

Peer-to-peer Telemetry in Disaggregated Optical Networks for Autonomic Transceiver Operation [Invited]

FRANCESCO PAOLUCCI^{1,*}, ANDREA SGAMBELLURI², ALESSANDRO PACINI², PIERO CASTOLDI², AND FILIPPO CUGINI¹

¹ CNIT, Pisa, Italy

² Scuola Superiore Sant'Anna, Pisa, Italy

* Corresponding author: francesco.paolucci@cnit.it

Compiled August 24, 2022

Autonomic networking and monitoring will drive the evolution of next generation Software Defined Networking (SDN) optical networks towards the zero touch networking paradigm. Optical telemetry services will play a key role to enable advanced network awareness at the device and component granularity. Optical disaggregation is the main driver of efficient and telemetry services thanks to vendor-neutral control and service management YANG models for optical systems and devices. In addition, open source micro services along with mature machine learning tools and platforms allow to elaborate huge amount of data streams from optical devices related to Quality of Transmission (QoT) parameters and transceivers digital signal processing key performance indicators. However, currently envisioned centralized telemetry collectors may pose scalability issues, limitations in the interactions with the SDN controller and suboptimal soft failure recovery due to operational mode limitations and the inability of tuning finer or proprietary transmission parameters, often conveniently achievable directly at the transceiver level. The paper proposes a novel Peer-to-peer telemetry (P2PT) service ready for local processing and soft failure recovery at the transceiver agent level. The P2PT architecture, workflow and subscription extensions are conceived to enable direct and fast recovery at the transceiver level resorting to optical signal retuning and adaptations. Experimental evaluations are provided in a multi-vendor disaggregated optical network testbed to assess different soft failure use cases and P2PT service scalability.

© 2022 Optical Society of America

<http://dx.doi.org/10.1364/ao.XX.XXXXX>

1. INTRODUCTION

Data deluge is characterizing the evolution of telecommunication networking [1]. Each next generation networking segment, technology and architecture will be designed with a special focus on full, open and online monitoring data extraction and analysis. This trend is evident in the optical networking ecosystem, as well. In fact, optical networks data, control and management systems are evolving in parallel with novel methods to achieve autonomic optical networking disclosed by accurate network awareness. Detailed optical performance monitoring is becoming a crucial aspect opening the way to the era of Zero Touch Networking [2].

One of the most significant and key instruments to achieve optical network awareness is the optical network telemetry service [3]. With respect to traditional monitoring platforms based

on closed vendor-locked alarming systems and traditional protocols such as Simple Network Management Protocol (SNMP), novel telemetry services have been conceived in the last years to support a programmable and versatile system for collecting optical key performance indicators. Advanced monitoring systems allow improving network control, failure and anomaly detection and to forecast incoming major faults in advance, thus minimizing critical network downs [4–7]. Telemetry service is expected to become a standard tool in the control/management plane domain, able to assist big data analytics and Artificial Intelligence (AI) platforms.

The main drivers of the advent and the feasibility of telemetry services in the optical networking ecosystem may be identified in three main parallel threads. First, the Elastic Optical Network (EON) disaggregation at the optical layer is pushing common

models to break the vendor lock-in and handle multi-vendor optical transceiver and pluggable cards, Reconfigurable Add Drop Multiplexers (ROADM) components such as Wavelength Selective Switches (WSS), amplifiers, optical degree stages, and Optical Line Systems (OLS) [8, 9]. Collaborative initiatives such as OpenConfig [10], OpenROADM [11], and the Telecom Infra Project (TIP) [12], are promoting open YANG models suitable for both control and management of multi-vendor white boxes and network components, including a unified method to set the configuration and retrieve the current device KPI status. Moreover, thanks to the standardization effort spent in the hardware and electronic driver design of optical devices [13], it becomes possible to directly access an increasing number of parameters, such as channel and aggregate optical power levels at optical amplifiers, spectra profile and in-band power at WSS, optical spectrum analyzer sensors, and data available from the digital signal processing units of optical coherent receivers (Digital Coherent Optics - DCO). The latter is extremely relevant since enables the disclosing of detailed levels of Optical Signal to Noise Ratio (OSNR) and pre Forwarding Error Correction Bit error Rate (pre-FEC BER) at the receiving cards. The second driver is the large availability of open source tools for data monitoring and processing. Such tools include serialization libraries (e.g., Google protobuf), Interface Definition Languages and Remote Procedure Calls frameworks (e.g., gRPC, Thrift), monitoring and collector systems (e.g., Prometheus, Kafka), time series databases (e.g., InfluxDB, Timescale) and graphical tools (e.g., Grafana). Such large availability is enabling the design and development of large trials and complex platforms deployable using such tools as blocks of microservices. The third driver is the availability of mature AI platforms and algorithms for big data analytics (e.g., Keras, PyTorch, TensorFlow), enabled by cloud/edge storage capabilities and dedicated computational capabilities (e.g., Graphic Processing Units - GPU).

This paper is divided into two main parts. In the first one, we review and envision the current state and efforts of optical telemetry services, including disaggregated architectures, workflows, implementation, and standardization initiatives. In particular, we highlight the possible limitations and drawbacks of current centralized-oriented telemetry services. In light of this, in the second part we propose a novel telemetry service framework, conceived to work in combination with centralized telemetry service. The novel service is designed to enable network awareness locally at the DCO cards agents, ready for next-generation pluggable DCO possibly equipped with lightweight AI chipsets. The proposed Peer-to-peer telemetry (P2PT) enables the collection of monitoring data to selected streaming data produced by a set of devices of interest for the optical transceiver agent (e.g., those crossed by the originating lightpath, or adjacent channels KPIs) and the data elaboration at the card for possible automatic transmission tuning optimization. The proposed scheme is particularly attractive for the detection of sporadic soft failures, optical signal deviations, filter shifts and amplifier gain degradations at the card. In this way, automatic failure recovery can be performed (e.g., automatic power or central frequency tuning) without the involvement of the SDN controller.

The paper is an extended version of the conference paper in [14]. This paper expands the related works on disaggregated optical telemetry, the centralized and proposed P2PT architectural solutions, details the P2PT service gRPC model and provides a novel set of results including additional soft failure use case applications and a detailed scalability study.

2. TELEMETRY IN DISAGGREGATED OPTICAL NETWORKS: STATE OF THE ART , ARCHITECTURES AND MODELS

According to the IETF draft on telemetry framework and RFC 7799, telemetry is defined as "any information that can be extracted from networks (including data plane, control plane, and management plane) and used to gain visibility or as basis for actions is considered telemetry data. It includes statistics, event records and logs, snapshots of state, configuration data, etc. It also covers the outputs of any active and passive measurements" [3, 15]. The document extends telemetry usage to control, management, forwarding and external data planes, with a focus on packet switched networks. In this work, we consider the optical data plane and optical telemetry services as the set of architectures, workflows and protocols enabling the augmented real-time and highly configurable monitoring of optical data sources originated by disaggregated optical devices in a Software-Defined networking (SDN) control. The telemetry service triggers a time series collection of multiple, optionally parallel, data source samples. Typically, the service relies on a streaming oriented connection, i.e., instead of resorting on polling-based or asynchronous notifications mechanisms, utilizes a streaming protocol to retrieve data with a given periodic sampling time. Typically, streaming data are collected and consumed by a centralized Analytics Handler (AH), possibly implementing machine learning tools to detect, localize and predict failures, alarms and anomalies. The service may be triggered by the SDN controller or by the management system in an automatic or manual fashion. Moreover, it may be activated permanently or using on-demand ticket with pre-defined service duration. Unsubscription events may also be defined to terminate a permanent streaming or stop in advance a telemetry with given service duration.

The main telemetry service instance, as defined by the OpenConfig and the OpenROADM models [10, 11], is referred to as *service subscription*, defining the type of service, the monitors to activate and the destination of the sampled data. In particular the subscription defines the following main elements:

- the subscription type, either *persistent* (configured locally and maintained across device restart/modifications) or *dynamic* (configured via RPC channel, not persisting across restart and channel reset or tear down);
- the subscription ID, unique identifier of the instance;
- the producer devices, expressed as sensor profiles, groups and paths associated to the source data to be retrieved, as disclosed by related SDN device agents;
- the collector devices, expressed as destination groups. If addresses are specified, the device initiates the streaming (dial-out session), otherwise the device waits for inbound connection referencing the subscription id (dial-in session);
- the heartbeat-config, the maximum time elapsing two subsequent telemetry messages;
- the stream-protocol-config, the type of streaming protocol (e.g., SSH, gRPC, JSON-RPC, Thrift-RPC, WebSocket-RPC);
- stream-encoding config, the wire format and the RPC framework (e.g., XML, IETF-based JSON, protobuf version3);

- stream-frequency-config, the device data sample interval (if set to zero, the sensor is asynchronous upon any data source value change).

A. State of the art: literature

A relevant research effort has been pushed in the last years to propose innovative monitoring and telemetry deployments, specifically designed for the disaggregated scenario and ready to support AI-based data analytics. The main research topics focused on efficient telemetry protocols, disaggregation flavour workflows, network agent extensions, innovative hardware drivers along with open interfaces, and high resolution sensors able to measure optical parameters with short sampling time.

The first activities related to optical telemetry have emerged from the concept of autonomic networking, where the monitoring plane has been conceived as an automatic framework participating at the observe-analyze and act loop [16], typical of the current zero-touch networking framework [2]. Optical autonomic networking [5] including spectra telemetry of optical spectrum analyzer resorting to the IPFIX protocol has been proposed to detect optical filter malfunctioning and type of failure (e.g., filter shift, filter tightening), including an extended work extension covering also multi-layer aspects [17].

One of the first works exploring the architectures, the workflows and the impact of telemetry protocols in a disaggregated scenario was presented in [6]. In particular, the work highlighted the potential of the gRPC protocol to provide programmable, lightweight and scalable telemetry service in comparison with the standard NETCONF protocol.

Focusing on the optical device agent, the work in [18] adopts a novel network operating system that includes, besides gRPC-based streaming telemetry, a dedicated threshold-based telemetry service suitable for specific network verification purposes.

The work in [19] proposes a telemetry-based workflow to test and select the most suitable transmission operational modes to optical card transmitters in a partially disaggregated scenario. The proposal aims at overcoming the issue of the opaque operational mode attribute in the OpenConfig model.

Highly reconfigurable telemetry systems have been proposed for disaster recovery exploiting the OpenConfig model [20]. The study has targeted the dynamic reconfiguration of control plane using wireless low-bandwidth network segments.

Fully open telemetry-enabled agents have been deployed recently. For example, the works in [21] and [22] envision open disaggregated ROADMs with filterless add/drop module and photodetector tap arrays. The co-located control agent exploits gRPC telemetry up to 20Hz sample update frequency.

Channel probing procedures for both narrowband and wideband connection configurations in the context of optical spectrum as a service have been detailed at the disaggregated network data plane level [23].

The application of machine learning to optical monitoring of disaggregated whiteboxes and telemetry has been covered in a number of recent works. In particular, [24] describes the monitoring workflows and the application of deep neural networks to infer diagnosis results from data observations retrieved by optical coherent transceivers and optical line systems, showing an offline and online evaluation for fiber bending event detection. Moreover, the recent work [25] proposes the adoption of machine learning techniques to enable soft failure detection and localization techniques resorting to partial telemetry, in which a disaggregated network exposes a number of limited shared monitors.

Significant improvements at the agent and device driver levels have been carried out recently to allow the remote monitoring of many parameters with effective open API and streaming protocols [26]. The work reported in [27] has deployed a power monitor blade capable of providing up to 400 μ s telemetry streaming period using direct memory access to be consumed by AI platforms. One of the most comprehensive open ROADM implementations has been extended to support telemetry streams with high-resolution data from internal WSS, capable of scanning the whole C-band transit spectrum at the resolution of 312.5 MHz using YANG Push streaming mode [13].

Concerning the optical line system, among the latest initiatives, the proposal of a specific OLS controller in charge of monitoring and proactively tuning the optical line parameters have been detailed and evaluated in [28].

In the multi-layer scenario, a coordinated per-layer telemetry system has been proposed and analyzed in [29]. The system merges optical telemetry, related to wavelength-switched disaggregated network, with in-band telemetry of packet-switched tributary traffic, performed by a programmable P4 switch.

The demonstration of the different API and protocols combining both control and monitoring functions have been detailed in [7], offering a careful performance evaluation of gRPC, gNMI, NETCONF, RESTCONF.

The relationship between disaggregated optical networks and existing monitoring tools (e.g., Kafka, Prometheus) utilized in different scenarios and production networks have been analyzed in the very recent works in [30] and in [31]. The former introduces and adapts the use of the Kafka framework in a disaggregated network to provide the telemetry service, while the latter utilizes the SONiC operating system to deploy a gNMI-based telemetry written in the Go language, in combination with the Prometheus open source monitoring platform.

B. Disaggregated telemetry service architectures

The telemetry service architecture is strictly related to the kind of disaggregation selected in the considered network scenario. Two main disaggregation flavours are hereafter considered, each one disclosing a telemetry service architectural configuration.

Partial disaggregation, promoted by the OpenConfig model, relies on the disaggregation of multi-vendor xPonder cards. The OLS is single vendor (including ROADMs and links), while multi-vendor transceivers may be hosted and activated in a pair fashion (i.e., the optical connection is assumed to be established between two cards belonging to the same vendor). This degree of disaggregation is considered a good trade-off for telco operators, since vendor lock-in is opened at the card level, while the OLS infrastructure is conveniently handled by a single vendor.

Full disaggregation, promoted by the OpenROADM model, opens multi-vendor coexistence also at the OLS level. This way, the ROADM becomes the central disaggregation point in which multiple multi-vendor optical devices are handled by device agents and may form hierarchical relationships at the ROADM node level. The underline model is more complex due to the strict interaction of each component at the same node level, however hierarchical structure may lead to distributed and offloaded control. For example, this model may disclose a ROADM controller, acting as intermediate control hierarchy layer between the central SDN controller and the SDN agent.

Depending on the selected disaggregation flavour, different telemetry service architectures may be conceived, as shown in Fig. 1, where the telemetry subscription and streaming phases are detailed. In partial disaggregation, the central SDN controller




Fig. 1. Telemetry subscriptions and streaming at disaggregated nodes/agents.

is the only element that may coordinate the subscriptions. Thus, all the different APIs are needed at the controller level, thus implying potential scalability issues due to the different nature and configuration of the single telemetry sources. To cope with this, telemetry subscriptions may require the adoption of data bundling feature (i.e., use the same streaming instance to collect different data source originated from the same device/agent) and data compression (e.g., use of implicit headers and field formats). Such option imply a slightly increased processing complexity at the telemetry agent. As an example, in Fig. 1 subscription S1 from SDN controller to the device agent triggers telemetry T1 to central collectors.

In full disaggregation, telemetry service offload is feasible. Such offloading implies a unique group subscription to a whitebox controller (WC) related to data sources belonging to different whitebox devices. Node controller is then in charge of activating each subscription to the different devices. Moreover, another interesting feature is that the whitebox controller may activate independent local telemetry. In parallel, collectors may be centralized (SDN controller level) or decentralized (node controller level). Such architectural degree of freedom improves service scalability at the controller. However, it may imply an increased number of connections and computational resources at the whitebox level. In the example of Fig. 1, SDN controller issues a subscription S2 to the WC. WC then issues subscription S3 retrieving telemetry T3 towards local collectors and AH. The result of T3 (either raw data or processed data) are sent to central AH using telemetry T2. Note that WC may issue independent subscriptions to monitor selected device agents (e.g., S3 and T3 may be the result of independent telemetry sessions at the WC).

Telemetry services realize the first step of the zero touch networking and in general of the autonomic networking framework loop (observe, analyze, act). The following steps involve the telemetry data analysis, typically performed at centralized Analytics Handlers (AH) and the feedback configurations performed by the SDN controller. The full loop is described in Fig. 2. The telemetry subscriptions are coordinated by the SDN controller in order to live monitor a set of optical parameters from a number of disaggregated devices related to one or more installed connections (i.e., lightpaths). In the figure, lightpath L_a is monitored thanks to telemetry subscriptions performed at the DCO receiver and intermediate ROADMs R2. In particular, the

subscription to the DCO agent includes the retrieval of OSNR, BER and received power P_{in} , while the subscription to R2 includes the power values P_{in} and P_{out} . Services may be deployed for both disaggregation flavours. In particular, R2 streaming may be activated resorting either to the OLS agent (in the partial disaggregated scenario) or to the R2 node controller (in the full disaggregation), responsible to trigger streaming instances to the selected whitebox devices (e.g., WSS or degree component). In addition, the subscription set may be enlarged to fully or partially co-routed adjacent lightpaths (e.g., lightpath L_b DCO is monitored as well). Streams are collected by AH and, in the case of value anomalies, soft failure detection and localization is performed and feedbacks are provided to the SDN controller, in charge of performing the most suitable recovery procedure (e.g., elastic operation, rerouting) including finer procedures at the transmitter, such as launch power tuning, signal shaping, central frequency detuning.

C. Telemetry models, protocols and encodings

The effectiveness of telemetry services resides in the accurate selection of the most suitable models, methods, protocols and data encoding to convey optical data plane samples in a telemetry instance session.

Telemetry models are converging towards the unified Open-Config model. This framework, defined in the device/system hierarchy of the model is mature and permits a high service configurability, offering streaming sessions, flexible definitions of collectors (groups), types of subscriptions, dial-in or dial-out triggering mechanisms, heartbeat, stream rate and encoding.

Two main streaming telemetry channels are emerging. The former, promoted by IETF, conveys telemetry streaming sessions within the YANG/NETCONF API, resorting to the YANG push stream options. The latter, promoted by Google, resorts to the unified gNMI management and control and resorts to gRPC. In addition, several implementations are using dedicated protocols to separate the control API from the management/telemetry channels. The most considered protocols are gRPC, IPFIX and Thrift. In particular, gRPC and Thrift are programmable to define the methods, the protocol fields and the encoding, along with specific features aiming to reduce the streaming rate. Indeed, the usage of implicit headers, data bundling and compression, which requires a limited additional complexity at the telemetry agent, makes such interfaces suitable for a scalable distributed micro stream system, such as the telemetry service.

Micro streams are typical of telemetry of optical parameters. If compared with traditional 15-minutes average statistics of last generation optical legacy systems, the availability of optical value samples at the second sampling period granularity is sufficient for most monitoring applications, in particular for lightpath health monitoring. Even though emerging open platforms are able to monitor optical parameters with a very high rate (e.g., 50ms telemetry update of power values have been demonstrated recently in an open whitebox implementation [13]), current digital signal processing units at the DCO are not able to sustain update rates of few Hz. Thus, the 1 second granularity is considered enough for the optical layer data plane. Such telemetry pace implies that a large amount of micro streams are easily handled by a single collector. This is generally not true for packet-switched telemetry (e.g. as in in-band network telemetry - INT), in which the telemetry streams rates scale with the rate of connections, implying significant scalability issues at the collector system [32].

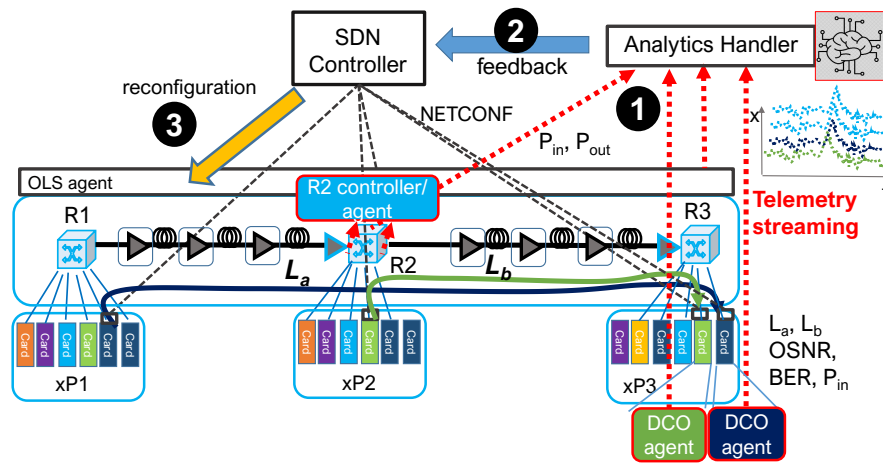


Fig. 2. Telemetry Services in disaggregated optical network and full observe-analyze-act loop.

3. PEER-TO-PEER TELEMETRY: MOTIVATIONS, WORKFLOWS AND MODEL IMPLEMENTATION

A. Motivations and technology trends

The scalability of centralized optical telemetry services does not necessarily imply the effectiveness in terms of reactivity of the whole telemetry-handler-controller closed loop when a disaggregated network is considered.

First, the communication between the Analytics Handler and the SDN controller lacks of design and standardization efforts. In particular, this feedback channel pace and granularity has not been defined nor addressed yet. This point is particularly critical since SDN controllers are not designed to handle high frequency notifications and re-optimization requests.

Second, the contribution of SDN controller re-computation, re-configurations in terms of reaction speed is not negligible. Such loop delays may be justified during major recovery or re-optimization operations, especially if resorting to make-before-break restorations or using temporary auxiliary backup channels (e.g., 1+1, 1:1 protection). In the case of soft failure these delays may lead to ineffective recovery procedures, especially if affecting the re-configuration change of ROADM degrees, WSS and other devices taking even minutes to become effective.

Third, SDN controller may be unaware of the proprietary DCO operational mode features. This aspect is controversial in the disaggregation framework. On the one hand, open models enable a vendor-neutral control and management template. This means that advanced proprietary transmission capabilities of a DCO can not be controlled by the SDN controller with a complete awareness. Operational modes have been defined in the Open-Config model to hide the specific optical transmission formats and attributes. The result is that the controller has a summarized view of the possible transmission modes, not being able to enforce a fine transmission tuning. As a consequence, the SDN controller recovery decisions may be sub-optimal, implying possible re-routing or adaptation of multiple lightpaths.

In light of the aforementioned aspects, since a number of soft failures may be recovered conveniently at a lower level, a mechanism enabling a local Quality-of-transmission aware fine tuning, also resorting to proprietary solutions not disclosed by open models, and soft failure recovery at the DCO transmitter could help to optimize the effectiveness of the recovery and the reaction speed. Such mechanism require a novel telemetry

service definition.

The recent advances in transmission technologies have pushed the development of pluggable coherent optical transceivers at rates of 400Gb/s with configurable transmission parameters such as modulation format and Forward Error Correction (FEC) [33]. For example, Digital Coherent Optics (DCO) transceivers currently relying on 7nm technology have shown the capability to cover up to 1500km at 400Gb/s using 16 Quadrature Amplitude Modulation (QAM) at 69 Gbaud in a 75 GHz-spaced system with probabilistic constellation shaping and soft-decision forward error correction [34]. In addition, excellent interoperability demonstrations have been already performed by DCO-CFP2 modules provided by different vendors and compliant with the OpenROADM multi source agreement (MSA) [35, 36]. Further improvements towards 800Gb/s and beyond are expected with the fast approaching 5nm technology. Furthermore, hybrid platforms are nowadays commercially available supporting, besides multiple high-speed pluggable modules, also packet forwarding capabilities (e.g., Cassini and Galileo platforms) or COM-Express module and x86 processor for advanced applications to run at the network node (e.g., Wedge platform). This trend is expected to continue and future network platforms are expected to natively encompass AI processors to enable local AI data elaborations. Such trend is evident in the IoT ecosystem, where gateways are already equipped with lightweight ARM and even AI chip modules dedicated for online data processing.

B. P2PT Architecture and Workflow

In light of these motivations, we propose the adoption of an auxiliary telemetry service, referred to as Peer-to-Peer Telemetry (P2PT) service. The service is conceived to enable a disaggregated DCO card to analyze the online health of its originating lightpath thanks to selected distributed monitors along the whole disaggregated network. In combination with centralized telemetry, P2PT is conceived to bring network awareness at the pluggable card. The main goal of the service is to break the rigid centralized telemetry hierarchy improving the monitoring loop efficiency and scalability. In addition, and most important, the goal is to enable the DCO agent to take independent optical transmission decisions based on the data analyzed via P2PT without the direct involvement of the SDN controller, at the same time committed within the SDN configuration mandate. Fig. 3

shows the modules involved in the DCO. The telemetry agent is responsible for providing the standard producer telemetry (i.e., data originated by the DCO, for example the DCO launch power level) to centralized collectors. Moreover, it is extended to receive the novel P2PT data originated at remote monitors. P2PT data are stored and analyzed locally using a dedicated processing module (the AI module in the figure). The processing module is aware of the monitors, the retrieved data, and runs algorithms to derive soft failure detection and localization. Based on the type of failure, the module selects and enforces the most suitable optical signal tuning/adaptation. The SDN agent is updated about modifications involving the control model status (e.g., the NETCONF database) or to notify higher control levels (i.e., node controller, SDN controller) about failure detection or unacceptable QoT level in the case local adaptations are not sufficient to recover the desired signal quality. As a general SDN applicability constraint, the local adaptation should not imply a frequency slot occupancy exceed with respect to the one assigned by the controller (e.g., elastic operations are not permitted [37]). In addition, specific frequency and power adaptation margins have to be considered to prevent cascaded issues to adjacent signals or to avoid ringing effects and power overshoot at the OLS amplifiers.




Fig. 3. Peer-to-peer telemetry at the next-generation transmitter DCO: modules and adaptation workflow.

Fig. 4 shows the P2PT service activation and features. The P2PT related to an active lightpath is initiated by the SDN controller. The controller is aware of all the installed lightpath attributes (i.e., end nodes, reserved spectrum, path, hops and crossed devices) and is able to select, in a permanent fashion or on-demand based on specific network status and events, a pool of remote monitors of interest related to the considered lightpath (i.e., source lightpath). The monitors are selected along the lightpath route, including the receiver DCO and transit monitors (e.g., OLS EDFA amplifiers, ROADM degrees, ROADM WSS). Moreover, the SDN controller may consider additional monitors pool, related to installed lightpaths for which a possible adaptation may induce interference and QoT degradation. For example, the receiver data of spectrum adjacent lightpaths, co-routed for significant network segments, are key data used to ensure that source adaptations do not affect adjacent channels or trigger data plane instabilities. Enabling P2PT, the SDN controller delegates the transmitter DCO agent to perform self adaptations over its source lightpath. Optionally, adaptation margins may be specified by the controller to allow adaptations within pre-defined ranges. Possible DCO self adaptations include fine central frequency tuning, proprietary FEC/constellation shaping

parameters, launch power variations.

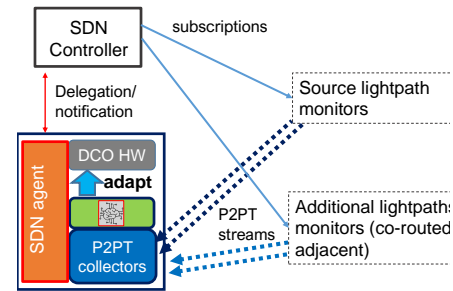


Fig. 4. P2PT service activation and streams to the DCO agent.

Fig. 5 shows the P2PT workflow with respect to the centralized telemetry service shown in Fig. 2. In step 1, the same monitors of Fig. 2 are directed towards the agent transmission side of the source DCO by means of dedicated P2PT subscriptions, including source lightpath L_a monitors and adjacent lightpath L_b monitor (i.e., OSNR and BER from green DCO receiver agent). Step 2 performs local signal adaptation in the case of soft failure detection, step 3 provides final notification to the SDN controller.

The P2PT service does not replace the traditional optical supervisory channel (OSC), where receiver and transmitter exploit a dedicated channel to exchange quality of signal parameters. Indeed, with respect to OSC, P2PT focuses on control plane operations, enabling a higher degree of configurability, flexibility and awareness. Moreover, P2PT breaks the rigidity of the Optical Transport Network (OTN) hierarchy, typically exploited by OSC, enabling the use of out-of-band control channels already available at each disaggregated node or device agent. Such choice simplifies the service deployment. However, a scalability evaluation needs to be assessed to estimate the potential P2PT traffic load impact in control channels.

C. Implementation

The proposed P2PT has been implemented as a parallel gRPC channel with respect to existing control API. This way the device agent YANG models (e.g., OpenConfig, OpenROADM) are not extended, thus not requiring any modification of the agent at the control layer. Future model extensions may take advantage of the information carried out by the specific P2PT service subscription. The gRPC protocol has been extended easily to provide key stream information related to the source, the type of sampled data and other specific information used by the DCO agents to collect and analyze data. The P2PT gRPC subscription request, performed by the SDN controller to the involved monitor agents, besides the standard parameters defined in the OpenConfig model and described in Sec. 2 defines the following novel parameters:

- the *template_id*, identifying the optical device source (e.g., DCO, WSS, ROADM, OLS EDFA);
- the *correlation*, identifying the relationship of the monitored data with respect to the transmitter DCO signal;
- the *positioning*, identifying the sorted hop positioning of the monitor device with respect to the monitored lightpath;

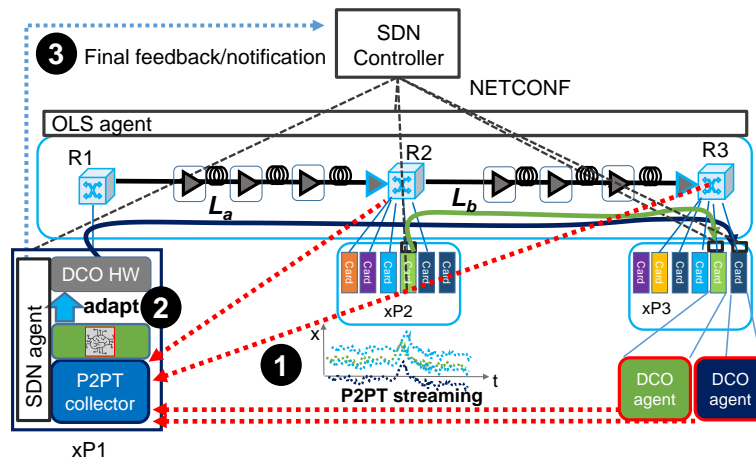


Fig. 5. P2PT service workflow in disaggregated optical network and local observe-analyze-act loop.

- the *resources*, the actual streamed data sample, identified by the specific model path of the considered device.

macro:

macro:

In the gRPC telemetry scheme, the structure of the messages is defined using protocol buffer language. Fig. 6 shows the data model, extracted from the protobuf file, of the telemetry subscription request. In particular, the subscription request includes all the required fields to activate the data stream (i.e., collector IP address, collector port, the list of resources to be collected, encoding, period, duration, template_id, correlation, positioning). The telemetry subscription request is sent to each traversed device by the SDN controller. In fact, the SDN controller is the only element that knows all the required details about the path and the devices ordering. The parameters included in the telemetry subscription request are then used by the telemetry server to tag the data, during the streaming process.

```

message SubscriptionRequest {
  string observation_point = 1;
  bool suppression = 2;
  uint32 interval = 3;
  uint32 duration = 4;
  uint32 template_id = 5;
  uint64 subscription_id = 6;
  string correlation = 7;
  uint32 positioning = 8;
  repeated Collector collectors = 9;
  repeated Resource resources = 10;
}

```

Fig. 6. P2PT: main gRPC implementation parameters.

The example of Fig. 7 explains the novel parameters included in the P2PT subscriptions, that are mapped as gRPC fields of the P2PT stream packet. The template-id identifies the disaggregated device sourcing optical data samples. Such info is key since soft failure recovery may be different based on the affected device. For example, a WSS failure may be recovered using narrow filtering techniques (e.g., signal re-shaping or fine central frequency retuning), while EDFA failures are recovered conveniently using launch power or soft FEC adaptations). The correlation parameter, with possible values *self*, *adjacent*, *aggregate* specifies whether monitored data are referred to the source lightpath in-band channel (e.g., as for receiving DCO and in-channel monitor inside disaggregated degrees or in the case

of WSS in-channel power monitors P_{in}^c and P_{out}^c), to adjacent in-band channels (e.g. as for adjacent lightpaths DCO receiver monitors), or to aggregate line (e.g., as for OLS EDFA, monitoring the input power levels P_{in} of the whole OLS line). The positioning info is used to identify the sorted list of devices crossed by the monitored lightpath. This info is crucial for failure localization algorithms and for optimized recovery policies. For example, a failure affecting the first lightpath OLS EDFA span may be easily recovered using DCO power launch adaptation. Finally, the resources are the set of time series data under monitoring and are identified through the path referred to the specific open source model.

The described gRPC subscription extensions do not require additional complexity at the monitor agents (i.e., improved control and topology awareness). All the subscriptions are handled by the controller, fully aware of device types and inventory, relative position and signal correlation with respect to the monitored lightpath. Subscriptions model is designed to allow dynamic P2PT activation. To this purpose, the positioning id is selected by the controller with independent policies and is not necessarily a set of subsequent and consecutive id, in order to enable time-dynamic subscriptions of additional devices and allow to detect the relative ordered position of each device. For this reason, the implementation assumes a max default value (i.e., 1000) for each receiver DCO monitor as the last monitors in the P2PT lightpath chain. Finally, all the mentioned subscription parameters are conveyed in the gRPC stream message, including the optical data samples and the timestamp, so that lightweight implementations at the DCO Telemetry Agent are feasible, resorting to gRPC servers and time series databases.

4. P2PT PERFORMANCE EVALUATION

The proposed P2PT implementation has been deployed and evaluated in a real disaggregated and multi-vendor optical network testbed shown in Fig. 8. The testbed includes both data and control plane.

The data plane reproduces partially the scenario in Fig. 5 and is composed by a pair of Ericsson SPO 1400 100G xPonders (exploiting Polarization Multiplexed Quadrature Phase Shift Keying - PM-QPSK modulation format), acting as source and receiver end points respectively, 2 OLS EDFA amplifiers connected to 80km fiber spans and an intermediate ROADM, emulated us-

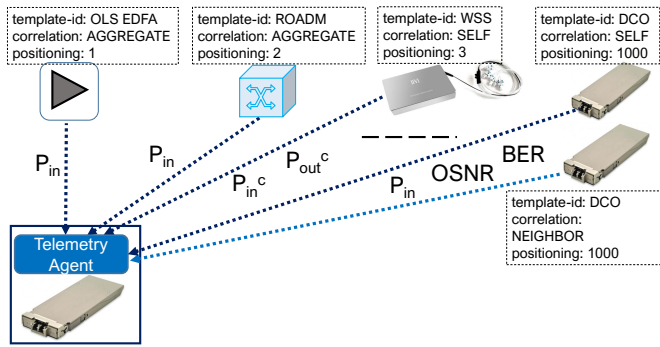


Fig. 7. P2PT: main gRPC implementation parameters.

ing two Lumentum Greybox ROADM-20 degrees. The SDN control plane is realized in a 10Gigabit Ethernet network. A central Mellanox SN2010 switch connects the SDN controller (Open Network Operating System - ONOS with Open and Disaggregated Transport Networks - ODTN module release 2.2, based on the implementation carried out in our works in [38] and in [19]) with the data plane device agents. Since the considered commercial optical devices (i.e., Ericsson SPO1400) are not natively NETCONF-based, in order to fully exploit the SDN paradigm, we adopted ad-hoc defined SDN agents able to access the proprietary APIs exposed by the devices (i.e., HTTP based) for configuration and monitoring purposes. The considered SDN agents run as containerized apps in a co-located Linux-based PC capable of access the control/management interface of the devices.

The SDN agents include a NETCONF agent module, based on the ConfD tool [39] and a gRPC telemetry agent module. Moreover, the transmitter agent A1 runs additional software modules: a time series database module, based on InfluxDB, and a graphic user interface dashboard, based on the Grafana tool. Finally, A1 includes a lightweight python-based analysis and soft failure detection module. The module implements a detection scheme exploiting, for each monitored P2PT parameter, a short term window-based differential strategy. In addition, for the sake of stability, it considers configurable thresholds to identify the values associated with the failure and react applying the appropriate re-configuration using a look up table. More specifically, it is able to filter and correlate the data, running a detection algorithm based on a three samples window and thresholds, in order to identify the presence of a failure, recognize its type and, optionally, perform the proper reaction.

A. P2PT-driven soft failure adaptation tests

Two different soft failures have been tested using P2PT monitoring: a filter shift soft failure at the Lumentum node (test 1) and a reduced power launch at the transmitter card (test 2). The two tests have been selected to describe the detection of different physical soft failures occurring at different BER devices in the lightpath chain, along with the most appropriate recovery adaptations driven by the different P2PT data set.

In both tests a lightpath intent has been configured by the ONOS along the OLS line including the Lumentum degrees transit node. The lightpath is configured with 193.7 THz central frequency, 37.5 GHz channel width (corresponding to three allocated 12.5GHz frequency slots in the flexible DWDM grid) and FEC1. After provisioning, the ONOS activates P2PT subscriptions to the Lumentum agent A3, the OLS EDFA agent A2/A4

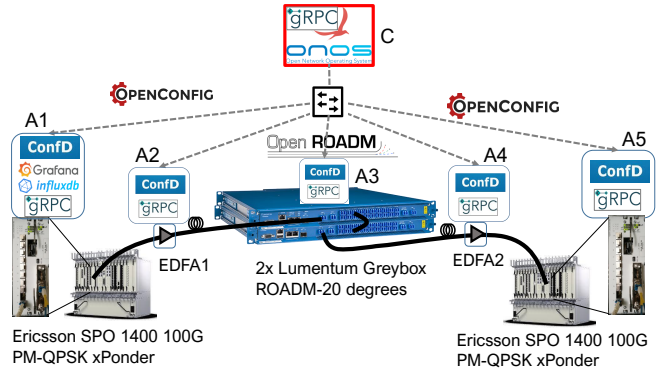


Fig. 8. Disaggregated optical network testbed.

and the SPO receiver A5 having A1 as collector. The selected monitors are the in-channel input and output power (before amplifier stage) centered at 193.7 THz at A3 (1s sampling time), the input aggregate power at A2/A4 and the estimated OSNR and BER values at A5 (5s sampling time due to BER computation at the card). Active lightpath steady state provided by P2PT monitors at A1 indicate 22.3dBm OSNR and 3×10^{-6} BER.

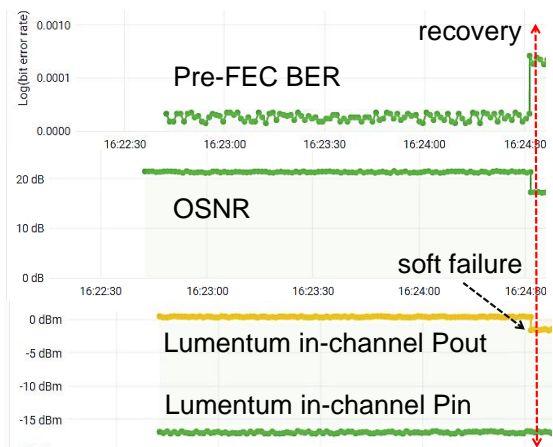


Fig. 9. Test 1 (filter shift): Grafana P2PT monitors at A1 agent.

In test 1 (Fig. 9), the ROADM filter soft failure is emulated by inserting a WSS between the two degrees. The WSS has been configured as a 37.5GHz detuned filter with 5GHz shift central frequency offset. The filter activation induces a degradation of the monitored A3 in-channel output power (around -1.5 dB), impacting A5 OSNR (around -2dB) with a BER degradation (1.5×10^{-4}). Agent A1 detects the anomalous pattern at A3. Since P2PT upstream monitors at A3 (in-channel input power) continue to stream steady state values, agent A1 detects a failure at the ROADM related to filtering issues. Agent A1, using the python script, after three A3 degraded samples (3s), triggers signal adaptation to FEC2, more robust to narrow filtering. FEC2 was not originally selected at the first provisioning since it requires higher processing complexity compared to FEC1. In the case of centralized telemetry, detailed results conducted on the same network testbed devices are reported in [19]. In particular, results report the case of hard link failure notification, with consequent trigger towards the SDN controller to re-route the existing lightpath along with the IP tributary switching operation and the related flow entry configurations. In light of

this, the only transit from the SDN controller implies that the same procedure would have taken at least 3s more due to the controller path re-computation and API configuration involving NETCONF messages [19]. Moreover, due to internal ONOS controller timers, NETCONF-based flow entries used for optical layer agents require additional acknowledgement time, thus incrementing the full rerouting procedure up to 30 seconds more. Since flow entries are not enforced during P2PT-based adaptation, the P2PT gain in terms of reactivity time is up to ten times faster. The effectiveness of P2PT will further improve when current limitations of commercially available transponders will be overcome, enabling quasi-hitless operational mode adaptation (see for example the hitless baud-rate adaptation of [40]).

In test 2 (Fig. 10), reduced power launch failure have been emulated by inserting a WSS at the output of the card and configuring flat 3 dB attenuation in the three channel slots. The WSS attenuation induces a degradation of both in-channel input and output power at Lumentum, impacting receiver OSNR and BER accordingly (19.5dB, 4×10^{-5}). This means that the failure is localized at the first link, possibly including the source itself. As a consequence, A1 agent triggers a 1.5 dB increase of the signal launch power (i.e., moving the TX power from 0dBm to +1.5dBm). The recovery is shown in the P2PT monitors as a function of the time, OSNR and BER are recovered to 20.3dB and 1.5×10^{-5} . The whole recovery procedure takes around 16s, including the 3 degraded samples check (15s), configured in the Python script to avoid data plane instability. The recovery time may be reduced up to 4s excluding BER telemetry and using OSNR 1s sampling time.

Fig. 11 shows the capture of gRPC streams received at A1 collector (IP address 10.30.2.37) in test 2. The capture shows the P2PT gRPC streams and the REST command (HTTP packet 3626) used to perform the power launch adaptation to the transmitter hardware driver (management IP 10.30.2.24). Each agent gRPC stream rate is around 1.6kbit/s average and around 3kbit/s peak in the producer-consumer direction. Fig. 12 details the P2PT gRPC fields of A5 stream, with `template_id` set to `card` (i.e., 3), correlation set to `self`, positioning set to default receiver (i.e., 10000), and key-value fields with the BER and OSNR sample values. The A3 stream (not shown) includes `template_id` set to `ROADM degree` (i.e., 2), positioning set to 10, `self` correlation and key-value fields with the two Lumentum power levels.

B. Scalability assessment

To evaluate the impact of P2PT stream in the control plane network and processing load, at monitor producers and transceiver collectors, a scalability assessment with a set of measurements as a function of the number of generated P2PT stream subscriptions has been performed.

Focusing on the resulting control network throughput, the scenario with one telemetry collector and multiple docker-based gRPC telemetry servers has been considered. The worst case configuration has been evaluated, implementing the gRPC encoding with no compression option [6]. Table 1 summarizes the obtained results at the varying of the number of triggered gRPC server (i.e., from 1 to 20 servers). The network throughput at the telemetry collector grows linearly with the number of telemetry streams to be received, presenting around 2.8Kb/s per stream. This result meets the requirements of a general limit condition, with a collector receiving up to 20 telemetry streams related to the same lighthouse (i.e., type self) and the ones on the adjacent optical channels (i.e., type other). Scaling the obtained results to a 10G control plane network, up to few millions of telemetry

#Telemetry Streams	Network Throughput
1	2.832Kb/s
5	13Kb/s
10	27Kb/s
15	35Kb/s
20	44Kb/s

Table 1. Network scalability tests at the varying of the number of active telemetry streams.

streams, distributed among the different telemetry collectors, can be sustained by the proposed system.

Focusing on the CPU load at the gRPC server, a scenario with one collector and one docker-based gRPC server has been considered. The gRPC server run in a Linux server equipped with an Intel Core(TM)2 Quad CPU Q8200 @2.33GHz and 6GB of RAM. Table 2 summarizes the obtained results at the varying of the number of gRPC streams generated by the server (i.e., from 1 to 200 streams). Also in this case, the CPU load grows linearly with the number of telemetry streams to be generated, remaining below the 10% of occupancy also in the case of 200 streams generated by the gRPC server. This result shows the efficiency of the gRPC protocol for the data streaming, presenting a limited impact on a node also in the case of 200 telemetry streams to be served.

#Telemetry Streams	CPU%
1	0.08
10	0.52
50	2.3
100	4.7
200	9.6

Table 2. CPU scalability tests at the gRPC server.

5. CONCLUSION

This paper surveyed the state-of-the art of SDN disaggregated optical network telemetry, including architectures, standard open models, protocols and standardization efforts to support highly programmable data stream of optical transmission parameters at the device level.

The peer-to-peer telemetry service was proposed in parallel with existing centralized monitoring to improve soft failure recovery reaction and effectiveness at the transceiver level, ready for next generation whitebox hosting optical pluggables and local AI platforms.

Experimental results showed that P2PT is effective to enable different failure detection/localization events and to speed up local lighthouse QoT recovery. The service is scalable due to the efficient gRPC encoding, sustaining few millions of active streams in the 10G control channel with a limited impact in the SDN agent processing load.

Further benefits are expected when the limitations affecting current generation of commercially available transponders




Fig. 10. Test 2 (TX power reduction): Grafana P2PT monitors at A1 agent, showing failure and recovery events.

No.	Time	Source	Destination	Protocol	Length	Info
3376	136.684330706	10.10.255.202	10.30.2.37	GRPC	155	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3389	137.709313748	10.10.255.202	10.30.2.37	GRPC	156	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3391	137.743813264	10.10.255.201	10.30.2.37	GRPC	143	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3410	136.748662385	10.10.255.202	10.30.2.37	GRPC	155	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3451	139.673647842	10.10.255.202	10.30.2.37	GRPC	155	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3467	140.698750524	10.10.255.202	10.30.2.37	GRPC	156	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3482	141.733685604	10.10.255.202	10.30.2.37	GRPC	156	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3493	142.254214490	10.10.255.201	10.30.2.37	GRPC	143	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3510	142.789852711	10.10.255.202	10.30.2.37	GRPC	155	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3535	143.698437466	10.10.255.202	10.30.2.37	GRPC	155	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3573	144.724415352	10.10.255.202	10.30.2.37	GRPC	156	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3591	145.757512982	10.10.255.202	10.30.2.37	GRPC	155	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3610	146.726495810	10.10.255.202	10.30.2.37	GRPC	155	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3614	146.846114749	10.10.255.201	10.30.2.37	GRPC	144	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3623	147.015570412	10.30.2.24	10.30.2.24	TCP	74	35731 → 5080 [SYN] Seq=0 Win=2920 Len=0 MSS=1460
3624	147.015474700	10.30.2.24	10.30.2.24	TCP	66	5080 → 35731 [SYN, ACK] Seq=0 Ack=1 Min=65535 Len=0
3625	147.015509713	10.30.2.24	10.30.2.24	TCP	54	35731 → 5080 [ACK] Seq=1 Ack=1 Min=29312 Len=0
3626	147.015554713	10.30.2.37	10.30.2.24	HTTP	258	PUT /Configuration/SPO2/AmplGain/Fixed/4/23 HTTP/1
3627	147.062408536	10.30.2.24	10.30.2.37	TCP	60	5080 → 35731 [ACK] Seq=1 Ack=205 Min=1051136 Len=0
3645	147.760619691	10.10.255.202	10.30.2.37	GRPC	155	DATA[1] (GRPC) (PROTOBUF) Telemetry.TelemetryData
3652	148.044879168	10.30.2.24	10.30.2.37	TCP	71	5080 → 35731 [PSH, ACK] Seq=1 Ack=205 Min=1051136 Len=0
3653	148.044111332	10.30.2.37	10.30.2.24	TCP	54	35731 → 5080 [ACK] Seq=205 Ack=18 Min=29312 Len=0
3656	148.044344779	10.30.2.24	10.30.2.37	HTTP/1	196	HTTP/1.0 200 OK, JavaScript Object Notation (appli

Fig. 11. Test 2: Wireshark capture at transmitter agent A1.

will be overcome, e.g. enabling quasi-hitless operational mode adaptation while fully supporting open models and transparent description of hardware capabilities.

Future works will investigate additional use cases where the proposed P2PT technology is expected to provide significant benefits, for example in the case of super-channels supporting control loops providing joint optimization of power level and sub-channel spacing for every sub-channel. Indeed, such optimization can not be efficiently performed by other devices (e.g., the first traversed ROADM) or by the SDN Controller.

ACKNOWLEDGMENTS

This work received funding from the ECSEL JU project BRAINE (grant agreement No 876967). The JU receives support from the EU Horizon 2020 research and innovation programme and the Italian Ministry of Education, University, and Research (MIUR).

REFERENCES

1. K. Slavakis, G. B. Giannakis, and G. Mateos, "Modeling and optimization for big data analytics: (statistical) learning tools for our era of data deluge," *IEEE Signal Process. Mag.* **31**, 18–31 (2014).
2. ETSI, "Zero-touch network and service management (zsm); reference architecture," ETSI GS ZSM 002 V1.1.1 (2019-08) (2019).
3. H. Song, F. Qin, P. Martinez-Julia, L. Ciavaglia, and A. Wang, "Network telemetry framework," IETF draft-ietf-opsawg-ntf-07 (2021).
4. B. Shariati, M. Ruiz, J. Comellas, and L. Velasco, "Learning from the optical spectrum: Failure detection and identification," *J. Light. Technol.* **37**, 433–440 (2019).
5. L. Velasco, A. Sgambelluri, R. Casellas, L. Gifre, J. Izquierdo-Zaragoza, F. Fresi, F. Paolucci, R. Martínez, and E. Riccardi, "Building autonomous optical whitebox-based networks," *J. Light. Technol.* **36**, 3097–3104 (2018).

```

Transmission Control Protocol, Src Port: 52318, Dst Port: 50083, Seq: 763, Ack: 228,
HyperText Transfer Protocol 2
GRPC Message: /Telemetry.OCReply/StreamData, Request
Protocol Buffers: /Telemetry.OCReply/StreamData,request
  Message: Telemetry.TelemetryData
    [Message Name: Telemetry.TelemetryData]
    > Field(1): timestamp = 1626972714809 (uint64)
    > Field(2): observation_point = card (string)
    > Field(3): template_id = 3 (uint32)
    > Field(4): sequence_number = 5 (uint32)
    > Field(5): subscription_id = 11 (uint64)
    > Field(6): correlation = 1 (uint32)
    > Field(7): positioning = 10000 (uint32)
    > Field(8): kv (message)
      kv: (22 bytes)
        Message: Telemetry.KeyValue
          [Message Name: Telemetry.KeyValue]
          > Field(1): key = spo_BER (string)
          > Field(3): str_value = 1.60386e-05 (string)
        Message: Telemetry.KeyValue
          [Message Name: Telemetry.KeyValue]
          > Field(1): key = spo_OSNR (string)
          > Field(3): str_value = 20.3 (string)
  
```

Fig. 12. P2PT gRPC stream fields related to receiver monitor.

6. F. Paolucci, A. Sgambelluri, F. Cugini, and P. Castoldi, "Network telemetry streaming services in SDN-based disaggregated optical networks," *J. Light. Technol.* **36**, 3142–3149 (2018).
7. R. Vilalta, C. Manso, N. Yoshikane, R. Casellas, R. Martínez, T. Tsuritani, I. Morita, and R. M. noz, "Experimental evaluation of control and monitoring protocols for optical SDN networks and equipment [Invited Tutorial]," *J. Opt. Commun. Netw.* **13**, D1–D12 (2021).
8. E. Riccardi, P. Gunning, O. G. de Dios, M. Quagliotti, V. López, and A. Lord, "An operator view on the introduction of white boxes into optical networks," *J. Light. Technol.* **36**, 3062–3072 (2018).
9. J. A. Hernandez, M. Quagliotti, E. Riccardi, V. Lopez, O. G. d. Dios, and R. Casellas, "A techno-economic study of optical network disaggregation employing open source software business models for metropolitan area networks," *IEEE Commun. Mag.* **58**, 40–46 (2020).
10. "OpenConfig," <http://openconfig.net>.
11. "Open ROADM Multi Source Agreement," <http://www.openroadm.org>.
12. "Telecom Infra Project," <https://telecominfraproject.com>.
13. J. Kandrát, M. Vaško, R. Krejčí, V. Kubernát, T. Pecka, O. Havliš, M. Šlapák, J. Jedlinský, and J. Vojtěch, "Opening up ROADMs: streaming telemetry [Invited]," *J. Opt. Commun. Netw.* **13**, E81–E93 (2021).
14. F. Paolucci, A. Sgambelluri, P. Castoldi, and F. Cugini, "Telemetry solutions in disaggregated optical networks: an experimental view," in *Optical Fiber Communications Conference (OFC)*, (2021), pp. 1–3.
15. A. Morton, "Active and passive metrics and methods (with hybrid types in-between)," IETF RFC 7799 (May 2016).
16. L. Velasco, A. C. Piat, O. Gonzalez, A. Lord, A. Napoli, P. Layec, D. Rafique, A. D'Errico, D. King, M. Ruiz, F. Cugini, and R. Casellas, "Monitoring and data analytics for optical networking: Benefits,

- architectures, and use cases," *IEEE Netw.* **33**, 100–108 (2019).
17. L. Gifre, J. Izquierdo-Zaragoza, M. Ruiz, and L. Velasco, "Autonomic disaggregated multilayer networking," *IEEE/OSA J. Opt. Commun. Netw.* **10**, 482–492 (2018).
 18. A. Sadasivarao, S. Syed, D. Panda, P. Gomes, R. Rao, J. Buset, L. Paraschis, J. Brar, and K. Raj, "Demonstration of extensible threshold-based streaming telemetry for open dwdm analytics and verification," in *Optical Fiber Communications Conference (OFC)*, (2020), pp. 1–3.
 19. A. Sgambelluri, A. Giorgetti, D. Scano, F. Cugini, and F. Paolucci, "OpenConfig and OpenROADM automation of operational modes in disaggregated optical networks," *IEEE Access* **8**, 190094–190107 (2020).
 20. S. Xu, Y. Hirota, M. Shiraiwa, M. Tornatore, S. Ferdousi, Y. Awaji, N. Wada, and B. Mukherjee, "Emergency opm recreation and telemetry for disaster recovery in optical networks," *J. Light. Technol.* **38**, 2656–2668 (2020).
 21. J. Kundrat, O. Havlis, J. Jedlinsky, and J. Vojtech, "Opening up roadms: Let us build a disaggregated open optical line system," *J. Light. Technol.* **37**, 4041–4051 (2019).
 22. J. Kundrat, O. Havlis, J. Radil, J. Jedlinsky, and J. Vojtech, "Opening up roadms: a filterless add/drop module for coherent-detection signals," *IEEE/OSA J. Opt. Commun. Netw.* **12**, C41–C49 (2020).
 23. K. Kaeval, T. Fehenberger, J. Zou, S. L. Jansen, K. Grobe, H. Griesser, J.-P. Elbers, M. Tikas, and G. Jervan, "QoT assessment of the optical spectrum as a service in disaggregated network scenarios," *J. Opt. Commun. Netw.* **13**, E1–E12 (2021).
 24. T. Tanaka, T. Inui, S. Kawai, S. Kuwabara, and H. Nishizawa, "Monitoring and diagnostic technologies using deep neural networks for predictive optical network maintenance [invited]," *J. Opt. Commun. Netw.* **13**, E13–E22 (2021).
 25. K. S. Mayer, J. A. Soares, R. P. Pinto, C. E. Rothenberg, D. S. Arantes, and D. A. A. Mello, "Machine-learning-based soft-failure localization with partial software-defined networking telemetry," *J. Opt. Commun. Netw.* **13**, E122–E131 (2021).
 26. A. Sadasivarao, S. Jain, S. Syed, K. Pithewan, P. Kantak, B. Lu, and L. Paraschis, "High performance streaming telemetry in optical transport networks," in *2018 Optical Fiber Communications Conference and Exposition (OFC)*, (2018), pp. 1–3.
 27. K. Ishii, S. Yanagimachi, A. Tajima, and S. Namiki, "Submillisecond control/monitoring of disaggregated optical node through a direct memory access based architecture," in *2019 Optical Fiber Communications Conference and Exhibition (OFC)*, (2019), pp. 1–3.
 28. G. Borraccini, A. D'Amico, S. Straullu, A. Nespola, S. Piciaccia, A. Tanzi, G. Galimberti, S. Bottacchi, S. Swail, and V. Curri, "Cognitive and autonomous QoT-driven optical line controller," *J. Opt. Commun. Netw.* **13**, E23–E31 (2021).
 29. A. Sgambelluri, F. Paolucci, A. Giorgetti, D. Scano, and F. Cugini, "Exploiting telemetry in multi-layer networks," in *2020 22nd International Conference on Transparent Optical Networks (ICTON)*, (2020), pp. 1–4.
 30. A. Sgambelluri, A. Pacini, F. Paolucci, P. Castoldi, and L. Valcarenghi, "Reliable and scalable kafka-based framework for optical network telemetry," *IEEE/OSA J. Opt. Commun. Netw.* **13**, E42–E52 (2021).
 31. A. Sadasivarao, Y. Yin, K. Balasubramanian, R. Morgan, P. Gomes, S. Syed, and L. Paraschis, "Demonstration of extensible disaggregated ne management based on a sonic-compliant software-agent optical openconfig/gnmi telemetry go implementation," in *2021 Optical Fiber Communications Conference and Exhibition (OFC)*, (2021), pp. 1–3.
 32. I. Pelle, F. Paolucci, B. Sonkoly, and F. Cugini, "Latency-sensitive edge/cloud serverless dynamic deployment over telemetry-based packet-optical network," *IEEE J. on Sel. Areas Commun.* **39**, 2849–2863 (2021).
 33. J. Pedro and S. Pato, "High-performance versus power-efficient coherent optical interfaces: Spectral efficiency and hardware count comparison," in *2020 22nd International Conference on Transparent Optical Networks (ICTON)*, (IEEE, 2020), pp. 1–6.
 34. <https://www.neophotonics.com/press-releases/?newsId=12486> (2021). [Online; accessed 15-July-2021].
 35. E. Pincemin, Y. Loussouarn, Y. Pan, G. Miller, A. Gibbemeyer, B. Mikkelsen, A. Gaibazzi, W. Way, T. Yamazaki, A. Hayashi *et al.*, "Interoperable cfp-dco and cfp2-dco pluggable optic interfaces for 100g wdm transmission," in *2019 Optical Fiber Communications Conference and Exhibition (OFC)*, (IEEE, 2019), pp. 1–3.
 36. M. Newland, R. Schmogrow, M. Cantono, V. Vusirikala, and T. Hofmeister, "Open optical communication systems at a hyperscale operator," *J. Opt. Commun. Netw.* **12**, C50–C57 (2020).
 37. F. Paolucci, A. Castro, F. Fresi, M. Imran, A. Giorgetti, B. B. Bhowmik, G. Berrettini, G. Meloni, F. Cugini *et al.*, "Active PCE demonstration performing elastic operations and hitless defragmentation in flexible grid optical networks," *Photonic Netw. Commun.* **29**, 57–66 (2015).
 38. A. Sgambelluri, J. . Izquierdo-Zaragoza, A. Giorgetti, L. Gifre, L. Velasco, F. Paolucci, N. Sambo, F. Fresi, P. Castoldi, A. C. Piat, R. Morro, E. Riccardi, A. D'Errico, and F. Cugini, "Fully disaggregated roadm white box with netconf/yang control, telemetry, and machine learning-based monitoring," in *2018 Optical Fiber Communications Conference and Exposition (OFC)*, (2018), pp. 1–3.
 39. A. Sgambelluri, A. Giorgetti, F. Paolucci, P. Castoldi, and F. Cugini, "Open source implementation of OpenConfig telemetry-enabled NETCONF agent," in *ICTON 2019; International Conference on Transparent Optical Networks*, (2019), pp. 1–4.
 40. A. Dupas, P. Layec, E. Dutisseuil, S. Belotti, S. Bigo, E. H. Salas, G. Zervas, and D. Simeonidou, "Elastic optical interface with variable baudrate: Architecture and proof-of-concept," *J. Opt. Commun. Netw.* **9**, A170–A175 (2017).