Designing Disaggregated Optical Networks

Filippo Cugini
CNIT
Pisa, Italy
filippo.cugini@cnit.it

Francesco Paolucci, Andrea Sgambelluri,
Alessio Giorgetti, Davide Scano, Piero Castoldi
Scuola Superiore Sant'Anna
Pisa, Italy
fr.paolucci@santannapisa.it

***Abstract*— Disaggregation is driving the design of next generation metro/regional networks. This work presents the different disaggregation flavors and discusses the current maturity level and open issues, highlighting future opportunity for disaggregation beyond current metro optical scenarios.**

***Keywords*— white box, SDN, OpenConfig, OpenROADM, pluggable, SONiC, ONIE, Open Networking Linux (ONL)**

I. Introduction

Disaggregation in optical networks has attracted significant Telecom operator's interest since it avoids vendor lock-in, leading to potential CAPEX savings [1-5]. This paper first provides an overview of the status of the two main flavours of optical disaggregation that have moved to consistent implementations [6-9]: partial and full disaggregation, the former being much more attractive than the latter, which however is actively supporting standardization initiatives. Then, future directions for optical network disaggregation are highlighted, mainly driven by the emerging availability of pluggable modules and open network operating software.

II. Disaggregation Flavours

Two main flavours of optical disaggregation have moved to consistent implementations: partial and full disaggregation.

A . Partial Disaggregation

The first flavour, called partial disaggregation, considers the optical transport infrastructure (e.g., optical line systems – OLS – including ROADMs and amplified links) managed by a single vendor, while transponders, provided in pairs, rely on vendor-neutral NETCONF/YANG control [10]. The most relevant standardization initiative providing YANG models for partial disaggregation is OpenConfig, led by Google. In OpenConfig, transponders are described by the terminal-device YANG model, relying on the concept of logical channels to define the mapping among client ports and line ports. At the line side, the model specifies OTN config and state parameters (e.g., pre/post Forward Error Correction - FEC). The OpenConfig terminal device model is augmented with a platform component defining config parameters such as (i) frequency, (ii) target output power, and (iii) operational modes (OP modes). No other transmission parameters are defined (e.g., modulation format, FEC type, constellation shaping, etc). Indeed, all these parameters have to be defined within the OP modes, which are vendor-specific. This way the model remains stable despite of technological evolutions and capable of including proprietary advanced transmission solutions maximizing throughput performance. As a drawback, this approach requires definition and implementation of specific workflows to manage the OP modes that go beyond the model itself [11-12].

B . Full Disaggregation

The second flavour, called full disaggregation, further disaggregates the OLS, considering ROADMs as white boxes to be controlled in a vendor-neutral way. The most relevant standardization initiative providing full disaggregation is OpenROADM, led by AT&T.

OpenROADM is a complete multi-source agreement (MSA) defining both data and control aspects. At the data plane side, Single-wave (W) interfaces define the optical specifications for the full C-band tunable DWDM optical line interface of the transponder that connects to a (Wr) add/drop port on the ROADM device. Line-side pluggable type have to be CFP-DCO, CFP2-ACO or CFP2-DCO with LC connectors. At the control side, OpenROADM defines models to describe the device, the network and the service level. For example, devices as transponders and ROADMs and amplifiers are fully detailed in the YANG model using a tree structure which includes node identification information (i.e., node-id, node-number, node-type, vendor, model, serial-id), a list of shelves, a list of circuit-packs, a list of interfaces, two lists for internal and external links and finally a list of degrees. Each shelf presents general descriptors (i.e., shelf-name, shelf-type, rack, administrative-state, vendor, model, serial-id, hardware-version, operational-state) and a list of slots. Each circuit-pack, besides general information, includes a list of ports with several descriptors (i.e., port-name, port-direction, label, circuit-id, administrative-state, operational-state, partner-port info used to include bidirectional connection details). In the list of interfaces all the available virtual-interfaces are exposed, including the related general details, the interface type (i.e., OTS, OMS) and the physical/virtual port supporting it. The external link list exposes all the topological information related to external connections with other ROADMs.

C . Partial vs. Full Disaggregation

So far, partial disaggregation has attracted significant interest from operators and vendors, since it neglects most of the optical data plane complexity without significantly compromising on transmission performance. Indeed, transponders are provided in pairs and can implement even

978-3-903176-21-8

proprietary transmission solutions by defining specific operational modes [11-12]. For these reasons, OpenConfig has gained most of the consensus for transponder control.

The business model that drives partial disaggregation considers that the optical transport (i.e., OLS) is a mature technology that, once deployed, can last for more than ten years without relevant upgrades. Moreover, OLS is an analog complex system and disaggregating it may lead to critical implementation, control and management, particularly considering the entire life-cycle of the system. On the other hand, transmission technology is evolving at extremely high pace, with impressive and continuous increase of the symbol rate. In some network scenarios, to cope with the continuous increase of data traffic, transponders can be replaced with higher rate versions even every three years. Thus, partial disaggregation enables Operators not to be bounded to a single vendor for new transponder deployments, still enabling a single vendor to have full responsibility of the OLS. In order to be practical and sustainable, typical partial disaggregated deployments are expected to involve, at least in the first phase, just two vendors per metro network: the one providing the OLS and most of the transponders and a second vendor providing the remaining percentage of transponders (e.g., 20-30%). The control and maintenance will be in charge of the first OLS vendor, which will have to take, from the operator perspective, full responsibility of the entire optical metro network operations. This approach would not require the Telco to have internal skills and effort to manage the disaggregated optical network, relaxing one of the most relevant aspect that concerned the Operators about disaggregation.

The business model that drives full disaggregation appears to be less evident at the moment, and potentially beneficial only as long term approach. Mainly for this reason, multivendor OpenROADM demonstrations have been so far presented by few vendors only, in controlled lab environments, and still relying on proprietary solutions for OAM functions. However, the OpenROADM full disaggregation initiative has already provided remarkable results, particularly in several data plane aspects of the MSA, for example accelerating the definition of standardized pluggable solutions.

III. Packet-Optical White Box

Although partial disaggregation flavour and OpenConfig control of transponders have reached good consensus, many different aspects are still at early standardization phase.

Indeed, the presence of multiple parallel standardization initiatives (e.g., IETF, Telecom Infra Project (TIP), in addition to OpenConfig and OperROADM) has certainly decelerated the adoption of the whole set of open solutions, with key players sometimes just waiting for the scenario to consolidate or making selective choices on the solutions to support. On the other hand, the clear separation between hardware and software enabled by disaggregation is facilitating the emergence of new players providing cost-effective white box or advanced SDN solutions. In addition, new technologies and solutions are rapidly emerging, rediscussing the future directions of optical networks.

The most relevant innovation that has attracted significant attention is the pluggable coherent module (e.g., CFP2-DCO, supporting 100 Gb/s PM-QPSK and 200 Gb/s 16QAM transmissions, e.g. [13]). Multi-vendor interoperability tests have shown excellent results, even with negligible performance degradation, a strong advantage towards disaggregation. Such pluggable technologies are expected to drive the removal of transponder modules as stand-alone elements, with clear benefits in cost, latency, and power consumption. Moreover, they will enable a tight integration with packet forwarding elements. First products encompassing both coherent DCO/ACO pluggable and Ethernet ASIC are already on the market. Much bigger impact is expected by the availability of even smaller form factor pluggable since they will be suitable for most commercially available high speed switching platforms. For example, ZR pluggable is rapidly emerging, supporting 400G Ethernet (multi-vendor) interconnections and, in the expected ZR+ version, also flexible bit rate adaptations according to reach and network conditions.

The removal of stand-alone transponders together with the adoption of pluggable coherent modules within bare metal switches (i.e., packet-optical white box, see Fig. 1) is expected to drive a complete redesign of optical networks, achieving tight integration with the packet switching layer.

However, the operating system of bare metal switches does not encompass yet the capability to control and monitor flexible optical transmission. For this reason, in the near future, significant activities will be carried out to enhance such node operating systems.

Unfortunately, also in this case, multiple parallel initiatives are in place, not facilitating rapid deployments.

Open Networking Linux (ONL) is a widely adopted basis for switching platforms, i.e., the underlying software that abstracts the hardware complexity. On top of it, Stratum could be adopted as an open source silicon-independent switch operating system designed for software defined networks, including novel protocols as P4Runtime and OpenConfig. Alternatively, typically on top of ONIE (an underlaying software abstraction), SONiC (Software for Open Networking in the Cloud) could be adopted. SONiC is a mature, complete operating system which relies on the switch abstraction interface (SAI) to interface the underlying hardware and supports traditional internet protocols (e.g., BGP). Designed by Microsoft, it is nowadays supported by most of the networking hardware vendors, and adopted by hyperscalers in their production cloud networks. However, SONiC can not be easily adapted on any hardware (sometimes SONiC over ONL over ONIE is considered in case of new HW) and it does not support new SDN solutions as P4Runtime. To this purpose, SONiC over Stratum over ONL over ONIE would be needed, but Stratum is not extremely mature and not widely supported yet. Clearly, the overall SW scenario for net operating systems requires further consolidation and progresses.

In this scenario, the first concrete steps to support pluggable optics on network operating systems are related to the definition of the transport application interface (TAI), i.e. a library that enables the pluggable parameters to be configured in a vendor-agnostic fashion, similarly as SAI controls vendor-agnostic Ethernet ASICs. TAI has been recently demonstrated on a containerized SONiC implementation over ONL [14].

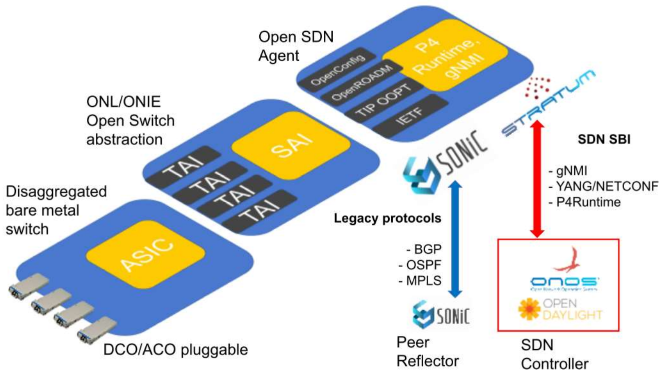

Figure 1: Packet-optical white box architecture and initiatives

IV. Future dircetions of Disaggregation

In the following, possible future directions in the design of disaggregated optical networks are discussed.

Open Network Operating System. As described in the previous section, relevant effort is in place on the implementation of open network operating systems for packet white boxes. So far, the optical community has focused on the standardization of YANG models. The next step, taking advantage of the above efforts, will be the implementation of open software for all network devices. Operators use network hardware from multiple vendors in different domains, but even a single vendor may supply different HW, firmware and software versions. In this scenario, the overall network management becomes complex. Moreover, adding new features to a traditional monolithic or proprietary software typically takes too long, due to the design, develop, test, and deploy phases. A common, structured, and open SW stack for network devices would remarkably simplify system integration, provide reliable management, and fasten new feature deployments on top of the underlying abstraction basis.

Southbound interfaces. The interface from Controllers to network devices, called southbound interface (SBI) is widely agreed to be based on YANG models and NETCONF protocol. While, as anticipated, different YANG models unfortunately exist to describe the optical network devices and architecture, NETCONF has emerged as the unique protocol adopted for optical device control. This significantly facilitates the adoption of disaggregated solutions. However, NETCONF does not represent the most efficient protocol technology, and alternative solutions are rapidly gaining consensus particularly at the IP level. This includes for example gNMI (gRPC Network Management Interface), an open source framework developed by Google and based on the remote procedure call (RPC) framework built on top of HTTP/2, which allows for a large variety of efficient options including unary, server streaming, client streaming bi-directional streaming RPCs, and multiplexing of RPCs over a single channel provided by libraries.

SDN Controller: Initially, OpenDayLight (ODL) gained significant consensus and it still represents the basis of several commercial SDN Controllers. However, its complexity has subsequently limited its further open development and the open source community has shifted effort towards ONOS, an open SDN Controller defined by the ODTN working group lead by the Open Networking Foundation (ONF). However, also in this case, progresses are rather slow and open solution are far from being ready for deployments in optical networks. Moreover, adding new functionalities typically brings in complexity and scalability issues. A radically new implementation approach seems to be required to enable remarkable progresses while guarantee scalability. One option could be the decoupling of intent operations and database management [15], avoiding monolithic SDN Controller implementations as of today. Database technologies nowadays provide ultra-high reliable and scalable performance and thanks to the adoption of YANG models, they could be implemented as stand-alone solutions able to directly collect data (e.g., through telemetry [9]), exchange network parameters (for reliability and information sharing), and even perform simple operations on the collected data. This way, the SDN Controller implementation would just implement queries to the DB and enforce intent-related operations.

Acknowledgment

This project has received funding from the ECSEL Joint Undertaking (JU) under grant agreement No 876967. The JU receives support from the European Union's H2020 research and innovation programme and from Italy Ministry of Education, University and Research (MIUR).

References

[1] E. Riccardi, P. Gunning, O. G. de Dios, M. Quagliotti, V. Lopez, and A. Lord, “An operator view on the introduction of white boxes into optical networks,” J. Light. Technol. 36, pp. 3062–3072, 2018.

[2] João Santos, Nelson Costa, and João Pedro, “On the Impact of Deploying Optical Transport Networks Using Disaggregated Line Systems”, J. Opt. Commun. Netw. 10(1) A60-A68, 2018

[3] C. Xie, et al, “Open and disaggregated optical transport networks for data center interconnects [Invited]”, J. Opt. Commun. Netw. 12(6), 2020

[4] F. Paolucci et al, “Disaggregated edge-enabled C+ L-band filterless metro networks”, J. of Optical Communications and Networking 12 (3), 2019

[5] E Riccardi, M Schiano, “Optical Node Disaggregation Management and Interoperability”, Optical Fiber Communication (OFC) Conf., 2020

[6] R. Casellas et al., “Abstraction and Control of Multi-Domain Disaggregated Optical Networks with OpenROADM Device Models”, Journal of Lightwave, 2020

[7] A Campanella et al., “ODTN: Open Disaggregated Transport Network. Discovery and control of a disaggregated optical network through open source software and open APIs”, OFC 2019

[8] A Giorgetti, et al. “Control of open and disaggregated transport networks using the Open Network Operating System (ONOS)”, Journal of Optical Communication. Netwowrk, 2019.

[9] F. Paolucci, A. Sgambelluri, F. Cugini, and P. Castoldi, “Network telemetry streaming services in SDN-based disaggregated optical networks”, J. Light. Technol. 36, pp. 3142–3149, 2018.

[10] M Dallaglio, N Sambo, F Cugini, P Castoldi, “Control and management of transponders with NETCONF and YANG”, JOCN 9 (3), 2017

[11] A. Sgambelluri, F. Paolucci, F. Cugini “Telemetry-driven Validation of Operational Modes in OpenConfig Disaggregated Networks”, ECOC ‘19

[12] D. Scano et al, “Operational Mode and Slicing Adaptation in OpenConfig Disaggregated Optical Networks”, OFC 2020

[13] https://www.lumentum.com/en/products/cfp2-dco-100g-200g-pluggable-transceiver

[14] V. Lopez et al., “Enabling fully programmable transponder white boxes [Invited]”, J. of Optical Comm. and Networking Vol. 12, Issue 2, 2020.

[15] F. Paolucci, F. Cugini, G. Cecchetti, and P. Castoldi, “Open Network Database for Application-Based Control in Multilayer Networks”, Journal of Lightwave Technology, 2017