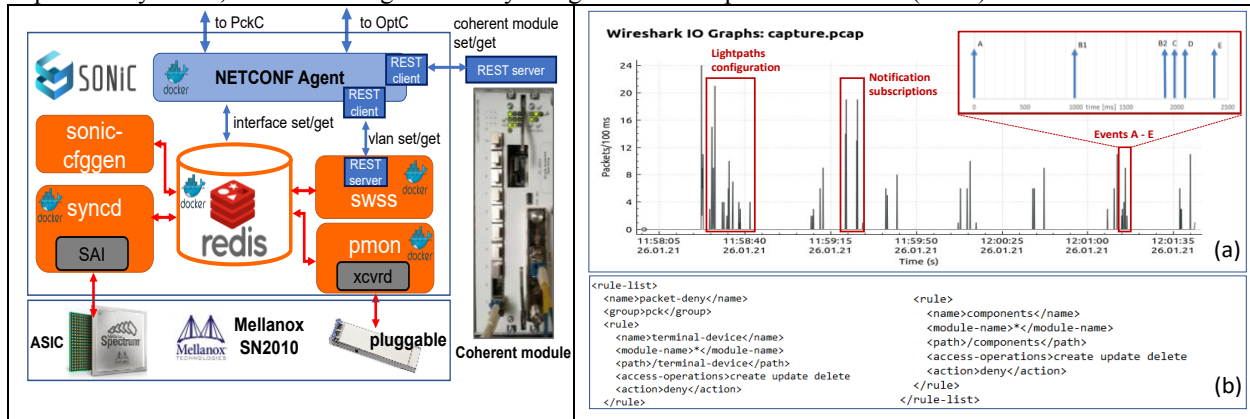


Fig. 3: packet-optical node internal architecture based on Mellanox SN2010 and SONiC.

Fig. 4: NETCONF messages captured between the agent and the controllers; xml scheme implementing RFC 8341.

Fig. 4 shows the Wireshark capture report of the NETCONF messages exchanged by the SONiC container and the two controllers. First, edit-config messages are exchanged to establish the two lightpaths, second the two controllers subscribe to the notification stream of the agent, third the soft failure occurs and it is notified to the controllers that employ the proposed workflow. Other messages are periodically generated by the controllers for synchronization purpose. Workflow events are zoomed in the figure inset. In particular, failure notification is generated at time 0 (step A). Then PckC elaborates the rerouting options, e.g., exploiting the alternative path through port 3. After 996 ms (step B1), the container applies the received VLAN configuration thus rerouting the traffic on the backup lightpath (traffic rerouting is performed without loss of packets). The previous VLAN configuration is deleted at Step B2 (time: 1881 ms), triggering a new notification message (Step C time: 1984 ms) that informs the OptC that no traffic is anymore forwarded through port 2. Thus, OptC generates an edit-config message configuring a new operational mode (e.g., Forward Error Correction - FEC - adaptation) on port 2; this message is received at the agent at time 2072 ms (Step D). This configuration is applied at time 2372 ms (Step E), however the corresponding data plane operation requires around 38 seconds. During this time, port 2 becomes unavailable. Therefore, without the proposed coordination and traffic rerouting significant traffic loss would have been experienced, while, thanks to the coordinated workflow, no traffic disruption is experienced. The overall measured time between step A and E is always less than 2.5 seconds (out of ten experiments), with around 2 seconds required by the data plane configuration implemented in SONiC, while around 0.5 seconds are due to the control plane operations.

## Conclusions

SONiC-based packet-optical nodes including pluggable transceivers are successfully exploited in a multi-layer recovery scenario involving two SDN controllers for packet and optical resources. Fast event coordination is achieved without experiencing traffic disruption, completing the workflow in less than 2.5 seconds mainly due to hardware configuration.

**Acknowledgment.** This work has been partially supported by TIM under the UniversiTIM project and by the BRAINE Project, funded by ECSEL Joint Undertaking under grant agreement No. 876967.

## References

- [1] E. Riccardi, et al., "An Operator view on the Introduction of White Boxes into Optical Networks", JLT 2018.
- [2] Chongjin Xie, et al., "Open and disaggregated optical transport networks for data center interconnects [Invited]" JOCN 2020.
- [3] J. Hernandez, et al., "Comprehensive model for technoeconomic studies of next-generation central offices for metro networks", JOCN 2020.
- [4] A Sgambelluri, A Giorgetti, D Scano, F Cugini, F Paolucci, "OpenConfig and OpenROADM Automation of Operational Modes in Disaggregated Optical Networks", IEEE Access 2020.
- [5] A. Giorgetti, et al., "Control of open and disaggregated transport networks using the Open Network Operating System (ONOS)", JOCN 2020.
- [6] B. Mirkhanzadeh et al., "Demonstration of Joint Operation across OpenROADM Metro Network, OpenFlow Packet Domain, and OpenStack Compute Domain", OFC 2020
- [7] M. Birk et al, "The OpenROADM initiative [Invited]", JOCN 2020.
- [8] R Casellas et al, "Abstraction and Control of Multi-Domain Disaggregated Optical Networks with OpenROADM Device Models", JLT 2020.
- [9] V. Lopez, et al., "Enabling fully programmable transponder white boxes [Invited]", JOCN 2020.