

Autonomic Disaggregated Multilayer Networking

Lluís Gifre,  Jose-Luis Izquierdo-Zaragoza,  Marc Ruiz,  and Luis Velasco 

Abstract—Focused on reducing capital expenditures by opening the data plane to multiple vendors without impacting performance, node disaggregation is attracting the interest of network operators. Although the software-defined networking (SDN) paradigm is key for the control of such networks, the increased complexity of multilayer networks strictly requires monitoring/telemetry and data analytics capabilities to assist in creating and operating *self-managed (autonomic)* networks. Such autonomicity greatly reduces operational expenditures, while improving network performance. In this context, a monitoring and data analytics (MDA) architecture consisting of centralized data storage with data analytics capabilities, together with a generic node agent for monitoring/telemetry supporting disaggregation, is presented. A YANG data model that allows one to clearly separate responsibilities for monitoring configuration from node configuration is also proposed. The MDA architecture and YANG data models are experimentally demonstrated through three different use cases: *i) virtual link creation* supported by an optical connection, where monitoring is automatically activated; *ii) multilayer self-configuration* after bit error rate (BER) degradation detection, where a modulation format adaptation is recommended for the SDN controller to minimize errors (this entails reducing the capacity of both the virtual link and supported multiprotocol label switching-transport profile (MPLS-TP) paths); and *iii) optical layer self-healing*, including failure localization at the optical layer to find the cause of BER degradation. A combination of active and passive monitoring procedures allows one to localize the cause of the failure, leading to lightpath rerouting recommendations toward the SDN controller avoiding the failing element(s).

Index Terms—Active and passive monitoring; Autonomic networking; Disaggregated multilayer networks.

I. INTRODUCTION AND MOTIVATION

Node disaggregation (*whitebox*) focuses on radically reducing capital expenditures, allowing one to create real multi-vendor networks, which are not just based on control plane interoperability [1], without compromising features and/or network performance. Different levels of disaggregation might be conceived, either at the hardware or software level or both simultaneously; for the sake of generalization, we denote a *node* as a self-contained

element exposing certain levels of programmability to software-defined networking (SDN) controllers.

Today, packet-layer whiteboxes are becoming a reality, where extensions for optical devices are still in progress [2]. Several initiatives are working on the definition of optical interoperability in wavelength division multiplexing (WDM) and elastic optical networks scenarios [3] to produce specifications and YANG data models [4], such as OpenROADM [5] and OpenConfig [6]. However, it is unclear how to operate a disaggregated network and whether such disaggregation will increase the complexity of the network, thus increasing operational expenditures (OPEX). Here, the role of monitoring and telemetry is gaining traction as an enabler for real-time self-managed (*autonomic*) networking [7]. Such autonomic networks will reduce human intervention, ultimately reducing OPEX; thus, because network performance information based on measurements is used to identify, isolate, and solve potential issues as quickly as possible, networks will run smoothly. Then, to build autonomic networks, monitoring and telemetry data must be not only gathered from networking devices but also analyzed for knowledge discovery from data (KDD), so that *recommendations* to the SDN controller can be produced aiming at improving network performance.

In this regard, passive monitoring techniques are the most-common choice in optical networks because they entail using noninvasive methods to obtain measurements; examples include measuring bit error rate (BER) in optical transponders, optical power in fiber links, and acquiring the optical spectrum in the whole C-band in links using optical spectrum analyzers (OSAs). Nonetheless, some active monitoring techniques are also available at the optical layer; for instance, the authors in [8] proposed to use in-band optical supervisory channel (OSC) devices to monitor the BER of 100 Gb/s dual-polarization quadrature phase shift keying lightpaths at transit points, i.e., not in the transponders. This technique consists of overmodulating the target lightpath with a low-speed low-bandwidth on-off keying (OOK) signal in an OSC_{TX} device at the ingress node. Then, the BER of the OOK signal is measured at intermediate locations by OSC_{RX} devices using low-speed components, and the BER of the lightpath is estimated through BER correlation curves obtained *a priori*. A similar technique was recently used in [9] for lightpath commissioning testing.

Regarding architectures supporting monitoring [10], the authors in [11] define a distributed data analytics framework, so-called CASTOR, including *extended nodes*, which run close to the network nodes, and a big data centralized

Manuscript received January 2, 2018; revised March 5, 2018; accepted March 5, 2018; published April 5, 2018 (Doc. ID 318799).

L. Gifre is with Universidad Autónoma de Madrid (UAM), Madrid, Spain. J.-L. Izquierdo-Zaragoza, M. Ruiz, and L. Velasco (e-mail: lvelasco@ac.upc.edu) are with the Optical Communications Group (GCO), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain.
<https://doi.org/10.1364/JOCN.10.000482>

monitoring and data analytics (MDA) system running in the control and management plane. Extended nodes collect monitoring data records from configured *observation points* (OP) in the nodes, which can be aggregated therein to reduce the amount of data sent toward the MDA system and for KDD to proactively notify the MDA system about network anomalies and degradations. The result from local KDD processes could be additionally exploited for network-wide control loops, including recovery and re-optimization [12,13].

In this paper, we extend the architecture in [11]. The MDA system includes the operational databases (e.g., topology and connections) retrieved from the SDN controller, which entails having a complete view of the network. A generic multi-node agent (*hypernode*) has been conceived to support multilayer disaggregated scenarios; it manages monitoring configuration, data collection, and KDD. An autodiscovery function allows retrieving the available monitoring capabilities for each node as well as already configured OPs. This novel functionality is of paramount importance in brownfield scenarios, where the network is already in operation. Particularly, the contribution of this paper is threefold:

- 1) The reference multilayer multiprotocol label switching-transport profile (MPLS-TP)-over-optical network architecture is first introduced in Section II, where intralayer and interlayer reference interfaces are defined. Three types of nodes are considered to create a partial disaggregated data plane: MPLS-TP switches, optical transponder nodes, and optical switches. The reference interfaces allow defining not only a multilayer network topology but also physical intralayer and logical interlayer connectivity. To this respect, reference interfaces provide primitives to monitor both themselves and associated connections, which can be also enriched with measurements from dedicated monitoring devices such as OSAs and OSCs.
- 2) Section III is devoted to monitoring and data analytics, where the proposed architecture is detailed. A hierarchy of monitoring modules bringing distributed and centralized data analytics to the network. A YANG data model specifically designed to manage monitoring, allowing one to separate responsibilities when accessing to node controllers, is proposed. Finally, generic workflows putting together all elements at data/control/management planes are presented.
- 3) To demonstrate the concepts proposed in the paper, three use cases are defined in Section IV, built on top of the workflows presented in Section III. First, a virtual link supported by an optical connection is created. Next, a certain degree of BER degradation in the optical connection is detected, and two different workflows are presented for *i*) multilayer self-configuration and *ii*) optical layer self-healing, including a failure localization procedure based on the combined use of passive and active optical monitoring/telemetry.

The discussion is supported by the experimental demonstration presented in Section V.

II. MULTILAYER NETWORKS: INTERFACES AND MONITORING

As aforementioned, we consider a MPLS-TP-over-optical multilayer network, where the partially disaggregated data plane consists of three elements nodes: the MPLS-TP (L2) switch, the transponder (TP) node, and the optical (L0) switch. For the sake of simplicity, we denote as L2 (electronic) both the Ethernet and the MPLS-TP network sublayers, whereas we denote as L0 the whole optical layer, which includes the channel sublayer (TPs) as well as multiplex/transmission sublayers (L0 switches) [14].

Let us now define the interfaces within the multilayer network architecture. Three types of interfaces are generically considered in every network (sub-) layer [Fig. 1(a)]: *i*) *server interfaces* (SI), providing access to a given network layer; *ii*) *network interfaces* (NI), connecting nodes in the same network layer (NIs can be physical or logical interfaces, as will be described below); *iii*) *client interfaces* (CI), which access the server network layer. Pairs CI-SI between network (sub-) layers, as well as NI-NI in the same network (sub-) layer, are connected through an interlayer or network layer link, which creates the multilayer network topology.

Regarding interlayer adaptation, we assume that it is performed inside the network nodes; specifically, Ethernet to MPLS-TP and MPLS-TP to the optical layer adaptation are performed inside the MPLS-TP switches, whereas optical (gray) to WDM (colored) adaptation is performed in the optical transponders inside TP nodes. Finally, network connections (LSPs) are created between two nodes in the same network (sub-)layer [Fig. 1(b)].

Figure 1(c) shows an example of a multilayer partially disaggregated data plane. The L2 switch includes Ethernet interfaces (L2-SI), being the access of client traffic to the network. Ethernet frames are tagged creating

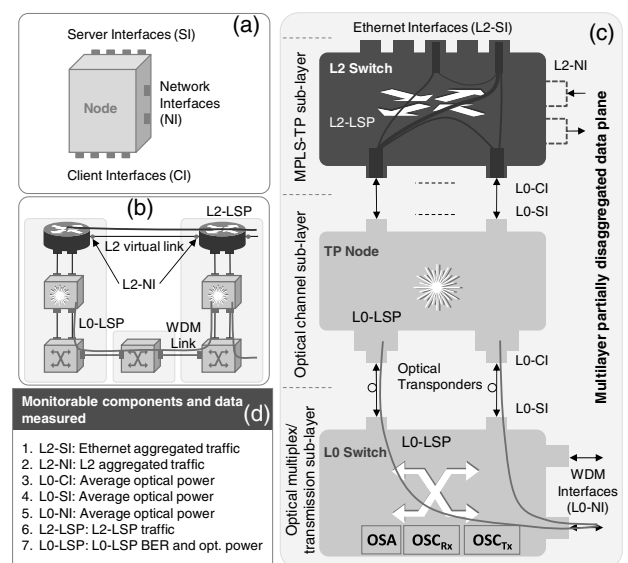


Fig. 1. Multilayer partially disaggregated data plane architecture, interfaces, and monitorable data.

L2-LSPs that are switched and leave the switch through outgoing L2 virtual links (*vlink*). Note that in such cases, end NIs are logical interfaces. To access the network, other Ethernet switches or even IP routers might exist connected to the MPLS-TP switch through Ethernet client interfaces (L2-CI). The TP node integrates optical transponders, so vlinks are supported by optical connections, i.e., L0-LSPs, in the optical layer.

In the presented model, outgoing L2 traffic leaving the L2 node through a given vlink is first aggregated then adapted to the optical layer and sent through a L0-CI toward the TP node. The optical signal is converted back to a flow in the electrical domain and injected into an optical transponder that converts the flow to a WDM signal using a pre-determined modulation format. The WDM signal is sent to the collocated L0 switch, which includes wavelength selective switches (WSS) and forwards the optical signal through the corresponding WDM interface toward a neighboring L0 switch. When the WDM signal arrives to the remote location, the inverse procedure is performed to bring the electronic flow to the egress L2 switch. From the perspective of the L2 layer, the two end L2 switches appear connected through a L2 link (the vlink).

Each of these interfaces and network layer connections are denoted *monitorable components*, as they can provide monitoring data regarding their performance when observation points (OP) are activated on them. The table in Fig. 1(d) details the data that can be retrieved from any of the monitorable components. Specifically: *i*) incoming/outgoing Ethernet traffic can be monitored in L2-SIs; *ii*) L2 traffic for the corresponding LSP can be monitored in L2-LSPs; *iii*) aggregated L2 traffic can be monitored in L2 vlink endpoints (L2-NIs); *iv*) average optical power can be monitored in L0 interfaces; and *v*) BER, power, and other parameters of the corresponding L0-LSP can be monitored at the subcarrier module of the transponders.

In addition to transmission and switching devices, network operators are deploying specific monitoring devices that enrich the available monitoring data, as mentioned above. Examples include OSAs and OSCs at the optical layer. Although other configurations are possible without loss of generality, in this paper we assume that such optical monitoring devices are deployed inside L0 switches to maximize its utilization. In such cases, an optical spectrum can be obtained from L0-NIs, L0-SIs, and L0-LSPs inside L0 switches, whereas in-line BER can be measured in L0-LSPs also inside L0 switches.

In the case of OSAs, the spectrum of the optical links can be acquired and used to analyze the signal of every single L0-LSP, as proposed in [9]. Algorithms analyzing the optical spectrum can find *i*) deviations in the central frequency, which might be a result of a laser drift problem in the TP node, or *ii*) distortions in the shape of the measured signal, which might be a consequence of a filter problem in an intermediate L0 switch. Therefore, a simple failure localization procedure can be implemented by retrieving signal diagnosis information from those node-wide algorithms analyzing the spectrum in every L0 switch along the route of an L0-LSP.

A different failure localization procedure can be implemented using an OSC device. By over-modulating the incoming optical signal in the first L0 switch with an OSC_{TX} and measuring the BER in the OSC_{RX} in every other intermediate L0 switch, an algorithm can follow the evolution of the measured BER and compare it against the expected value obtained using an analytical model (e.g., using [15]). Discrepancies between the measured and the expected BER might help localizing failures. The difference between using OSAs and using OSC resides in the way resources are used for the measurements; OSAs are periodically acquiring the optical spectrum of links, whereas a set of OSCs in L0 switches along the route of a given L0-LSP need to be allocated to perform BER measurements. In fact, OSC-based measurement creates continuous streams of the BER data collected using telemetry.

In conclusion, an MDA system is required to manage networking and monitoring devices, collect measurements, and analyze data.

III. MONITORING AND DATA ANALYTICS

In this section, we present our architecture to support active and passive monitoring, telemetry, and data analytics assuming partially disaggregated multilayer scenarios. A new YANG data model will be needed to facilitate auto-discovery, monitoring, and telemetry operations, including active and passive schemes. Operation of the proposed architecture and data model will be eventually described.

A. Architecture

The proposed monitoring and data analytics architecture shown in Fig. 2 consists of a multi-node agent, referred to as *Hypernode*, managed by the centralized MDA system in the control and management plane. The number of hypernodes may vary depending on the size of the network, geographical extension, and/or any other criteria.

Hypernodes are designed to configure monitoring and telemetry, as well as to collect measurements from one or more nodes in the disaggregated data plane. While each node controller usually controls one single node and exposes one single interface toward the SDN controller, hypernodes are designed to be in charge of monitoring and telemetry of a set of nodes. Hypernodes augment *extended nodes* from [11] and consist of two building blocks, i.e., the *local configuration* module and the *local KDD* module. The local configuration module is in charge of receiving configuration and exposes a RESTCONF-based [16] north-bound interface (NBI) to the MDA system. The NBI is based on a YANG data model, which includes two subtrees: *i*) *applications*, for configuring the applications in the local KDD module inside the hypernode and *ii*) *nodes*, for configuring the underlying nodes. Finally, a number of node adapters (one per node) are used to implement the specific protocols exposed by every node controller for node configuration and for monitoring data collection.

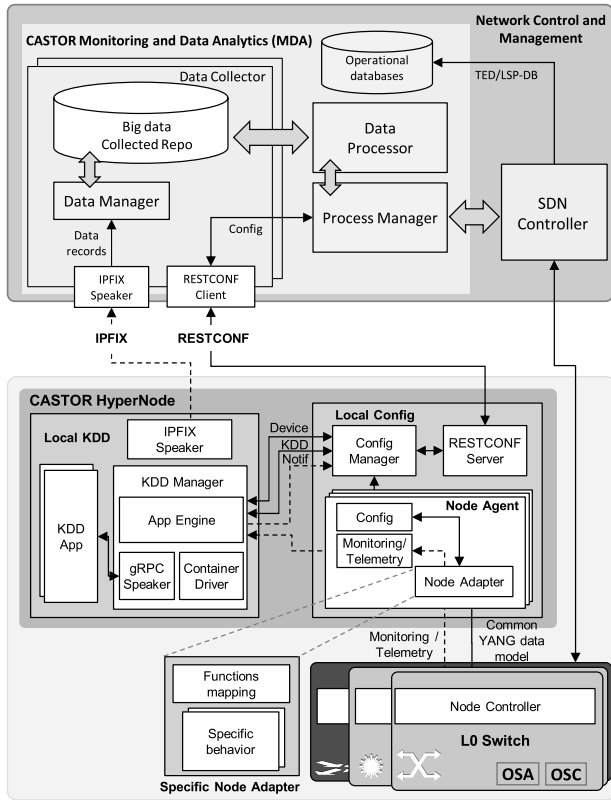


Fig. 2. CASTOR architecture and interfaces.

Although we assume that node controllers' NBIs are based on a common YANG data model, different protocols for configuration, monitoring, and telemetry might be considered (e.g., NETCONF [17], IPFIX [18], and gRPC-based protocols [19] for configuration, monitoring, and telemetry, respectively, NETCONF for everything, or any other combination). For this very reason, node agents include bespoke node adapters implementing specific protocols and function mapping for the underlying node controller.

Regarding the local KDD module, its original scope in the *extended nodes* was to apply data analytics to the measurements received from the nodes focused on the KDD process. Output of the data analytics procedure was forwarded to the MDA system to implement network-wide control loops. Two types of messages are supported: *i*) IPFIX-based monitoring messages, including processed monitoring samples (i.e., values are averaged over the selected monitoring period, e.g., 15 min); and *ii*) through asynchronous notifications using the RESTCONF NBI. Regarding telemetry, measurements are locally processed by specific KDD processes in the hypernode to reduce data exchange with the MDA system. Note that telemetry measurements might be taken in a sub-second basis, so analysis is performed locally in the hypernode, and results can be conveyed to the MDA system for decision making.

Similar to the hypernode, the MDA system has been extended from a similar module, as described in [11]. Extensions are related to the configuration of the monitoring and telemetry, defined in the new YANG data models.

In addition, operational database (e.g., topology and LSP databases) synchronization from the SDN controller is supported. Full database synchronization as well as topology or LSP update notifications are now supported. With operational databases synchronized, the CASTOR system is able not only to collate measurements from the nodes (as in [11]) but also to correlate them with the route of a LSP for failure localization purposes. Therefore, more sophisticated procedures can be developed correlating measurements with topology and LSP data.

B. YANG Data Model and Monitoring Data

Figure 3 presents the structure of the proposed YANG data model that is assumed to be supported by every node controller. The model defines two differentiated subtrees with the objective to separate configuration from monitoring and telemetry responsibilities; the SDN controller is focused on the *configuration* subtree, whereas hypernodes are mainly focused on the *monitoring* subtree. The configuration subtree includes every programmable or monitorable component in the node, while the monitoring subtree includes monitoring capabilities and Ops, and it is specifically designed to facilitate autodiscovery by hypernodes.

A key element in the model is the *component*, representing any configurable or monitorable element in the node and is locally identified by its *component-id*. Nodes are assumed to feed a data store compliant with this model during bootstrapping, whereas dissemination toward hypernodes and the SDN controller can be made per update notifications or by polling. Although components under the configuration and monitoring subtrees are related, we relax such condition to allow obtaining only the monitoring subtree during the autodiscovery process. In fact, the configuration subtree is not stored in the hypernode to avoid synchronization issues; whenever configuration of some component is required by a KDD application, the local configuration module retrieves it directly from the node controller.

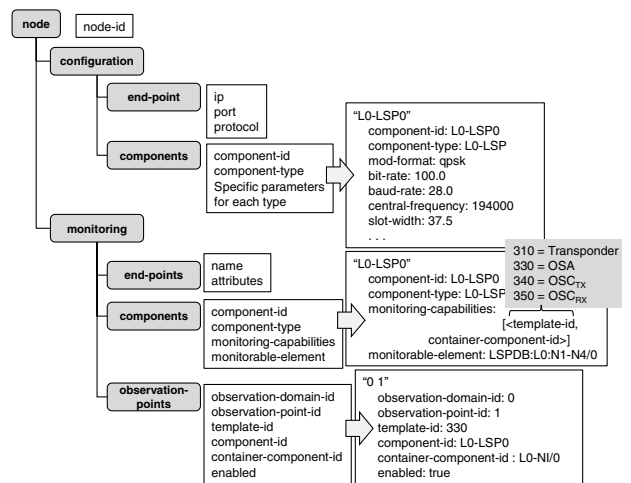


Fig. 3. Proposed YANG data model and examples.

Monitoring/telemetry in monitorable components can be activated by creating and enabling OPs and deactivated by disabling and/or deleting OPs. Monitorable components include a list of tuples with supported monitoring template identifiers (*template-id*) and the container, e.g., an interface, where measurements will be done (*container-component-id*). Monitoring templates represent different measurement methods, so each OP is associated to one single monitoring template and interface tuple. We use the *template-id* to identify the monitoring method used for the measurements, as well as their data structure; specifically, we define the following identifiers: *i*) *template-id* 310 identifies monitoring of BER and optical power in the transponders; *ii*) *template-id* 330 identifies optical spectrum monitoring performed by OSAs, which send sequences of tuples (frequency, optical power); *iii*) *template-id* 340 identifies OSC_{TX} devices that report heartbeat/keepalive messages; and *iv*) *template-id* 350 identifies OSC_{RX} devices that send estimated BER values.

The monitoring subtree also includes every existing OP (enabled, e.g., performing measurements or disabled). Specifically, the ability to enable and disable existing OPs supports allocation of monitoring devices in active monitoring schemes. Nodes are expected to allocate monitoring resources by translating tuples (*component-id*, *template-id*, *container-component-id*) into device-specific commands. Observation domain/point identifiers are internal identifiers used by KDD processes inside hypernodes to isolate different data sets, and they are included in monitoring messages. Multiple OPs can be defined for a single component, each reporting monitored data formatted as per the selected supported template.

Figure 3 presents an example for a lightpath (L0-LSP): configuration parameters include, but are not limited to, modulation format, bit/ baud rate, and the allocated frequency slot specified by *central-frequency* and *slot-width* parameters. In the monitoring subtree, components include two relevant attributes: *i*) *monitorable-element* (“what”), which is used for correlation purposes between the operational databases in the SDN controller and the MDA system; and *ii*) *monitorable-capabilities* containing a set of pairs with the supported *template-id* (“how”), representing the monitoring method that such component supports in the specific container component, identified by its *container-component-id* in the network node (“where”). For illustrative purposes, all the templates have been included in Fig. 3, but only those supported by the component in the network node will be advertised. Finally, an OP is currently active; specifically, L0-LSP0 is being monitored in interface L0-NI/0 using an OSA.

C. Operation

Figure 4 presents the three basic workflows supported: *i*) *autodiscovery*; *ii*) *LSP setup and monitoring/telemetry activation*; and *iii*) *KDD process*. These basic workflows will be part of the use cases presented and experimentally demonstrated in the following sections. Additionally, the workflows give insight into the internal relations in the

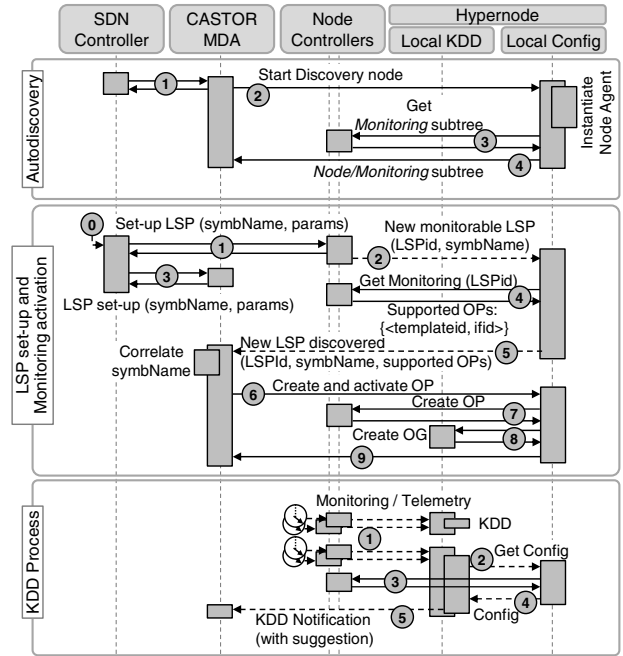


Fig. 4. Basic workflows supported by hypernodes.

hypernodes, which are represented by the local KDD and the local config internal blocks. Each exchanged message is identified with a number that is relative to the workflow it belongs to; note that numbers are reset in every workflow.

The autodiscovery workflow is initiated by the CASTOR MDA system at startup, when the operational databases are retrieved from the SDN controller, as well as by the SDN controller after a new node has been added to its topology database. In the second case, a notification is sent to the CASTOR MDA system (message 1), which updates its topology database and requests the corresponding hypernode to start the autodiscovery process (2); specifically, the *node-id* and the parameters for configuration, monitoring, and telemetry are needed to start the autodiscovery process. Upon receipt of the start autodiscovery message, the local configuration block in the hypernode instantiates a new node agent with the specific adapter, which connects to its peer node controller to obtain the current *Monitoring subtree* status (3), which is sent back to the MDA system (4).

The LSP setup and monitoring/telemetry activation workflow are triggered by a request to the SDN controller (message 0). After computing the specific routing algorithm (e.g., the routing, modulation, and spectrum allocation-RMSA-algorithm for L0-LSPs [12]), the SDN controller sends configuration messages (1) to the node controllers along the route of the LSP specifying a set of parameters, e.g., the spectrum allocation in the case of a L0-LSP; additionally, the LSP symbolic name is added to the configuration message so it can be used as an identifier in future messages. When the LSP has finally been set up, a notification is sent to the CASTOR MDA system (3) containing the route of the LSP, configuration parameters, and

symbolic name; however, the LSP needs to be also discovered by the hypernodes so as to ensure proper data synchronization. To that end, node controllers notify the corresponding hypernodes about resources allocated for a new LSP, identified by its local identifier and the identifier received from the SDN controller (2). In this case, the hypernode needs to obtain from the node controller the monitoring subtree for the specific LSP, and the node controller replies with the set of pairs $\langle template-id, container-component-id \rangle$ supported for such LSP (4). Next, the hypernode sends a notification toward the MDA system announcing that a new LSP has been discovered. The notification includes the SDN controller identifier for the LSP, the local identifier, and the monitoring capabilities (5), and is correlated with the one received from the SDN controller. Next, the MDA system creates and activates a subset of OP for the LSP (6). Requests to create and activate OPs are sent to the node controllers (7). In the local KDD, OP handlers, called observation groups (OG), are created inside a KDD application to collect, analyze, and aggregate measures from a set of OPs (8); once an OG has been created, add/remove OPs messages can be issued. The hypernode replies to the MDA system once the creation and activation processes are completed (9).

Once OPs are activated, monitoring/telemetry data are received and can be analyzed to discover patterns. The KDD process workflow starts when new measurements are received in the hypernode (message 1); monitoring data records are analyzed by the corresponding process in a KDD application and, when a certain pattern is detected, the process asks the local config block in the hypernode (2) to interrogate the node controller about the current configuration of some components (3). The KDD process analyzes the received component configuration to discern whether it can be modified and notifies the MDA system about the discovered pattern; a suggestion of parameter tuning might be included in the notification as a result of the configuration analysis (5). With such notification, the MDA system can make decisions, including suggesting a reconfiguration to the SDN controller.

IV. USE CASES

In this section, we consider the architecture proposed above and apply the defined workflows on a multilayer scenario. Here, we assume that the topology has already been discovered and systems are synchronized. The use cases use a L2 vlink supported by an L0-LSP. The evolution of the BER in the supporting L0-LSP is monitored in the hypernode and, in the case that BER degradation is detected, a notification is sent to the MDA system [20]. Two different alternatives are considered once the notification is received in the MDA system: *i) L0/L2 self-configuration*, entailing modification of the modulation format for the underlying L0-LSP and, consequently, reducing the available bandwidth in the vlink, as well as in every L2 LSP using such vlink. Note that different policies can be applied in this case, such as rerouting L2-LSPs to keep the allocated bandwidth [21] or even creating a parallel L0-LSP to keep invariant the vlink capacity. *ii) L0 self-healing*, where a

failure localization procedure using OSAs and OSCs is executed to localize the optical link responsible for the high BER. Upon completion, the L0-LSP can be restored avoiding that optical link.

A. L2 Virtual Link Creation

The workflow for a L2 vlink creation supported by an L0-LSP (WF1) is presented in Fig. 5. The workflow consists of two phases: first a L0-LSP between two end TP nodes is set up (messages 0-7); then, the vlink is created between the two end L2 switches (messages 8-14). Monitoring is automatically activated for the L0-LSP and the vlink.

The L0-LSP setup and monitoring activation phase are based on the generic LSP setup workflow described in the previous section. The only particularities are that the involved nodes are TPs and L0-switches. The parameters included in messages (1) and (3) are those related to the optical layer, such as central frequency (fc), slot width (bw), modulation format, etc. Regarding monitoring capabilities, the TP node controller reports that the L0-LSP can be monitored in the transponders (*template-id* 310), whereas the L0 switch controller reports that it supports OSA monitoring (*template-id* 330) as well as OSC active monitoring (*template-id* 340 and 350). Finally, the MDA system decides to create and activate OPs related to transponders and OSAs; note that active monitoring entails resource reservation, and it will be activated on-demand when needed for failure localization.

The vlink creation and monitoring activation phase is also based on the workflow described in the previous section. In this case, the parameters needed for the creation are the two end points and the capacity (bw). In this case, the involved nodes are the end L2 switches that report L2 traffic monitoring capabilities (*template-id* 256), and OPs are automatically created and activated.

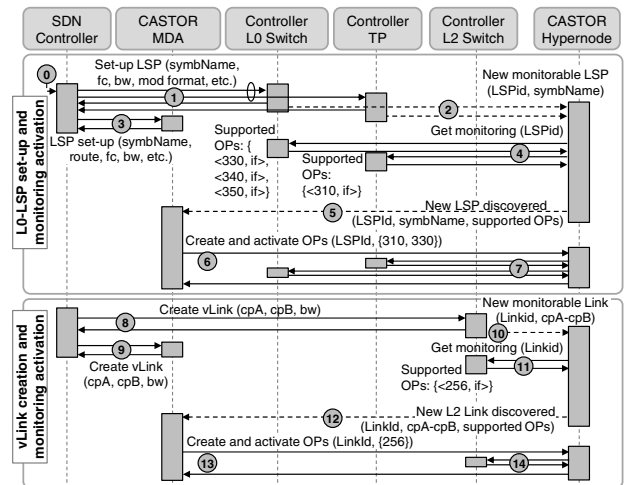


Fig. 5. WF1: L2 virtual link creation supported by an L0-LSP.

B. L0/L2 Self-Configuration

Once the vlink is in operation, L2 LSPs can be set up using its available capacity, and monitoring data records are collected and analyzed by the hypernodes. Figure 6 presents the workflow (WF2) for this use case, which is divided into two phases: BER degradation detection (messages 1-3 in Fig. 6) and multilayer reconfiguration (messages 4-10).

The BER evolution is analyzed for the L0-LSP (1), so a workflow similar to the KDD process described in the previous section is followed. In the case that BER degradation is detected, the KDD process requests the current configuration to the TP node controller, which in this case includes the optical parameters configured at setup time, including the modulation format (2). Let us assume that the QAM-16 modulation format is currently used for the L0-LSP; thus, changing the modulation format to QPSK would improve the quality of transmission, which would eventually reduce the BER. In consequence, the KDD process issues a notification to the MDA system reporting the high BER detected and includes the change to the QPSK modulation format as recommended action (3).

In the case that, upon the notification reception, the MDA system follows the recommended action, sending a notification to the SDN controller recommending the change in the modulation format of the L0-LSP (4). In view of that, the SDN controller follows an internal policy to first reduce the capacity of any L2-LSP using the affected vlink (traffic shaping is reconfigured in the ingress L2 switch of the LSPs) (5), then it reduces the capacity of the vlink in the end L2 switches (7), and it finally changes the modulation format of the supporting L0-LSP in the end TP nodes (9). The SDN controller informs about the changes to the MDA system (messages 6, 8, and 10).

C. L0 Self-Healing

Let us assume now that, when the notification with the BER degradation detection (message 3 in Fig. 6) is received

in the MDA system, the recommended action to change the modulation format is not followed, and a procedure to localize the failure affecting the L0-LSP is triggered instead. Figure 7 presents the workflow (WF3) for this use case, which includes the failure localization procedure at the optical layer (4-13) and the L0-LSP rerouting, excluding the identified optical link (14-20).

The failure localization procedure is triggered when the notification with BER degradation detection is received in the MDA system. We assume now that, in contrast with the previous use case, a specific procedure for failure localization is available as one of the data analytics processes in the data processor module in the MDA system (see Fig. 2). The failure localization procedure requests signal diagnosis to every hypernode in charge of an L0 switch in the route of the L0-LSP. A KDD process is in charge of analyzing the optical spectrum of the LSP using the received OSA data records from the local L0 switch (4); the received answers allow the MDA system to determine the cause of the failure (e.g., filter misconfiguration, laser drift, or other) as well as to localize the failure [9]. To confirm the localization of the failure, the procedure might decide to use the OSC monitoring devices to measure the BER along the route of the L0-LSP. To that end, it requests the hypernodes with OSC devices in the route of the L0-LSP to reserve an OSC_{TX} (in the first node) or an OSC_{RX} (in the rest of the nodes) by creating OPs with *template-id* 340/350 in the corresponding container components (5); in turn, hypernodes request the creation of the OPs (reserving the OSC devices) to the local L0 switch controller (6). When all the confirmations of the device reservation arrive in the MDA system, OP activation is performed (7 and 8). Note that, in this workflow, OP creation and activation are dissociated because monitoring devices need to be first reserved. Once OPs are activated, telemetry is used to send a continuous data stream with BER measurements to the hypernode (9); the corresponding KDD process analyzes

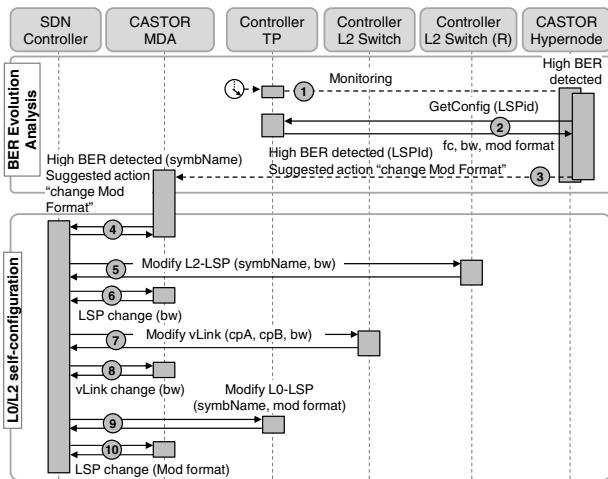


Fig. 6. WF2: High BER detection and L0/L2 self-configuration.

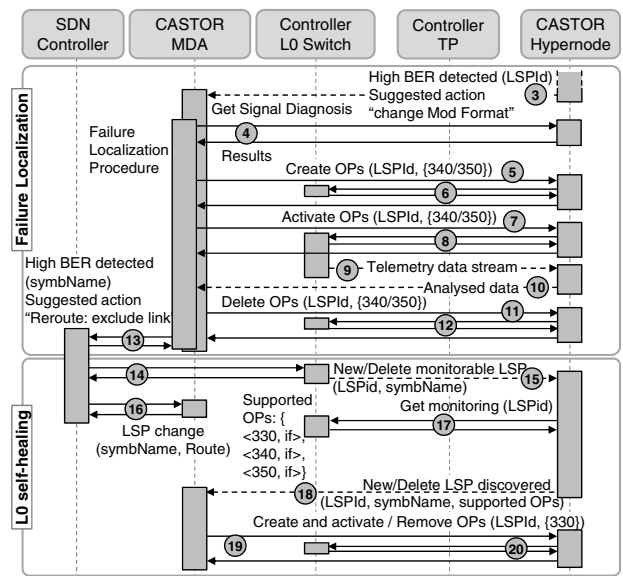


Fig. 7. WF3: Failure localization and L0 self-healing.

the telemetry data and reports the MDA system (10). The measured BER is compared with the expected BER computed analytically; this allows finding divergences between expected and measured BER, thus localizing the failure. Once the failure is localized, the OPs are deleted, so the OSC devices are released (11 and 12). Assuming that both localization methods return the same failure localization, e.g., in a link due to high noise in an optical amplifier, a notification is sent to the SDN controller informing about BER degradation in the L0-LSP with the suggested action to reroute the LSP excluding the identified link (13).

Upon receipt of the notification, the SDN controller follows the recommended action and finds an alternative route for the degraded L0-LSP. Note that more than one L0-LSP might be affected by the degradation in the link (this actually depends on the BER thresholds configured for each L0-LSP), so a planning tool might be used to compute optimal rerouting for the set of affected LSPs [21]. Once the route of the LSPs are determined, a similar workflow as the one for LSP setup is followed, and OPs are created in the L0 switches entering in the new route and deleted from those leaving the route.

V. EXPERIMENTAL RESULTS

This section is devoted to the experimental demonstration of the workflows defined in Section IV. The experiments have been carried out in the UPC's SYNERGY testbed. The scenario consists of a CASTOR MDA system and the ONOS SDN controller [22] in the control and management plane and a number of network locations, each with a hypernode in charge of one local L2 switch, TP node, and L0 switch. CASTOR MDA and hypernodes, as well as node controller emulators, have been developed in Python, where node controller emulators expose a NETCONF NBI based on the YANG data model presented in Section III.B. Communication between the MDA system and ONOS is performed through an ad hoc bidirectional REST API.

A. Starting Procedure and Autodiscovery

Before workflows in Section IV can be demonstrated, both ONOS and node controllers are started, so that the network is in operation. After starting hypernodes in the network locations and the MDA system in the control and management plane, an autodiscovery workflow (see Section III.C) needs to be carried out for synchronization purposes. This workflow is run each time a new node is added to ONOS. Figure 8 presents a capture with the

Source	Destination	Protocol	Info
① ONOS	CASTOR_MDA	HTTP	POST /rest/newNode HTTP/1.1
② CASTOR_MDA	ONOS	HTTP	HTTP/1.1 200 OK (application/json)
③ CASTOR_MDA	CASTOR_HN1	HTTP	POST /rest/rpc-autodiscovery HTTP/1.1
④ CASTOR_HN1	X1_Ctrl1	NetConf	NetConf
⑤ X1_Ctrl1	CASTOR_HN1	NetConf	NetConf
⑥ CASTOR_HN1	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)

Fig. 8. Exchanged messages for X1 L0 switch autodiscovery.

exchanged messages for X1 L0 switch autodiscovery; message numbering is consistent with that in Fig. 4. Remote procedure calls (RPC) are used in RESTCONF and NETCONF interfaces to facilitate data access, as shown in messages (2-3).

B. Virtual Link Creation

Figure 9 shows the exchanged messages for the creation of a L2 vlink between switches in locations 1 and 7; message numbering is consistent with that in Fig. 5, where only messages exchanged for location 1 are shown. Although notifications through the RESTCONF interface (messages 5 and 12 in Fig. 9) are not properly shown in the capture (Wireshark cannot parse HTTP event-streaming), their contents are properly decoded by the MDA system.

OSAs are continuously acquiring the optical spectrum in the corresponding outgoing link, and, when an OP is activated for a L0-LSP (message 7 in Fig. 9), its spectrum is encoded and sent in IPFIX messages using *templateId* 330, which includes a *subTemplateList* field to convey a structured list of data records [18]. This field is convenient to encode optical spectrum data because it allows including an ordered list of tuples $\langle \text{frequency}, \text{power} \rangle$; Fig. 10 presents the details of IPFIX messages for OSA monitoring.

The complete workflow WF1 was executed in less than 1 s, considering message exchange and control and management plane time, i.e., excluding device configuration.

C. L0/L2 Self-Configuration

Figure 11 shows the exchanged messages for WF2, when high BER is detected in the L0-LSP supporting the vlink between locations 1 and 7; message numbering is consistent with that in Fig. 6, where messages exchanged for

Source	Destination	Protocol	Info
① ONOS	X1_Ctrl1	NetConf	NetConf
② X1_Ctrl1	ONOS	NetConf	NetConf
③ X1_Ctrl1	CASTOR_HN1	NetConf	NetConf
④ ONOS	TP1_Ctrl1	NetConf	NetConf
⑤ TP1_Ctrl1	ONOS	NetConf	NetConf
⑥ TP1_Ctrl1	CASTOR_HN1	NetConf	NetConf
⑦ ONOS	CASTOR_MDA	HTTP	POST /rest/LSP/set-up HTTP/1.1
⑧ CASTOR_MDA	ONOS	HTTP	HTTP/1.1 200 OK (application/json)
⑨ CASTOR_HN1	TP1_Ctrl1	NetConf	NetConf
⑩ TP1_Ctrl1	CASTOR_HN1	NetConf	NetConf
⑪ CASTOR_HN1	X1_Ctrl1	NetConf	NetConf
⑫ X1_Ctrl1	CASTOR_HN1	NetConf	NetConf
⑬ CASTOR_HN1	CASTOR_MDA	HTTP	[TCP segment of a reassembled PDU]
⑭ CASTOR_MDA	CASTOR_HN1	HTTP	POST /rest/OP/activate HTTP/1.1
⑮ CASTOR_HN1	TP1_Ctrl1	NetConf	NetConf
⑯ TP1_Ctrl1	CASTOR_HN1	NetConf	NetConf
⑰ CASTOR_HN1	X1_Ctrl1	NetConf	NetConf
⑱ X1_Ctrl1	CASTOR_HN1	NetConf	NetConf
⑲ CASTOR_HN1	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)
⑳ ONOS	L2Sw1_Ctrl1	NetConf	NetConf
㉑ L2Sw1_Ctrl1	ONOS	NetConf	NetConf
㉒ ONOS	CASTOR_MDA	HTTP	POST /rest/vLink/set-up HTTP/1.1
㉓ CASTOR_MDA	ONOS	HTTP	HTTP/1.1 200 OK (application/json)
㉔ L2Sw1_Ctrl1	CASTOR_HN1	NetConf	NetConf
㉕ CASTOR_HN1	L2Sw1_Ctrl1	NetConf	NetConf
㉖ L2Sw1_Ctrl1	CASTOR_HN1	NetConf	NetConf
㉗ CASTOR_HN1	CASTOR_MDA	TCP	[TCP segment of a reassembled PDU]
㉘ CASTOR_MDA	CASTOR_HN1	HTTP	POST /rest/OP/activate HTTP/1.1
㉙ CASTOR_HN1	L2Sw1_Ctrl1	NetConf	NetConf
㉚ L2Sw1_Ctrl1	CASTOR_HN1	NetConf	NetConf
㉛ CASTOR_HN1	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)

Fig. 9. Exchanged messages for WF1.


```
'exportTime': '2017-Dec-27 18:53:21',
'observationDomainId': 0,
'dataSets': [{
  'setId': 330,
  'records': [{
    'observationPointId': 1,
    'subTemplateList': {
      'semantics': 'ordered',
      'templateId': 10000,
      'records': [{
        'frequencyTHz': 193.98,
        'rxPowerdBm': -60.01492
      }, ..., {
    } ] } ] } ] }
```

Fig. 10. Details of OSA monitoring.

locations 1 and 8 (ingress L2 switch of an L2-LSP using the vlink) are shown.

Figure 12 plots the evolution of the overall L2 traffic through the vlink as well as the individual traffic of each (two in this case) L2-LSP traversing it; the capacity of the vlink is also shown for reference. As a result of traffic shaping in the ingress of the L2-LSPs, its traffic is limited to half of the original one, which reduces the total traffic through the vlink. The capacity of the vlink is eventually reduced.

As mentioned in Section IV.B, we are considering an operational scheme in which all L2-LSPs are proportionally reshaped upon L0-LSP capacity changes. It is worth mentioning that bandwidth of any L2-LSPs cannot fall behind

the minimum bandwidth guarantees (sometimes referred to as “committed information rate” [CIR]) according to the associated service-level agreement (SLA) [23]. Otherwise, other mechanisms like L2 rerouting must be used to avoid CIR/SLA violation. In this regard, commercial devices allow different configurations of L2-LSPs depending on how vlink bandwidth is divided among them: absolute (e.g., 10 Gb/s) or proportional (e.g., 10% of vlink capacity). In our use case, we consider that bandwidth of L2-LSPs is defined following the absolute syntax, which is the reason why L2-LSP reconfiguration (5) is performed before vlink reconfiguration (7), to avoid the sum of CIRs to exceed total vlink capacity.

D. L0 Self-Healing

Figure 13 shows the exchanged messages for WF3; message numbering is consistent with that in Fig. 7, where only messages exchanged for location 1 are shown. This workflow uses active and passive monitoring techniques for failure localization purposes. Figure 14(a) presents the route of the L0-LSP established from TP-1 to TP-7 supporting the vlink between L2Sw1 and L2Sw7. In particular, the architecture of L0 switches is shown in Fig. 14(a), where pools of OSAs, each acquiring the whole C-band of an outgoing L0-NI, as well as pools of OSCs with OSC_{TX} and OSC_{RX}, are available.

When the MDA system is notified upon BER degradation detection in TP1 (message 3 in Fig. 13), it decides to run the failure localization procedure before notifying the SDN controller. To that end, the hypernodes along the route of the L0-LSP are interrogated to obtain a diagnosis of the spectrum of the optical signal received from the OPs activated in the OSAs (4). Figures 14(b) and 14(c) present two different failure cases, where graphs plotting the optical spectrum were captured in X1, X2, X5, and

Source	Destination	Protocol	Info
TP1_Ctrl	CASTOR_HN1	CFLOW	IPFIX flow (52 bytes) [Data:300]
TP1_Ctrl	CASTOR_HN1	CFLOW	IPFIX flow (52 bytes) [Data:300]
CASTOR_HN1	TP1_Ctrl	NetConf	NetConf
TP1_Ctrl	CASTOR_HN1	NetConf	NetConf
CASTOR_HN1	CASTOR_MDA	TCP	[TCP segment of a reassembled PDU]
CASTOR_MDA	ONOS	HTTP	POST /rest/LSP/performance HTTP/1.1
ONOS	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)
ONOS	L2Sw8_Ctrl	NetConf	NetConf
L2Sw8_Ctrl	ONOS	NetConf	NetConf
ONOS	CASTOR_MDA	HTTP	POST /rest/LSP/config HTTP/1.1
CASTOR_MDA	ONOS	HTTP	HTTP/1.1 200 OK (application/json)
ONOS	L2Sw1_Ctrl	NetConf	NetConf
L2Sw1_Ctrl	ONOS	NetConf	NetConf
ONOS	CASTOR_MDA	HTTP	POST /rest/vLink/config HTTP/1.1
CASTOR_MDA	ONOS	HTTP	HTTP/1.1 200 OK (application/json)
ONOS	TP1_Ctrl	NetConf	NetConf
TP1_Ctrl	ONOS	NetConf	NetConf
ONOS	CASTOR_MDA	HTTP	POST /rest/LSP/config HTTP/1.1
CASTOR_MDA	ONOS	HTTP	HTTP/1.1 200 OK (application/json)

Fig. 11. Exchanged messages for WF2.




Fig. 12. Virtual link capacity and L2 traffic evolution.

Source	Destination	Protocol	Info
CASTOR_HN1	CASTOR_MDA	TCP	[TCP segment of a reassembled PDU]
CASTOR_MDA	CASTOR_HN1	HTTP	POST /rest/LSP/diagnosis HTTP/1.1
CASTOR_HN1	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)
CASTOR_MDA	CASTOR_HN1	HTTP	POST /rest/OP/create HTTP/1.1
CASTOR_HN1	X1_Ctrl	NetConf	NetConf
X1_Ctrl	CASTOR_HN1	NetConf	NetConf
CASTOR_HN1	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)
CASTOR_MDA	CASTOR_HN1	HTTP	POST /rest/OP/activate HTTP/1.1
CASTOR_HN1	X1_Ctrl	NetConf	NetConf
X1_Ctrl	CASTOR_HN1	NetConf	NetConf
X1_Ctrl	CASTOR_HN1	NetConf	NetConf
CASTOR_HN1	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)
CASTOR_HN1	CASTOR_MDA	TCP	[TCP segment of a reassembled PDU]
CASTOR_MDA	CASTOR_HN1	HTTP	POST /rest/OP/delete HTTP/1.1
CASTOR_HN1	X1_Ctrl	NetConf	NetConf
X1_Ctrl	CASTOR_HN1	NetConf	NetConf
CASTOR_HN1	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)
ONOS	CASTOR_MDA	HTTP	POST /rest/LSP/performance HTTP/1.1
ONOS	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)
ONOS	X1_Ctrl	NetConf	NetConf
X1_Ctrl	ONOS	NetConf	NetConf
X1_Ctrl	CASTOR_HN1	NetConf	NetConf
ONOS	CASTOR_MDA	HTTP	POST /rest/LSP/config HTTP/1.1
CASTOR_MDA	ONOS	HTTP	HTTP/1.1 200 OK (application/json)
CASTOR_HN1	X1_Ctrl	NetConf	NetConf
X1_Ctrl	CASTOR_HN1	NetConf	NetConf
CASTOR_HN1	CASTOR_MDA	TCP	[TCP segment of a reassembled PDU]
CASTOR_MDA	CASTOR_HN1	HTTP	POST /rest/OP/create HTTP/1.1
CASTOR_HN1	X1_Ctrl	NetConf	NetConf
X1_Ctrl	CASTOR_HN1	NetConf	NetConf
CASTOR_HN1	CASTOR_MDA	HTTP	HTTP/1.1 200 OK (application/json)

Fig. 13. Exchanged messages for WF3.

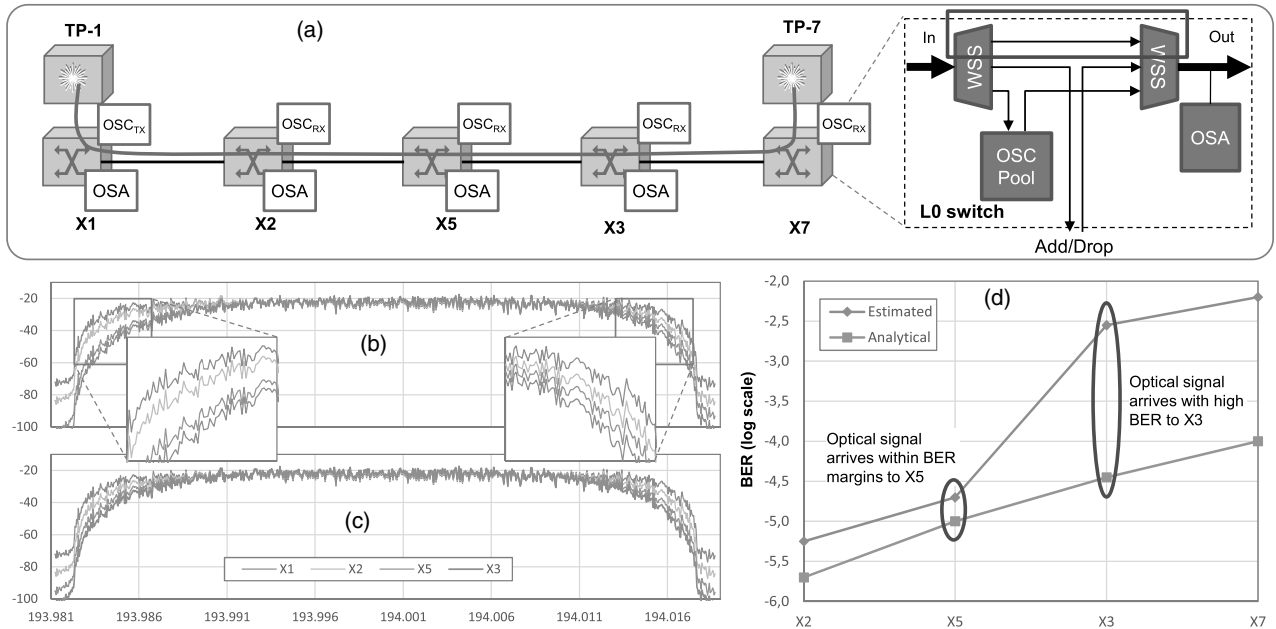


Fig. 14. Illustrative scenario for failure localization: (a) physical route for a reference lightpath; (b) OSA traces for a filter-shift event in X5; (c) OSA traces for an in-line amplifier problem in X5-X3; and (d) OSC measures for the same event as in (c).

X3 L0 switches. Figure 14(b) reproduces the case where a filter shift failure happens in X5; note that the signal is asymmetrically filtered before the OSA in X5, thus pointing out a misconfiguration of a WSS in such node [WSSs within the red box in Fig. 14(a)]. In this case, a specific procedure could be followed to solve the WSS misconfiguration. Figure 14(c) reproduces the case where no filtering problem is detected in the intermediate nodes. In such a case, OSAs monitoring cannot localize the failure, and OSC-based active monitoring can be used as an alternative.

Recall that OSC-based active monitoring requires us to reserve OSC resources in the intermediate L0 switches (i.e., X2, X5, X3, and X7); in particular, one OSC_{TX} is strictly needed. Assuming that OPs are created (OSC resources are reserved) and activated (5-8), a gRPC-based telemetry data stream is created from every L0 switch controller toward the local hypernode with the estimated BER from the measurements performed on the low-speed signal (9); OSC_{RX} produces BER data records at a subsecond rate. Telemetry data are analyzed by the corresponding KDD process in the hypernode to produce average BER values that are sent to the MDA system (10). Figure 14(d) presents an overall chart with BER estimations received from OSC_{RX} and expected BER values computed using analytical models. In view of the different slopes in the segment X5 to X3, it can be concluded that something happens in such an optical link (e.g., high noise generated by some optical amplifier).

VI. CONCLUSIONS

An architecture for autonomic operation of multilayer disaggregated networks has been presented; the

architecture is based on the combination of SDN controllers and MDA systems. On the one hand, the SDN controller enables programmability of the network, not only in terms of provisioning of connectivity services but also in terms of reconfiguration. On the other hand, the MDA system provides a distributed platform for network intelligence processes based on monitoring data analytics. Together, these two entities may enable the deployment of production-ready autonomic (self-managed) multilayer networks. Note that hypernodes enable such data analysis distribution, which brings many benefits such as *i*) the reduction of detection times for failures and anomalies, *ii*) the reduction of data conveyed to the centralized MDA system, *iii*) the improvement in the scalability of the control and management architecture, etc.

Here, we consider a disaggregated data plane composed of three types of network devices: MPLS-TP switches, optical transponder nodes, and optical switches. Aiming at identifying candidate components in which performance can be measured, a reference multilayer MPLS-TP-over-optical network architecture has been introduced, where intralayer and interlayer reference interfaces have been defined. Although network devices are usually able to perform measurements on those reference interfaces, dedicated monitoring devices can be used to enrich such measurements.

Besides, the architecture is complemented with a YANG data model specifically designed to manage monitoring, focused on separating responsibilities for monitoring configuration and network configuration when accessing the node controllers. General workflows that combine the above elements were presented, including autodiscovery, LSP setup and OP creation and activation, and monitoring and KDD processes.

To demonstrate the concepts presented in this paper, three workflows were defined based on the generic ones and experimentally demonstrated. First, a virtual link supported by an L0-LSP was created. Upon detection of the BER in the L0-LSP, two different workflows were presented for multilayer self-configuration and optical self-healing, including a failure localization procedure based on the combined use of passive and active optical monitoring/telemetry using OSA and OSCs. It was demonstrated that OSA and OSC monitoring/telemetry provide complementary views, and their combined use increased accuracy of failure localization.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Commission for the H2020-ICT-2016-2 METRO-HAUL project (G.A. 761727), from the AEI/FEDER TWINS project (TEC2017-90097-R), and from the Catalan Institution for Research and Advanced Studies (ICREA). This work was presented in part at OFC [24,25].

REFERENCES

- [1] O. González de Dios, R. Casellas, F. Paolucci, A. Napoli, L. Gifre, A. Dupas, E. Hugues-Salas, R. Morro, S. Belotti, G. Meloni, T. Rahman, V. López, R. Martínez, F. Fresi, M. Bohn, S. Yan, L. Velasco, P. Layec, and J. P. Fernandez-Palacios, "Experimental demonstration of multi-vendor and multi-domain elastic optical network with data and control interoperability over a PAN-European test-bed," *J. Lightwave Technol.*, vol. 34, pp. 1610–1617, 2016.
- [2] R. Nejabati, S. Peng, and D. Simeonidou, "Optical network democratization," *Philos. Trans. R. Soc. A*, vol. 374, p. 20140443, Aug. 2016.
- [3] V. López and L. Velasco, *Elastic Optical Networks: Architectures, Technologies, and Control*, 1st ed., Springer, 2016.
- [4] M. Bjorklund, "YANG—A data modeling language for the network configuration protocol (NETCONF)," IETF RFC 6020, 2010.
- [5] OpenROADM [Online]. Available: <https://www.openroadm.org>.
- [6] OpenConfig [Online]. Available: <https://www.openconfig.net>.
- [7] M. Behringer, M. Pritikin, S. Bjarnason, A. Clemm, B. Carpenter, S. Jiang, and L. Ciavaglia, "Autonomic networking: definitions and design goals," IETF RFC 7575, 2015.
- [8] D. Geisler, R. Proietti, Y. Yin, R. Scott, X. Cai, N. Fontaine, L. Paraschis, O. Gerstel, and S. J. B. Yoo, "Experimental demonstration of flexible bandwidth networking with real-time impairment awareness," *Opt. Express*, vol. 19, pp. B736–B745, 2011.
- [9] A. P. Vela, B. Shariati, M. Ruiz, F. Cugini, A. Castro, H. Lu, R. Proietti, J. Comellas, P. Castoldi, S. J. B. Yoo, and L. Velasco, "Soft failure localization during commissioning testing and lightpath operation [Invited]," *J. Opt. Commun. Netw.*, vol. 10, pp. A27–A36, 2018.
- [10] A. P. Vela, M. Ruiz, and L. Velasco, "Distributing data analytics for efficient multiple traffic anomalies detection," *Comput. Commun.*, vol. 107, pp. 1–12, 2017.
- [11] 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 [Invited]," *J. Lightwave Technol.*, vol. 36, pp. 135–141, 2018.
- [12] L. Velasco, A. P. Vela, F. Morales, and M. Ruiz, "Designing, operating and re-optimizing elastic optical networks [Invited Tutorial]," *J. Lightwave Technol.*, vol. 35, pp. 513–526, 2017.
- [13] L. Velasco and M. Ruiz, *Provisioning, Recovery and In-Operation Planning in Elastic Optical Networks*, 1st ed., Wiley, 2017.
- [14] "Architecture of optical transport networks," ITU-T Recommendation G.872, 2017.
- [15] P. Poggiolini, "The GN model of non-linear propagation in uncompensated coherent optical systems," *J. Lightwave Technol.*, vol. 30, pp. 3857–3879, 2012.
- [16] A. Bierman, M. Bjorklund, and K. Watsen, "RESTCONF protocol," IETF RFC 8040, 2017.
- [17] R. Enns, M. Bjorklund, J. Schoenwaelder, and A. Bierman, "Network configuration protocol (NETCONF)," IETF RFC 6241, 2011.
- [18] B. Claise, B. Trammell, and P. Aitken, "Specification of the IP flow information export (IPFIX) protocol," IETF RFC 7011, 2013.
- [19] "gRPC: a high performance, open-source universal RPC framework" [Online]. Available: <https://grpc.io>.
- [20] A. P. Vela, M. Ruiz, F. Fresi, N. Sambo, F. Cugini, G. Meloni, L. Poti, L. Velasco, and P. Castoldi, "BER degradation detection and failure identification in elastic optical networks," *J. Lightwave Technol.*, vol. 35, pp. 4595–4604, 2017.
- [21] A. Castro, R. Martínez, R. Casellas, L. Velasco, R. Muñoz, R. Vilalta, and J. Comellas, "Experimental assessment of bulk path restoration in multilayer networks using PCE-based global concurrent optimization," *J. Lightwave Technol.*, vol. 32, pp. 81–90, 2014.
- [22] Open Network Operating System (ONOS) [Online]. Available: <https://onosproject.org/>.
- [23] S. Alvarez, *QoS for IP/MPLS Networks*, 1st ed., Cisco, 2006.
- [24] L. Velasco, L. Gifre, and J.-L. Izquierdo-Zaragoza, "Toward multilayer disaggregated node telemetry and local decision making," in *Optical Fiber Conf.*, 2018.
- [25] L. Gifre, J.-L. Izquierdo-Zaragoza, B. Shariati, and L. Velasco, "Experimental demonstration of active and passive optical networks telemetry," in *Optical Fiber Conf.*, 2018.