

Network Telemetry Streaming Services in SDN-based Disaggregated Optical Networks

F. Paolucci, A. Sgambelluri, F. Cugini and P. Castoldi

Abstract—Accurate real-time availability of transmission parameters at the network controller has the potential to significantly improve the efficiency of control and management operations, particularly in the case of soft failures where detection and localization procedures are typically affected by the long monitoring time intervals usually adopted in current Elastic Optical Networks (EONs).

In this paper, we report on the design, implementation and experimental demonstration of a telemetry service exploiting the gRPC protocol to enable on-demand streaming of real-time monitoring parameters, dynamically retrieved from a configurable set of network devices. The telemetry service is efficiently introduced for Software Defined Networking in EONs, also accounting for disaggregated network elements through standard YANG-defined network models.

The implemented telemetry service is experimentally validated for partial and fully disaggregated architectures over networking scenarios where soft failures not addressable through traditional monitoring solutions are successfully detected and localized. Moreover, detailed performance evaluation is conducted varying the number of subscribed network elements, the type of communication (e.g., compressed vs. uncompressed, bundled vs. unbundled), showing remarkable scalability performance particularly in the case of compressed and bundled telemetry configurations.

Index Terms—Elastic Optical Networks, telemetry, monitoring, soft failure, disaggregated networks, white box, gRPC, YANG.

I. INTRODUCTION

ELASTIC Optical Networks (EONs) have gained a significant research interest in the last years and represent a mature technology to be deployed in core/transport networks [2]. The maturity of EON solutions is driving their adoption even in next generation metro networks, where relaxed optical reach constraints, along with advanced transmission adaptation capabilities (e.g., more spectrally-efficient modulation formats), are expected to successfully address the requirements for future bandwidth-hungry dynamic metro/5G services.

In both core and metro EONs, accurate and real-time knowledge of physical layer parameters and transmission impairments has the potential to significantly improve the efficiency of control and management operations [3]. Relevant benefits are expected in failure detection, localization, and

recovery mechanisms, particularly in the case of soft failures where just minor signal degradations are experienced (due to ageing, poor power supply quality, weather events, bad connecting), without causing relevant disruptions [4]. Detecting and localizing soft failures is extremely important since they may reduce system margin, introduce sporadic post-FEC errors or may reveal in advance possible major malfunctioning of network devices.

Effective soft failure detection is extremely important also in the context of the recently proposed disaggregated networks, where network elements as transponders, reconfigurable optical add drop multiplexers (ROADMs) and optical amplifiers are provided by different vendors (i.e., white box) and coexist in the same Software Defined Networking (SDN) control within a unique transparent optical domain [5]. That is, besides being fast and effective, EON monitoring solutions also have to provide standard and vendor-neutral functionalities, enabling on demand detailed analysis of transmission performance of the deployed network elements and activated communications.

However, current deployed monitoring devices and transponders typically provide statistics every 15 minutes. This time was historically selected to enable basic monitoring without overwhelming the management plane and the network controller with an excessive amount of data. Also network protocols deployed for carrying monitoring statistics, as SNMP, were designed and implemented to support such pace and amount of data. However, such basic monitoring and protocols, do not provide accurate and real-time knowledge of data plane info and are not really suitable for accurate and fast soft failure detection/localization. Recently, the availability of higher bandwidth in the control/management plane, combined with the introduction of big data analytics, machine learning and improved processing capabilities at the network controllers are overcoming the aforementioned limitations, opening the way for new and effective monitoring and management solutions, relying on the autonomic network concept [6] and on proactive, automatic event/forecast driven SDN control optimization and reconfiguration employing predictive models [7] [8].

In this paper, we propose, implement and experimentally demonstrate a telemetry service, providing on-demand, streaming-oriented, parallel, real-time monitoring statistics on a specifically selected set of network devices, particularly suitable for soft failure event detection and localizations. The telemetry service is proposed as additional control/management feature in the SDN control plane of EONs employing NETCONF protocol for device configuration and notification.

Manuscript received January 16, 2018.

F. Paolucci, A. Sgambelluri and P. Castoldi (email: {fr.paolucci, a.sgambelluri, castoldi}@santannapisa.it) are with Scuola Superiore Sant'Anna, Pisa, Italy. F. Cugini (email: filippo.cugini@cnit.it) is with CNIT, Pisa, Italy.

This paper is an extended version of the work presented in [1]. This work has been performed in the framework of the H2020-ICT-2017 project METRO-HAUL (Grant Agreement no. 671636), which is funded by the European Commission. Special thanks to Matteo Dallaglio who supported the authors in the design and setup of NETCONF agents.

First, two use cases are reported and discussed in order to motivate the adoption of a telemetry service. Then, the telemetry service is conceived in the framework of disaggregated optical networks, in particular the open-source YANG model enhancing the OpenROADM-based multi source agreement for multi-vendor interoperability [9]. Streaming-oriented telemetry service based on the gRPC protocol is proposed, providing efficient fine grained monitoring in the time domain. Two implemented telemetry services (partially and fully disaggregated) are then applied to a networking scenario where soft failure affects transmission performance of optical signals traversing multiple optical devices. In such scenario, traditional 15minute-based monitoring is not able to detect and localize the failure event. Instead, the demonstration shows that the proposed telemetry solution allows to identify the occurrence of a soft failure and to successfully localize the source of the issue. Finally, extensive scalability evaluation is provided to motivate the adoption of the gRPC protocol for large networks monitoring.

This paper extends the work in [1] by providing the following additional contributions:

- extends the telemetry service architectural proposal to disaggregated networks.
- proposes and discusses YANG and gRPC models for telemetry.
- presents extended results including an additional use case, the application to disaggregated networks and scalability evaluations in terms of network load and controller processing.

II. RELATED WORK

A number of studies have proposed and addressed the issue of optical performance monitoring, mainly focused on hard failures detection, end-to-end lightpath QoT continuous monitoring, techniques to retrieve and derive physical parameters statistics of optical signals that are not directly available by online monitors in currently deployed EON [3].

In the context of the Architecture on Demand programmable multi-layer optical network, a network monitoring system framework based on SDN network analytics for network re-planning and reoptimizations was proposed [10]. The system introduced a monitoring hub handler collecting optical performance monitoring data at different layers (e.g., Optical Signal to Noise Ratio - OSNR in the optical-switched layer, packet loss rate in the packet-switched layer) and triggering SDN reconfigurations upon changes in the network. The concept of living network with disaggregated components, with continuous monitoring, estimation of active connections QoT and provisioning computation and assignment (i.e., route, spectrum, code, modulation format) of new connections is based on such monitoring results [11].

Telemetry is gaining attention besides EON as well. A dedicated low-rate data-plane telemetry channel has been recently proposed and evaluated in the context of Passive Optical Networks (PON) to monitor the connections between an intelligent remote splitter and attached Optical Network Units (ONUs) [12].

In the very recent years, the improved availability of cloud-based storage and processing capabilities has raised the interest towards massive analytics of monitoring data to perform continuous planning, advanced lifecycle models and network proactive control based on forecast [13].

Focusing on SDN control and management plane, the work in [14] proposes an architecture supporting telemetry for YANG defined packet network parameters (e.g., one-way delay) in the framework of the Automatic Control Transport Network (ACTN). A related work focuses on YANG models for monitoring Traffic Engineering (TE) tunnels, providing end-to-end monitored data referred to a tunnel [15]. However, the work targets only end-to-end parameters and IP-layer variables.

Finally, the Distributed Network Analytics (DNA) model proposes the activation of YANG-driven massive analytics directly at equipment devices supporting processing and storage capabilities (e.g., fog nodes) and reporting functions [16].

However, none of the above studies has designed, implemented and experimentally demonstrated the telemetry service in the context of partially or fully disaggregated EON, also highlighting practical uses cases for effective soft failure detection and localization.

III. SOFT FAILURE USE CASES MOTIVATING TELEMETRY SERVICE

In order to motivate the need of telemetry services in EON, a couple of use cases are hereafter presented, analyzed and discussed. Both use cases are focused on soft failures affecting, as an example, optical amplifiers. In the first use case, direct telemetry of monitoring parameters is configured and activated on all optical devices. In the second use case, the amplifier is a black box not allowing direct communication through telemetry for streaming of parameter information and correlation is exploited taking advantage of telemetry information retrieved from other optical devices.

A. EON with EDFA monitoring

A portion of EON shown in Fig. 1 is considered. The network includes a 400km link $N3-N4$ of 5-span. A lightpath L_A generated at $N2$ and terminated at $N4$ is a 100G PM-QPSK provided by a commercial system. Under steady state normal conditions, the lightpath has OSNR statistics in the range between 22.3 and 22.0 dB, retrieved by the coherent receiver. As a typical commercial system [17], the 100G transponder provides such statistics every 15min (along with additional transmission parameters such as pre-FEC BER, received power, etc) through SNMP or NETCONF protocols to the central network controller. A minor failure then occurs on an intermediate device, such as the optical amplifier $D5$. In particular, the amplifier starts introducing limited (e.g., 1-2dB) and sporadic fluctuations on its launch power. Such fluctuations do not significantly impact the received power of L_A since it is recovered by the subsequent EDFA $D6$ (Pre-amplifier of node $N4$), typically configured to provide constant output power (i.e., automatic power control, APC). However, on span $D5-D6$, such fluctuations slightly affect the signal to

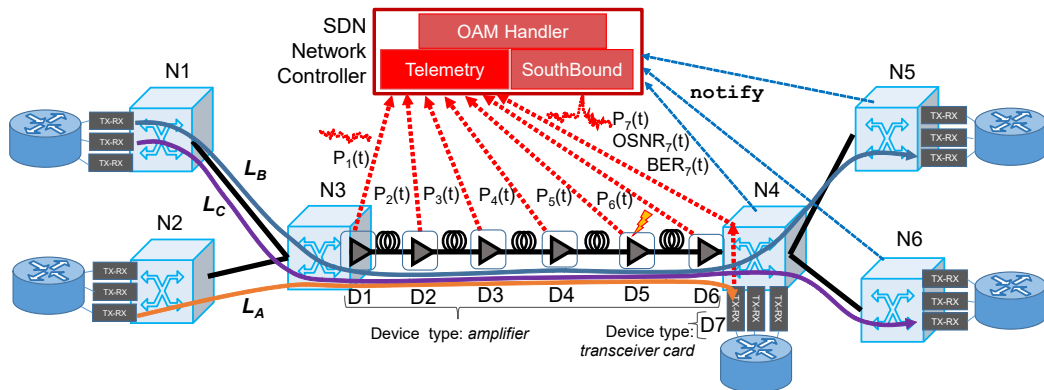


Figure 1. Soft failure use case with available line amplifiers monitoring.

noise ratio of the traversing signals, inducing a so called soft failure on lightpath L_A : the overall power at the receiver is not changed while the received OSNR statistics sporadically change in the range between 21.6 and 21.0 dB. The average statistics collected every 15min are typically not adequate to detect sporadic minor issues affecting the EDFA launch power or OSNR performance, which remains largely below warning thresholds. This implies that the minor malfunctioning of the EDFA (i.e., the soft failure) is traditionally not detected, even if it may be extremely useful to reveal in advance that a device is running out of specs, with high probability to experience in the near future a major fault, potentially affecting many end-to-end services. Thus, in this case, high rate signal correlations in the time domain between the EDFAs input power and the receiver OSNR, available at the controller in the case of telemetry service, may be extremely useful to identify the malfunctioning EDFA.

B. EON without EDFA monitoring

In the case EDFAs are black boxes not equipped with a monitoring interface, soft failure detection procedures may involve additional monitoring points. This is the case of Fig. 2, where lightpath L_A generated at $N1$ crosses $N2$ and terminates at $N3$ (card $D7$), while lightpath L_B generated at $N2$ terminates at $N3$ as well (card $D6$). Upon amplifier $D5$ soft malfunctioning, high rate correlation in the time domain between OSNR values of $D6$ and $D7$ may allow to detect whether the failure is affecting either amplifier $D5$ (crossed by both lightpaths) or the $D1$ - $D4$ chain (crossed only by L_A). In general, existing techniques (e.g., network kriging [18]) may be employed to predict a statistical behavior of a device signal not directly measurable, using probes in the network. For example, a high number of OSNR telemetry monitors (i.e., at receiver cards of active lightpaths) can be activated to detect and localize soft failures occurring in some black box. The feasibility of such service implies a strict scalability constraint in terms of network load, processing and memory, since a network controller has to handle a continuous parallel data retrieving process from a potentially high number of monitored devices.

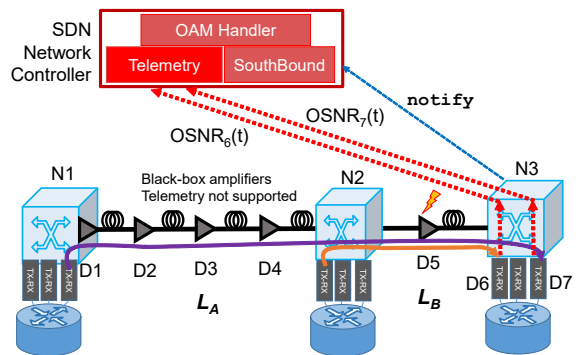


Figure 2. Soft failure use case with line amplifiers not supporting monitoring.

IV. SDN TELEMETRY SERVICE

Network telemetry, i.e. streaming of real-time sample values for in-depth analysis, with sampling rate of seconds or below, is proposed to be employed as additional finer monitoring service in the SDN control/management plane of EON. The telemetry service is conceived to be activated on-demand, either manually by the maintenance operator or automatically by network controllers, for a selected amount of time in parallel over a subset of target devices. The telemetry service operates as additional monitoring system with respect to the traditional monitoring (e.g., exploiting NETCONF protocol) which relies on statistics provided every 15mins and on notifications rapidly triggered in case of major faults/alarms. The set of data plane devices selected for telemetry is first identified by means of correlation methods. For example, in the use case of Fig. 1, the devices belonging to links shared by lightpaths L_A , L_B and L_C experimenting end-to-end signal degradations monitored at the receiver card are selected for telemetry activation. Correlation algorithms (e.g., regression test) in the time domain or Artificial Intelligence (AI)-based machine learning algorithms can then be applied on the telemetry samples to perform fast soft failure detection and localization.

Two different solutions can be considered for the implementation of telemetry services.

In the first solution, called *embedded*, the same control/management southbound protocol used for configuration and monitoring (e.g., NETCONF) is utilized also for the telemetry

service. Embedded telemetry services utilize YANG/NETCONF by opening additional XML-based connection (e.g., on different TCP port) for telemetry activation. On the one hand, this embedded solution simplifies the telemetry implementation since it relies on a unique southbound interface. On the other hand, it requires augmented YANG models and dedicated NETCONF procedures to be implemented within the control/monitoring system. Moreover, it induces bandwidth inefficiency since it exploits XML-based messages which typically include large per-message overhead.

In the second solution, called *independent*, a different dedicated protocol/API with respect to the southbound protocol utilized for control/management is utilized for telemetry service. Different candidate APIs may be utilized for the telemetry service, including IPFIX, JDB, Thrift, and gRPC. For example, IPFIX has been traditionally considered in the context of packet networks. Thrift and gRPC have been recently released by Facebook and Google respectively as a collection of software libraries and code generation tools that can be used to automatically generate the client and server code for distributed telemetry applications. The disadvantage of such independent solution is that additional API is required to support the telemetry service. However, such implementations provide lightweight API for fast data delivery at high rates, including support of streaming mode with extremely low overhead, suitable for distributed sensors.

Two independent telemetry service architectures based on the gRPC protocol are then hereafter proposed: a fully disaggregated (FD) architecture, suitable for white boxes, and a black-box (BB) architecture, suitable for proprietary nodes/agents not compliant with the white box framework.

A. Fully Disaggregated Architecture

Fig. 3 shows the proposed FD architecture exploiting telemetry for soft failure detection. Three hierarchy levels are considered, following the OpenROADM framework of optical white boxes in disaggregated networks [9]: the SDN-based Network Controller (NC), responsible for the entire network domain configuration and management; the Network Element Controller (NEC), responsible for the configuration of white boxes (e.g., ROADMs, xponder box), and the Device Agent (DA), controlling a single disaggregated optical device (e.g., line power amplifier, node degree, transponder, BV-WSS).

SDN network configuration and status retrieval control (e.g., controller-triggered lightpath setup and lightpath operational status) relying on southbound NETCONF-based interfaces is supposed, occurring in the NC-NEC-DA chain. Asynchronous notification is supported, allowing the bottom-up communication of lightpaths/signal parameter degradations or hard failures [19]. Typically, such notifications are configured to be forwarded when performance parameters (e.g., pre-FEC BER), averaged in a given time interval (e.g., 15 minutes), exceed a warning or a critical threshold or deviate from typical operational values. Upon NETCONF degradation notification, the upper layer OAM may decide to trigger a set of telemetry subscription instances for fine-grained live monitoring (step 1). The subscription specifies the target device (also called

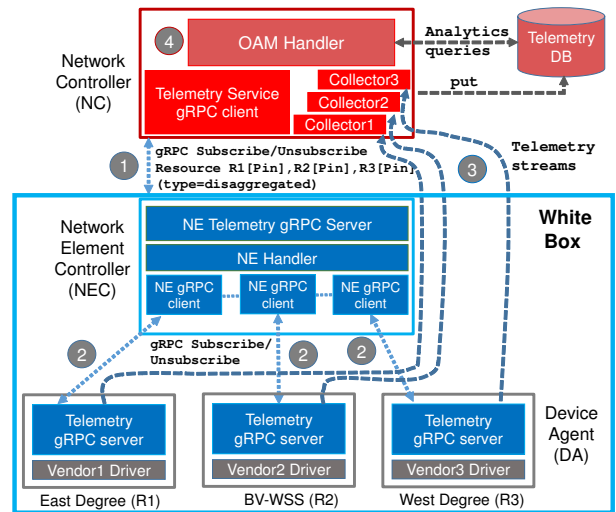


Figure 3. Telemetry service architecture for fully disaggregated (FD) white boxes.

resource), the physical parameter to monitor and the requested sampling interval. To this purpose, NC is equipped with a novel dedicated telemetry service module based on gRPC protocol clients [20]. Moreover, the NEC implements a gRPC server and a NE Handler, responsible for receiving subscription requests and direct them to the target DA through an array of NE gRPC clients (step 2). Finally, gRPC servers at DAs receive the telemetry subscription and start pushing real time streaming of the monitored data directly received by the hardware drivers (step 3). In Fig. 3 telemetry streams are directed to dedicated NC Collectors (created during the subscription and specified in the subscription request) for network-level correlation. Collectors store stream data in a Telemetry open database [21] to enable, if needed, off-line massive analytics and complex and more time-consuming processing (e.g., machine learning algorithms [22] [23]). Once the OAM Handler has performed time domain correlations and/or identified the soft failure source, it stops the telemetry streams by sending an unsubscribe request following the same procedure. Based on the discovered soft failure sources, the OAM Handler may raise management alarms for specific device replacement and the controller may take proactive actions to guarantee the current quality of transmission, e.g., SDN lightpath restoration [24] or adaptation [25].

With the FD architecture, independent subscriptions are allowed to each white box device with the maximum flexibility (i.e., subscriptions sampling interval may be different based on the sampling granularity of each device). The architecture also allows, for each device, a subscription to a configurable number of monitoring parameters, hereafter called *bundle subscription*, for which a single streaming service is activated providing telemetry of multiple parameters of the same device.

B. Black Box Architecture

Fig. 4 shows the BB architecture enabling the telemetry service for partially disaggregated or proprietary nodes or devices. In this case disaggregation is not supported, the

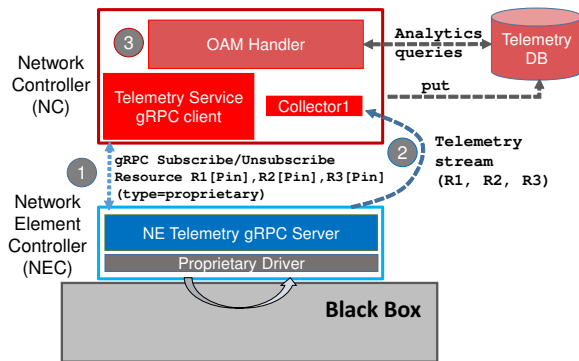


Figure 4. Telemetry service architecture for partially disaggregated or black box (BB) Network Elements.

hierarchy is collapsed and a single NE telemetry server is responsible for all the subscription requests towards that NE. In particular, subscription request for multiple physical parameters in the same black box trigger a single telemetry stream conveying all the requested parameters in a single gRPC sample message (i.e., bundle subscription). In this case, flexibility is limited and telemetry data serialization and subscription parameters such as the sampling interval strictly depend on the features of the proprietary driver.

V. TELEMETRY SERVICE: GRPC AND YANG MODELS

gRPC is a HTTP/2 open source high performance Remote Procedure Calling (RPC) framework developed by Google that can run in any environment [20]. It includes pluggable support for load balancing, tracing, health checking and authentication. The choice of gRPC as telemetry protocol is dictated by a number of reasons. First, its capability to support high performance server push streaming mode, allowing, besides independent YANG-based dataset and method definition, also binary data format compression and efficient wire encoding through protocol buffers serialization. Second, payload-agnostic, open source and vendor-neutral features enable its use in disaggregated networks and assures compatibility with existing and new data formats and encodings. Third, it provides a lightweight library platform designed to limit the impact of processing, networking and also energy consumption burden. Such capabilities make gRPC attractive to connect devices and sensors for telemetry streaming with respect to standard HTTP, RESTCONF, NETCONF interfaces based on native XML or JSON formats.

To activate gRPC telemetry on a target device, a set of specific planning and engineering parameters are included in the YANG model available at the controllers [9]. They include central office location, network element type, network element identifier assigned within the given central office, a unique network element ID within the entire network, Vendor identifier of the supplier and related model number identifier, IP address and domain subnetwork, rack/shelf ID in which the device/element controller is equipped, and the number of controlled devices under the same device controller (e.g., xponder ID within a shelf encompassing multiple xponders, all controlled by the same network element controller).

```

syntax = "proto3";
package Telemetry;
service SDNTelemetry {
  rpc telemetrySubscribe(SubscriptionRequest) returns (SubscriptionResponse) {}
  rpc cancelTelemetrySubscription(SubscriptionId) returns (NoMessage) {}
}
service TelemetryStream {
  rpc StreamData(stream TelemetryData) returns (NoMessage) {}
}
message SubscriptionRequest {
  bool disaggregated = 1;
  repeated Collector collectors = 2;
  repeated Resource resources = 3;
}
message Collector {
  string ip_address = 1;
  uint32 port = 2;
}
message Resource {
  Path path = 1;
  string path_filter = 2;
  int64 sampling_interval = 3;
}
message TelemetryData {
  string system_id = 1;
  uint32 component_id = 2;
  uint32 sub_component_id = 3;
  uint64 timestamp = 4;
  repeated KeyValue kv = 5;
}

```

Figure 5. Telemetry service: proposed gRPC model definition.

Thanks to the provided parameters, the OAM Handler within the SDN Controller can identify all devices of interest and trigger parallel gRPC subscriptions for telemetry to all related NECs targeting specific internal resources (i.e., devices).

The specific parameters to be selected for subscriptions are then identified in a standard and vendor neutral way by relying on a common YANG data model. For example, an optical channel interface of an xponder is defined by a specific list of optical channel attributes (expressed as leaves in the YANG tree) as *rate*, *frequency*, *width*, *modulation-format*, and *transmit-power*. Such terminology is univocally defined together with the related unit of measure (e.g., dBm for power). This way, the OAM service is able to identify and select all relevant network element and device parameters for telemetry subscription.

Fig. 5 shows an excerpt of the proposed gRPC model for telemetry. The model includes two so-called services, the *SDNTelemetry* and the *TelemetryStream* service. The former defines subscription and unsubscription procedures, the latter provides the streaming handler. In particular, the subscription request, besides the list of collectors (i.e., the streaming destination endpoints) and resources (i.e., the streaming originator, located at the devices), includes the *disaggregated* boolean field. This parameter is utilized by the NEC to implement either FD or BB procedure. In the FD case, *disaggregated* is set to `True` and the NEC triggers an additional subscription array hierarchy towards the controlled device agents, utilizing the same gRPC model. Otherwise, if set to `False`, it implements BB procedure collecting requested values directly from the proprietary driver. The resource model includes also the *sampling_interval*, specifying the requested telemetry rate for the specific target. Finally, the *TelemetryData* includes, besides OpenROADM-compliant device identifiers, an array of *kv* values (specified in the subscription request among the available values provided by that resource) and a timestamp. This way, each sampled data is provided with an absolute time reference, allowing time-domain correlations at the OAM Handler.

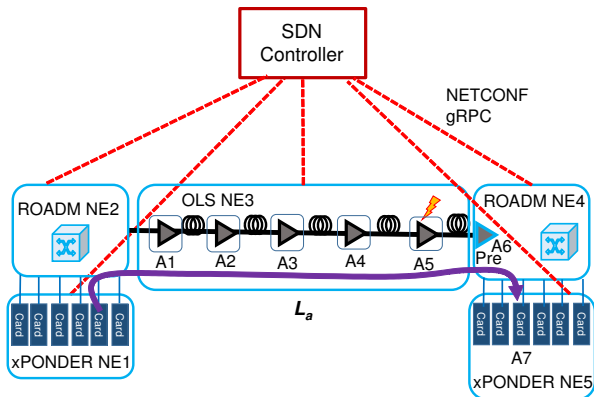


Figure 6. EON testbed employing OpenROADM-based network elements.

VI. DEMONSTRATION AND SCALABILITY RESULTS

To assess the proposed *independent* telemetry-enabled solution, an experimental demonstration has been carried out on the EON testbed of Fig. 6 representing a subset of Fig. 1 following the OpenROADM framework.

The testbed data plane includes 5 optical network elements. ROADMs and xPONDER nodes employ Ericsson SPO-1400 ROADMs (vendor A) equipped with commercial 100G PM-QPSK transponder cards. The NE3 Optical Line System (OLS) link includes 5 Alcatel EDFA line amplifiers (A1-A5) (vendor B) and 5 80km-long fiber spans. A Finisar Waveshaper Spectrum Selective Switch placed between A5 and the fiber span is inserted to emulate the soft failure. The control plane includes a NETCONF-enabled SDN controller [19], NETCONF NECs written in C++ based on the ConfD tool (a NETCONF server implementation made by Tail-f and CISCO DEVNET [26]) and device agents fully extended with gRPC modules. Each agent is a python-based monitoring system that leverages on proprietary protocols for the communication with data-plane commercial devices. Different types of agents are available. In particular, agents A1-A5 (type amplifier) of the OLS monitor the line power, whereas *Pre* A6 and *Card* A7 agents at ROADMs and xPONDER monitor the input received power, the pre-FEC BER and the estimated OSNR, respectively.

Lightpath L_a (bitrate 100G, central frequency $f_c=193.1$ THz, $m=3$ reserved frequency slots corresponding to 37.5GHz channel width) is setup from NE1 to NE5, reproducing the same OSNR conditions described in Sec. III-A. Standard NETCONF notification is activated when the NE5 card OSNR average value in the last 15 minutes falls below 21.7dB, applying 0.1dB margin with respect to the 21-21.6dB soft failure condition range detected in Sec. III-A. A soft failure of device A5 is induced by reducing its launch power of 1-2dB in a random pattern. The subsequent 15min slot triggers NE5 NEC NETCONF notification due to the slight OSNR degradation. Indeed, the controller identifies the NE3 link as a candidate degradation source. However, relying on just 15mins average statistics (i.e., without telemetry), it would not be able to localize the soft failure. Thus, to localize the soft failure, the controller triggers parallel gRPC telemetry subscriptions on the A1-A5 devices, NE4 *Pre* and NE5 card.

Both FD and BB telemetry service instances have been implemented and performed. Fig. 7 reports control messages captures exchanged by the NC (IP 10.1.1.250) and the NE4-NE5 NEC (IP 10.1.1.1) in the case of BB. In this case NE4 and NE5 are envisioned as a single proprietary black box. Thus, a single NEC provides a single telemetry stream conveying different monitors (i.e., the A6 *Pre* input power, the A7 card input power and OSNR). The capture shows the subscription step S (frames 33-41) triggering telemetry streaming t_0 directed to the NC. The subscription step is performed in less than 5ms. The BB sample streaming message is encoded in only 123 bytes. The same information collected through a NETCONF notification using the same YANG model [19] with the *embedded* architecture is encoded in 556 bytes. Fig. 7 reports also the FD implementation. In this case, NE4 and NE5 are seen as different elements, each one equipped with a dedicated controller. Thus, the NEC implements a hierarchical telemetry subscription towards NE4 and NE5. The capture shows the subscription step S (frames 43-51) with the disaggregated flag set to `True`, triggering two parallel telemetry subscriptions towards NE4 and NE5 (not appearing in the capture at NC), which are in fact initiated having NC as collector and originated by NE4 (IP 10.1.1.2, stream t_1 carrying *Pre* input power) and NE5 (IP 10.1.1.3, stream t_2 carrying card OSNR), respectively. The whole subscription chain is completed after around 15ms. The FD sample streaming message is encoded in only 110 bytes. The same information collected through a NETCONF notification is encoded in 456 bytes. Fig. 8 shows the telemetry monitors (1s sampling rate) in the time domain acquired in parallel for each device in 15 minutes at the NC. Telemetry analysis clearly shows that *Pre* P_{in} , i.e. $P_6(t)$, and OSNR at card are highly correlated, while all the other $P_x(t)$ monitors, including the receiver card P_{in} , do not reveal significant correlations. This means that the issue is not located at NE4 and NE5, since receiver card P_{in} does not reveal any degradation. Moreover, each upstream device with respect to uncorrelated power monitor is supposed to work properly. This result allows the controller to successfully identify amplifier A5 as the soft failure source.

Telemetry scalability has been investigated by setting lower sampling times, enlarging the number of monitored devices in the network and measuring the impact in terms of network load and processing burden at the network controller. To evaluate this, a variable number of agents has been replicated as Docker Containers and subscribed for parallel telemetry streaming, also utilizing the data compression option provided by gRPC. The network controller PC employs a Intel Core i5-2400 quad core processor (3.10GHz) equipped with 4GB RAM.

Fig. 9 reports streaming network load in terms of throughput (max rate and average rate) and the processor utilization using compressed (*zip*) and uncompressed (*raw*) gRPC as a function of the number of subscribed agents. In particular, each subscription is requested for a single parameter (i.e., no subscription bundling is performed) at 500ms sampling interval. Results show that the average traffic rate even for a large number of monitored agents (up to 100 devices) is always below 200kb/s, while the maximum rate is always 1Mb/s. Increased streaming burstiness is experienced employ-

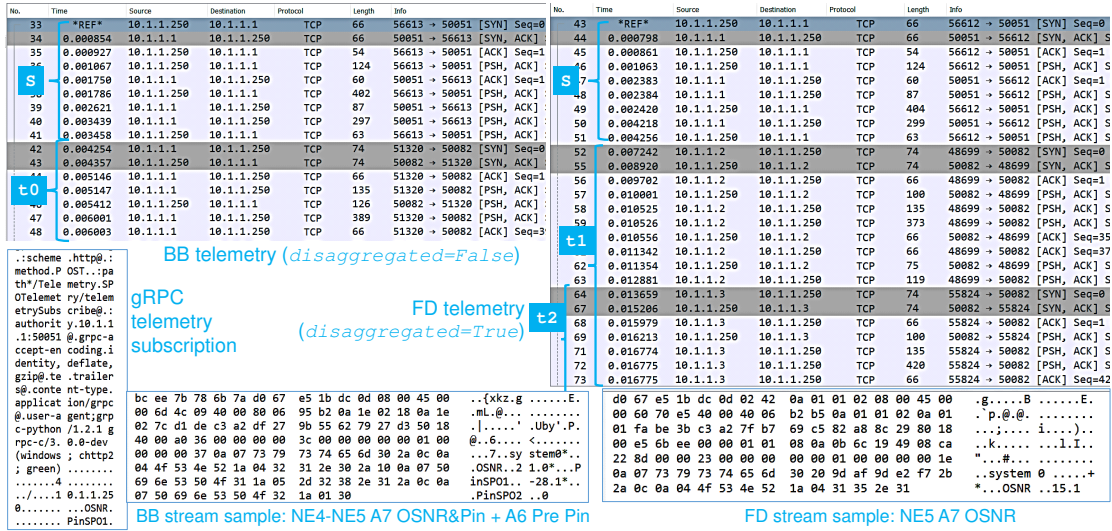


Figure 7. Captures of telemetry messages in the BB and FD configurations collected at the SDN network controller.

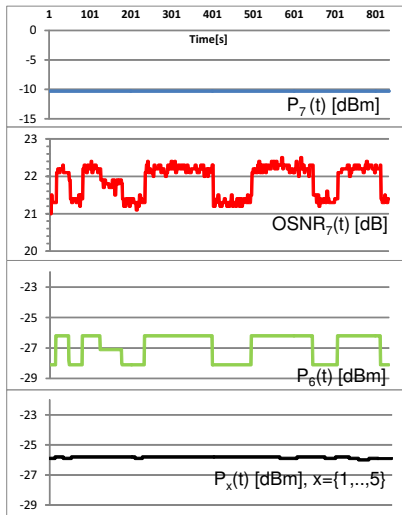


Figure 8. Telemetry monitors of devices A1-A5 ($P_x(t)$), NE4 A6 Pre power ($P_6(t)$) and NE5 A7 card power ($P_7(t)$) and OSNR ($OSNR_7(t)$) available at SDN network controller.

ing uncompressed gRPC, whereas using compressed gRPC the max rate is always below 400kb/s. CPU utilization follows a quasi-linear behavior, always below 5% of overall utilization, with no noticeable differences between the compressed and the uncompressed option, thus confirming the lightweight gRPC burden. Moreover, the compressed/uncompressed option does not substantially affect the average throughput since in this case a single parameter is monitored per subscription.

In order to evaluate the impact of the compressed option, the telemetry average throughput has been measured for a number $N=20$ of independent subscriptions targeting a single parameter (*single20*) and compared with the equivalent bundle subscription, i.e. an unique subscription targeting a bundle of $N=20$ parameters (*bundle20*) of the same device. Results depicted in Fig. 10 show that *single20* compression does not achieve a significant gain (i.e., around 5% rate reduction) and confirms results of Fig. 9, while *bundle20* compression




Figure 9. Telemetry scalability results: network and processor load at the controller as a function of the size of monitored agents (no bundling).




Figure 10. Telemetry bundling/unbundling average network load with compressed (*zip*) and uncompressed (*raw*) sampled data.

achieves more than 50% of throughput reduction, since the gRPC compression operates on the bundle of sampled values (i.e., in the payload of the gRPC message) and not on the message header. Indeed, results highlight the benefits of the bundling option itself. This means that telemetry of multiple

parameters originated by a single device (available for both BB and FD configurations) is extremely scalable providing around 75% throughput reduction with respect to independent subscriptions.

Such considerations raise to the conclusion that gRPC protocol is extremely effective to provide a high number of parallel telemetry streams without affecting control plane load and controllers processing resources in both FD and BB architectures. Moreover, gRPC achieves more than 4 times traffic load reduction with respect to NETCONF.

VII. CONCLUSIONS

A telemetry service exploiting the gRPC protocol was designed, implemented and experimentally demonstrated to efficiently enable on-demand streaming of real-time monitoring parameters, dynamically retrieved from a configurable set of network devices. The telemetry service, relying on standard YANG-defined network models is suitable for effective application in EONs also encompassing disaggregated network elements.

The implemented telemetry service was experimentally validated over two networking scenarios where soft failures affecting optical amplifiers are successfully detected and localized by retrieving and correlating fast OSNR and optical power variations not detectable through traditional 15-minute monitoring. The telemetry service was shown in the context of both fully or partially disaggregated scenarios, also providing scalability performance evaluations. In particular, scalability was assessed by varying the number of monitored devices in the network and measuring the network and processing load at the network controller, with and without data compression. Results showed that the traffic rate for even 100 subscribed optical devices is very low (average below 200kb/s and max rate below 400kb/s for compressed data). CPU utilization follows a quasi-linear behaviour, generally providing extremely low utilization (i.e., below 5%). Efficient gRPC encoding achieved more than 4 times telemetry traffic reduction with respect to XML-based NETCONF notifications.

Finally, results showed the benefits of the bundling option, enabling around 75% of throughput reduction compared to independent uncompressed subscriptions, leading to the conclusion that such bundling-based telemetry successfully supports high number of parallel telemetry streams without affecting control plane and processing load at the SDN Controller.

REFERENCES

- [1] F. Paolucci, A. Sgambelluri, M. Dallaglio, F. Cugini, and P. Castoldi, "Demonstration of gRPC telemetry for soft failure detection in elastic optical networks," in *Optical Communication (ECOC), 2017 European Conference on*, Sept 2017.
- [2] M. Jinno, "Elastic optical networking: Roles and benefits in beyond 100-gb/s era," *Journal of Lightwave Technology*, vol. 35, no. 5, pp. 1116–1124, March 2017.
- [3] Z. Dong, F. N. Khan, Q. Sui, K. Zhong, C. Lu, and A. P. T. Lau, "Optical performance monitoring: A review of current and future technologies," *Journal of Lightwave Technology*, vol. 34, no. 2, pp. 525–543, Jan 2016.
- [4] N. Sambo, F. Cugini, A. Sgambelluri, and P. Castoldi, "Monitoring plane architecture and OAM handler," *Journal of Lightwave Technology*, vol. 34, no. 8, pp. 1939–1945, April 2016.
- [5] P. Layec, A. Dupas, D. Verchère, K. Sparks, and S. Bigo, "Will metro networks be the playground for (true) elastic optical networks?" *Journal of Lightwave Technology*, vol. 35, no. 6, pp. 1260–1266, March 2017.
- [6] L. Velasco, L. Gifre, J. L. Izquierdo-Zaragoza, F. Paolucci, A. P. Vela, A. Sgambelluri, M. Ruiz, and F. Cugini, "An architecture to support autonomic slice networking," *Journal of Lightwave Technology*, vol. PP, no. 99, pp. 1–1, 2017.
- [7] Y. Li, W. Mo, S. Zhu, Y. Shen, J. Yu, P. Samadi, K. Bergman, and D. Kilper, "tsdx: Enabling impairment-aware all-optical inter-domain exchange," *Journal of Lightwave Technology*, vol. PP, no. 99, 2017.
- [8] F. Morales, L. Gifre, F. Paolucci, M. Ruiz, F. Cugini, L. Velasco, and P. Castoldi, "Experimental assessment of a flow controller for dynamic metro-core predictive traffic models estimation," in *Optical Communication (ECOC), 2017 European Conference on*, Sept 2017.
- [9] "Open ROADMs MSA device white paper v1.0," www.openroadm.org, Oct. 2017.
- [10] S. Yan, A. Aguado, Y. Ou, R. Wang, R. Nejabati, and D. Simeonidou, "Multilayer network analytics with sdn-based monitoring framework," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 9, no. 2, pp. A271–A279, Feb 2017.
- [11] S. Oda, M. Miyabe, S. Yoshida, T. Katagiri, Y. Aoki, T. Hoshida, J. C. Rasmussen, M. Birk, and K. Tse, "A learning living network with open ROADMs," *Journal of Lightwave Technology*, vol. 35, no. 8, pp. 1350–1356, April 2017.
- [12] P. Iannone, A. H. Gnauck, M. Straub, J. Hehmann, L. Jentsch, T. Pfeiffer, and M. Earnshaw, "An 8-x 10-gb/s 42-km high-split twdm pon featuring distributed raman amplification and a remotely powered intelligent splitter," *Journal of Lightwave Technology*, vol. 35, no. 7, pp. 1328–1332, April 2017.
- [13] L. Velasco, A. P. Vela, F. Morales, and M. Ruiz, "Designing, operating, and reoptimizing elastic optical networks," *Journal of Lightwave Technology*, vol. 35, no. 3, pp. 513–526, Feb 2017.
- [14] Y. Lee, R. Vilalta, R. Casellas, R. Martínez, and R. Muñoz, "Scalable telemetry and network autonomies in actn sdn controller hierarchy," in *2017 19th International Conference on Transparent Optical Networks (ICTON)*, July 2017.
- [15] Y. Lee, D. Dhody, R. Vilalta, D. King, and D. Ceccarelli, "YANG models for ACTN TE performance monitoring telemetry and network autonomies," *IETF, TEAS Working Group*, Mar. 2017.
- [16] M. Chandramouli and A. Clemm, "Model-driven analytics in sdn networks," in *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, May 2017, pp. 668–673.
- [17] M. Filer, J. Gaudette, M. Ghobadi, R. Mahajan, T. Issenuth, B. Klinkers, and J. Cox, "Elastic optical networking in the microsoft cloud [invited]," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 8, no. 7, pp. A45–A54, July 2016.
- [18] K. Christodoulopoulos, N. Sambo, and E. Varvarigos, "Exploiting network kriging for fault localization," in *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, March 2016.
- [19] M. Dallaglio, N. Sambo, F. Cugini, and P. Castoldi, "Control and management of transponders with netconf and yang," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 9, no. 3, pp. B43–B52, March 2017.
- [20] "gRPC: a high performance, open-source universal RPC framework," <https://grpc.io/>.
- [21] F. Paolucci, F. Cugini, G. Cecchetti, and P. Castoldi, "Open network database for application-based control in multilayer networks," *Journal of Lightwave Technology*, vol. 35, no. 9, pp. 1469–1476, May 2017.
- [22] J. Thrane, J. Wass, M. Piels, J. C. M. Diniz, R. Jones, and D. Zibar, "Machine learning techniques for optical performance monitoring from directly detected pdm-qam signals," *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 868–875, Feb 2017.
- [23] A. P. Vela, M. Ruiz, F. Cugini, and L. Velasco, "Combining a machine learning and optimization for early pre-fec ber degradation to meet committed qos," in *2017 19th International Conference on Transparent Optical Networks (ICTON)*, July 2017, pp. 1–4.
- [24] A. Giorgetti, F. Paolucci, F. Cugini, and P. Castoldi, "Dynamic restoration with GMPLS and SDN control plane in elastic optical networks [invited]," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 7, no. 2, pp. A174–A182, February 2015.
- [25] F. Cugini, F. Fresi, F. Paolucci, G. Meloni, N. Sambo, A. Giorgetti, T. Foggi, L. Potá, and P. Castoldi, "Active stateful PCE with hitless LDPC code adaptation [invited]," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 7, no. 2, pp. A268–A276, 2015.
- [26] "ConfD," developer.cisco.com/site/confD/.