Monitoring plane architecture and OAM Handler

Nicola Sambo, Filippo Cugini, Andrea Sgambelluri, and Piero Castoldi

***Abstract*—The increase of new services such as content distribution and inter-data center connectivity poses new requirements on the control and management of next generation networks. The Application-Based Network Operations (ABNO) is emerging as a paradigm integrating control and management functions for different types of services. Key functionalities lie in the ABNO Operation, Administration, Maintenance (OAM) Handler which is responsible for the verification and the actual maintenance of network services following specific service level agreements (SLAs). The OAM Handler collects and correlates alerts about potential problems and triggers actions to preserve services in case of failures or degradations.**

In this paper, a hierarchical monitoring architecture is presented. The architecture is proposed by building on the OAM Handler functionalities. Monitoring information can be filtered and correlated at each hierarchical level guaranteeing high scalability of the management plane. Actions can be taken at each hierarchical layer based on the type of failure and of affected service. Measurements on a commercial systems are carried on to identify the generated alarms upon failure. Then, simulations show the high scalability achieved by the proposed architecture.

***Index Terms*—Software defined network (SDN), ABNO, PCE, OAM.**

I. Introduction

CORE and metro networks are evolving to support ultra-high rate communication systems enabling elastic adaptation and optimization of transmission parameters while guaranteeing high reliability [2]. In the recent years, optical networking has experienced important advances both at the physical and the control layers. Emerging data plane technologies are going to support the increasing traffic demand with transponders enabling high bit rates such as 1 Tb/s, possibly looking at downscaling power consumption and costs per bit, and offering high spectral efficiency, thus increasing network life [3]–[7]. At the control plane layer, Software Define Networking (SDN) is emerging as an architecture to remotely set network devices, thus programming transmission characteristics (such as offered bit rate) and switching [8]–[12]. However, while data and control planes have experienced such advances, the innovations in the management plane have not followed these trends yet [13], leaving management as complex and decoupled at the IP and the optical layers. Network management includes the Operation, Administration, Maintenance (OAM) functions required to monitor and interpret the measurements on services and devices, and to recover from faults or degradations [13]. Because of such problematics, recently, the Application-Based Network Operations (ABNO) architecture [14] has been proposed within the IETF as a solution integrating control and management functionalities and orchestrating both applications and service provisioning at the client layer. Such architecture may have the prospects to offer to the operators an agreed way to properly control and manage networks and applications.

ABNO includes the OAM Handler, a key functional block to verify the actual quality of transmission (QoT) at the data plane layer and the service level according to specific agreements (SLAs) at the application layer. The OAM Handler is responsible for: i) receiving alerts about potential problems; ii) correlating them (e.g., for fault localization); iii) triggering other components of the ABNO, such as the Path Computation Element (PCE), to take actions to preserve the services that are interested by the fault or the degradation. In such a scenario, monitoring techniques are also fundamental. Again, the data plane layer offers new ways of monitoring. Current coherent systems permit direct monitoring of the optical connections through the digital signal processing into the receiver itself [2]. If needed, the OAM Handler has to trigger proper actions (e.g., adaptation of transmission parameters, re-routing) to react against soft or hard failures (e.g., link degradations or faults, respectively) which degrade QoT and, in turn, service level. While some functionalities of the ABNO architecture have been studied in the literature [15]–[19], the OAM Handler still requires to be deeply investigated, also considering scalability issues, e.g. given that the amount of alarms generated by an optical network may be huge. Moreover, alarms in next generation optical networks may also become much more frequent because of system margin reduction. Indeed, vendors and operators are now oriented to reduce system margins that account for aging, model inaccuracies, cross-phase modulation and other degradations [20]–[22]. Such margins cause the underestimation of the optical reach, thus, increases the number of regenerators in a network and in turns the costs. A reduction of system margins can decrease the number of installed regenerator, but, on the other hand, a more frequent generation of alarms may occur. Indeed, as described in [22], more conservative thresholds should be adopted to trigger alarm generation and, in this case, the number of alarms can become even much more higher than in current networks.

In this paper, we present the hierarchical monitoring architecture proposed within the framework of the EU project ORCHESTRA, by expanding the work of [1]. The monitoring infrastructure consists of virtual monitoring entities and agents with the OAM Handler at the root of the hierarchical infrastructure. Possible actions taken by the monitoring plane and by ABNO are summarized depending on the type of failure. The scalability of such architecture is evaluated. To this purpose,

Manuscript received December 17, 2015.

N. Sambo (email: nicola.sambo@sssup.it), A. Sgambelluri, and P. Castoldi are with Scuola Superiore Sant'Anna, Via Moruzzi 1, Pisa, Italy. F. Cugini is with CNIT, Via Moruzzi 1, Pisa, Italy.

This work was supported by the EC through the Horizon 2020 ORCHESTRA project (grant agreement 645360).

This paper is an extended version of the work presented in [1].

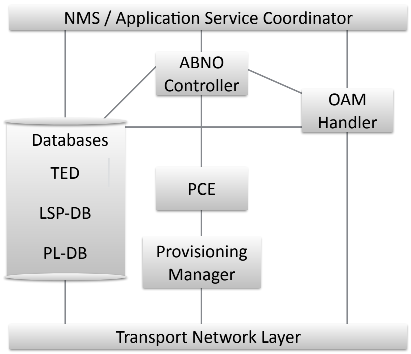

Figure 1. ABNO architecture.

first, measurements are carried out on a commercial system to recognize the amount and the type of generated alarms upon link degradation. Then, such measurements are used as input in the simulations, showing the improved scalability achieved by the proposed hierarchical architecture with respect to a centralized OAM Handler. Moreover, reactions to soft- or hard-failures are summarized considering new emerging technologies at the data plane.

The reminder of the paper is organized as follows. Sec. II presents the ABNO architecture proposed within IETF and related research works. Sec. III introduces the proposed hierarchical monitoring architecture and OAM Handler. Sec. IV summarizes possible actions and workflow taken by the monitoring plane and ABNO, depending on the fault and on the affected service. Sec. V presents a study on the scalability of the proposed architecture. Finally, Sec. VI concludes the paper.

II. ABNO ARCHITECTURE AND RELATED WORK

The ABNO architecture is composed of a set of functional modules enabling the control and the management of applications and network. It provides on-demand and application-specific reservation of network connectivity, reliability, resources, and management for several network applications [14].

Fig. 1 describes the ABNO architecture considering the main modules relevant for this work. For a complete ABNO architecture, the reader can refer to [14].

Hereafter, the main ABNO functional modules are described:

- The Network Management System (NMS) in ABNO may be considered as an entity that issues service requests to the ABNO Controller.
- Application Service Coordinator can be the alternative to NMS for requesting services on behalf of applications, such as a program on a host or server, a tool requiring end-to-end connection, or a virtual private network.
- The databases considered in the ABNO architecture are the Traffic Engineering Database (TED) storing traffic engineering information (such as the range of spectrum occupied per link) and the Label Switched Path DataBase (LSP-DB) containing information on the state of LSPs such as traversed interfaces, bit-rate, occupied frequency slot [23]. Besides these two databases included by the IETF RFC 7491 [14], we also consider a database storing physical layer information, named Physical Layer DataBase (PL-DB). Such database stores all the accessible information related to the physical layer, which are known from the devices' data sheets or accessible through monitoring. PL-DB entry may refer to an LSP or to a network device like a fiber. PL-DB may include bit error rate (BER) of lightapaths, amplifiers' information (e.g., noise figure), fiber information (e.g., attenuation, chromatic dispersion, average differential group delay, and effective area), and so on.
- The ABNO controller coordinates control and management functions in ABNO. The ABNO Controller governs the behavior of the network based on network conditions and application requirements.
- The Path Computation Element (PCE) is devoted to path computation (in some cases it also performs wavelength or spectrum assignment). The PCE may receive these requests from the ABNO Controller or from network elements themselves. The PCE performs computation based on the view of network topology stored in the TED. More advanced computations may be provided by a Stateful PCE that enhances the TED with LSP-DB or PL-DB [24]. Finally, an Active PCE allows additional *active* functionalities such as making provisioning requests to set up new services or to modify active services.
- The OAM Handler is devoted to detect faults or degradations and take actions to react to problems in the network. The OAM Handler is responsible for receiving alerts from the network about potential problems, for correlating them, and for triggering other components of the system to take actions to preserve or recover the affected services. The OAM Handler interacts with monitoring and testing devices. We also assume that the OAM Handler can also ask to the monitoring system to get specific monitored physical parameters on demand.
- The Provisioning Manager is responsible for the provision of requests, interfacing with the control plane implemented in the network or directly programming individual network devices.

Regarding research activities, some works in the literature have investigated ABNO [15]–[19]. In the experiment in [15], the ABNO controller efficiently manages two PCEs depending on the service to be accomplished: a PCE is invoked for service provisioning, while another PCE includes more complex algorithms for network re-optimization. In [16], ABNO orchestrates end-to-end multi-layer and multi-domain provisioning, considering optical packet and circuit switching. Authors in [17] demonstrate ABNO for a multilayer (IP and optical) and multi-vendor scenario. In [18], multi-layer path computation, connection provisioning and topology discovery/transfer are experimentally demonstrated. Authors in [19] present a testbed to assess re-optimization algorithms and ABNO workflow. OAM functionalities and architecture within ABNO still need to be deeply investigated.

III. The hierarchical monitoring architecture and the OAM Handler

An optical network (with fixed- or flexible-grid) is considered to be equipped with monitors, that are assumed for lightpaths (LPs), links, and nodes, as in Fig. 2. LP monitors are assumed integrated in the digital signal processing of each lightpath coherent receiver (e.g., pre-forward-error-correction bit error rate — pre-FEC BER — monitors), while, as an example, power monitors can be assumed for links and nodes. The proposed hierarchical monitoring architecture is presented in Fig. 3. Each entity provides the same OAM functions: i.e., collecting and correlating alerts and triggering actions to preserve services. Each one is responsible for a specific set of LPs or for a sub-set of nodes or links (e.g., nodes belonging to a given network area). Exchange of monitored information can be triggered due to physical layer degradations or faults and, in this case, information flows towards the upper layers. Each layer correlates and filters the received information efficiently sending less amount of monitoring information to an upper layer toward the OAM Handler. Alternatively, the OAM Handler on behalf of the ABNO controller can ask to the monitoring plane to get specific monitoring information (e.g., optical signal to noise ratio) that can be used to have a more updated PL-DB or to perform advanced path computations for new requests [25].

The architecture of a monitoring entity (e.g., $\mathrm{LP}_{\mathrm{group1}}$ at *level 1* of Fig. 3) is shown in Fig. 4. The Agent is responsible to disseminate monitoring information. Although not shown, the Manager at *level i* is connected to several monitoring entities of the *level i-1*. The Manager correlates and processes all the information coming from its associated agents at the *level i-1* and triggers actions for recovery. A *level i* Agent passes the information to a Manager at the *level i+1*. At the root, the Manager coincides with the OAM Handler.

The monitoring architecture is divided in three parts, one per monitored type of element: thus, it includes monitoring elements responsible for LPs, for links, and for nodes. Regarding LP monitoring architecture, the *level 0* is responsible for OAM of single LPs. Similarly, regarding link and node, the *level 0* is responsible for OAM of single links or nodes, respectively. Upper layers are responsible for a set of LPs (e.g., LPs starting from the same ingress node), links, or nodes. While the level in the hierarchy increases, the monitoring entities are responsible for a larger set of LPs or links and nodes, up to being able to correlate information of all LPs, links, and nodes.

As an example, by referring to Fig. 2 where four LPs are active, the *level 0* has four active LP monitors. By assuming an amplifier malfunction in link A-B, several alarms will be generated for the LP connecting A-B and for LP connecting A-C. Such alarms will be sent to the upper layer for correlation. The *level 1* is composed of monitoring entities, each one associated to a group of LPs. For example, each monitoring entity at the *level 1* can group all the lightpaths starting from the same ingress node (e.g., LP_{group1} is responsible for lightpaths A-B and A-C, while LP_{group2} for lightpaths F-E and F-D). Thus, such alarms will be sent to LP_{group1} of the *level 1* for correlation. This way, by correlating this

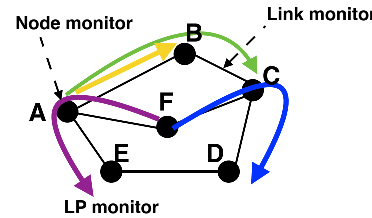

Figure 2. Example of network and monitors.

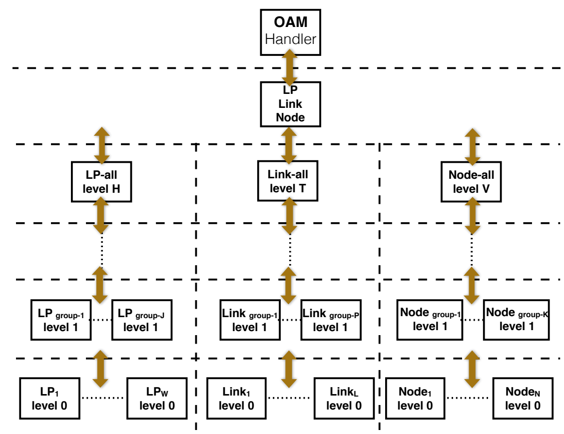

Figure 3. Hierarchical monitoring architecture.

information, a problem can be identified in the segment A-B. Then, LP *level 2* can group all the lightpaths whose ingress node belongs to a specific region of the network and so on up to a generic *level H*. This scheme is applied also to link and node monitors: e.g., *level 1* of link monitoring entities can correlate and filter information coming from links belonging to a specific area in the network. Once information is passed to the upper layers, the information is correlated so that failed or degraded network elements can be identified and the amount of alarms can be reduced while reaching the root of the architecture, i.e, the OAM Handler. This way, the architecture provides high scalability and the OAM Handler is not overloaded, nor the monitoring entities at each layer given that they are responsible for a limited set of network elements.

Regarding the LP part of monitoring architecture, each intermediate hierarchical level, besides collecting and correlating alarms and monitoring information, can take actions on the set of LPs for which is responsible. In case of protected LP, if the monitoring entity at a *level 0* reveals an excessive BER increase, such monitoring entity can automatically trigger the switch to the protection path. Then, upper layers are informed, e.g. for fault localization, and, finally, the OAM Handler will verify that the actual *service availability* is satisfied according to SLA. In case of not-protected lightpaths, a given level will trigger the restoration. In this case, the OAM Handler will communicate to the ABNO controller the needs of dynamic re-routing, and the ABNO controller will call PCE for path computation. Finally, the Provisioning Manager will take care of the new set up for dynamic restoration. In the next section, failures and actions are discussed based on the type of service.

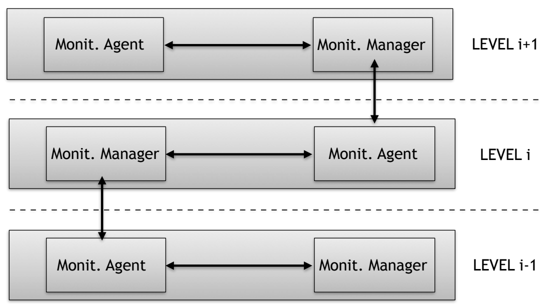

Figure 4. Monitoring entity at a generic level.

IV. Failures and actions

Monitoring Managers of Fig. 4 are responsible to trigger alarms in case of failures or degradations, identify the affected services, and take actions to preserve the affected services. Failures are classified in two types:

1) *hard*: they consist in the damage of a link or a node so that the connectivity is lost. As an example, they can be due to a fiber cut.
2) *soft*: they imply a transmission performance degradation, not necessary causing a service outage. As an example, if the pre-FEC BER threshold of a transponder is 2×10^{-3} and the soft failure only causes an increase on the BER up to 10^{-3}, no outage is experienced. Instead, an outage is experienced if the pre-FEC BER passes above the threshold. A soft failure may be due to a malfunction or aging of network devices like amplifiers, equalizers, fibers and so on. It can also be due to an excessive cross-talk with other channels. Such failures may be relevant when margins against aging or crosstalk or cross-phase modulation are under dimensioned [20], [21].

Thus, soft or hard failures impact services in a different way and it must be also considered that services may belong to different classes of traffic, thus requiring different actions for preservation. In order to provide indications on the workflow within the monitoring and the ABNO architectures, some classes of traffic are here considered:

1) *Gold*: protected lightpath
2) *Silver*: not protected lightpath and bit-rate reduction is not admitted after failure recovery (i.e., 100% bit rate recovery)
3) *Bronze*: not protected lightpath and bit-rate reduction is admitted after failure recovery (i.e., $\leq 100\%$ bit rate recovery)

A brief discussion is here provided on how a bit reduction could be experienced. A soft failure affecting a service can be overcome by switching to a more robust modulation format. This operation implies a reduction of the bit rate. As an example, with polarization multiplexing 16 quadrature amplitude modulation (PM-16QAM) at 28 Gbaud and 12% FEC, 200 Gb/s net rate is obtained; with polarization multiplexing quadrature phase shift keying (PM-QPSK) at 28 Gbaud and 12% FEC, 100 Gb/s net rate is obtained. Thus, if a reconfiguration from PM-16QAM to PM-QPSK is required, to keep the same information rate, the introduction of a new sub-carrier or the set up of a new lightpath must be performed. Similarly, a soft failure could be overcome by adapting the overhead (FEC) [26], thus increasing code redundancy. Also this operation may imply a reduction of the bit rate. Indeed, if a transponder works at the maximum baud rate and code bits are increased upon soft failure, net rate decreases.

Tab. I aims at summarizing possible actions and work flow taken once the monitoring architecture reveals an hard or a soft failure, depending on the class of traffic. In case of hard failure and protected lightpath (gold class), a switch to the protection path should be performed to recover the traffic. Thus, the monitoring layer identifying an affected lightpath triggers the switching to the protection path. Such event should be notified up to the OAM Handler to take records about possible loss of traffic taking care that SLAs are satisfied. In case of not-protected lightpath (silver and bronze), the monitoring plane informs the OAM Handler, which interacts with the ABNO controller. Thus, dynamic restoration can be triggered.

In case of soft failure such as amplifier multifunction, a BER increase above the pre-FEC threshold may be experienced. If the affected lightpath is protected, switch to the protection lightpath can be triggered. If the affected lightpath is silver class, FEC adaptation is taken into consideration, evaluating if the increase of the redundancy implies or not an increase of the ITU-T frequency slot width (i.e., the passband of traversed filters). If not, FEC adaptation is triggered by the monitoring layer identifying the affected service and then OAM Handler is notified. If the FEC adaptation implies a slot width increase, the OAM Handler has to be notified and the TED has to be consulted to check spectrum resource availability and, eventually, re-routing has to be ordered to preserve the whole traffic. In this case, the OAM Handler notifies the ABNO controller which requests the PCE to compute the recovery path. This also happens if FEC adaptation is not enough to provide transmission robustness and re-routing has to be performed. If the affected service is a bronze class, actions that imply a bit rate reduction could be taken. Thus, if FEC adaptation is expected to successfully recover from failure, this action is taken. Then, the ABNO controller will take or not actions related to the lost bit rate. Also the change to a more robust modulation format could be applied. Again, the ABNO controller will take or not actions related to the lost bit rate. If no modulation format or FEC adaptation are expected to recover from failure, the ABNO controller will trigger re-routing. Note that actions not involving the consultation of ABNO TED, LP-DB, and PL-DB or the invocation of PCE for re-routing – such as the switch to a protection path, modulation format adaptation, and, in many cases, FEC adaptation – are triggered by the layer detecting the failure (e.g., *level 0*) before notifying the upper layers in the monitoring hierarchy. Different actions can be taken if the soft failure is due to inter-channel interference (XPM). In this case, for silver and bronze classes, ABNO controller could order to shift lightpath in frequency, or power re-equalizzation, or re-routing. Indeed, channel interferences such as XPM typically depend on the distance in frequency of the interfering channels (the close the channels the largest the impact of XPM) and on the optical power (the highest the power of the interfering channel the largest the impact of XPM).

V. Performance evaluation

The performance of the proposed hierarchical monitoring architecture is evaluated in terms of scalability. First, measurements are carried out on a commercial system to identify the generated alarms in case of link degradation. Then, such alarms are used as input for simulations in a Spanish topology composed of 30 nodes and 55 bi-directional links [27].

A. Generated alarms on measurements

To identify the alarm generated upon link degradation considering commercial cards, we performed measurements on the testbed in Fig. 5. A lightpath has been actived at frequency 195.30THz. A variable optical attenuator (VOA) has been placed on the line interconnecting the two nodes and has been used to emulate link degradation. When we started attenuating the power over the optical link, the power at the receiver decreased: at -28 dBm some alarms appeared. From the receiver, 8 alarms were generated due to the link loss. They were mainly related to the OTU2/ODU2/ODU0 (4 alarms), while others to OCh layer(2 alarms), OMS(1 alarm), OTS (1 alarm). From the transmitter, 5 alarms appeared: 1 for OTS layer, 1 for OMS and the others 3 were OTU2/ODU2/ODU0. Thus, a single link degradation generated 13 alarms related to a single lightpath. In the experiment, no hierarchical monitoring architecture has been implemented.

B. Simulations

Simulations have considered two management architectures: i) the proposed hierarchical monitoring architecture; ii) a centralized OAM receiving all monitoring information and correlating them. Then, two scenarios have been considered: a) link hard-failure, where all lightpaths traversing the failed link are disrupted and each one is source of alarms; b) soft-failure, where performance of a network element – such as an amplifier – are degraded thus causing the OSNR decrease of traversing lightpaths. In the latter scenario, some lightpaths may suffer from this degradation while others not. For that, we assumed several values of OSNR penalty due to the degradation and the model for coherent PM-QPSK in [28]. OSNR degradation may imply a BER increase over a threshold (10^{-3} in the simulations) or not. If the threshold is exceeded, alarms are generated, otherwise not. Link hard- or soft-degradation is randomly generated on a single link and the generated number of alarms is counted. Results are recorded until the confidence interval of 5% at 95% confidence level is achieved. The network operates with Poisson traffic. The holding time of each connection is exponentially distributed. The hierarchical monitoring architecture consists in *level 0 , 1, 2,* and OAM Handler. *Level 0* is composed of monitors. *Level 1* is composed of functional entities, each one correlating monitoring information of LPs starting from the same ingress node. Thus, at *level 1*, there is a monitoring entity for each network node. *Level 2* is composed by a single entity which gathers all the info coming from *Level 1*. The root is the OAM Handler at *level 2*. In both centralized and hierarchical scenarios, the number of affected lightpaths generating the alarms is the same, but in the centralized scenario, 13 alarms per affected lightpath are sent to the OAM Handler (see Subsection A) while, in the hierarchical scenario, 13 alarms per affected lightpath are sent to the *level 1* and each entity at level 1 is responsible for a subset of lightpaths (thus, providing higher scalability).

Table II
Number of received alarms per monitoring entity at each *Level* in case of link hard failure.

	Level 1	Level 2	OAM Handler
Centralized	not present	not present	420.03
Hierarchical	47.97	9.2	1

Simulations in case of hard failure are reported in Tab. II that shows the average number of alarms received by the OAM Handler in the centralized scenario or by each monitoring entity at each level in the hierarchical scenario. In the centralized scenario, an average of 420 alarms reaches the OAM Handler which has to process all the information, identify the degraded network element, and take actions for the lightpaths interested by the degradation. On the contrary, with the hierarchical architecture, 48 alarms in average are received by each monitoring entity at *level 1*. At *level 2*, around 9 alarms are received, while 1 reaches the OAM Handler.

Simulations in case of soft failure are also reported. The three classes presented in Sec. IV are assumed. In particular, if the impacted service belongs to silver or bronze class, the ABNO controller has to invoke PCE for computation: indeed, for a silver-class service, a restoration path is computed for the whole 100 Gb/s impacted rate, while for a bronze-class service, a restoration path is computed for the half rate service (see Tab. I). Fig. 6 shows the average number of received alarms by the centralized OAM or by the hierarchical levels at varying the OSNR penalty introduced by the soft failure. The number of alarms increases with the OSNR penalty since the BER of more lightpaths passes above the threshold (10^{-3} in the simulations) thus, generating alarms. The figure shows that the hierarchical architecture is a more scalable solution with respect to the centralized OAM Handler (e.g., at level 2 of the hierarchy, the number of alarms is almost constant with the OSNR penalty). The number of alarms in the centralized solution increases much more rapidly than the ones at each *levels 1* and *2*. Fig. 7 shows the amount of traffic to be re-routed, for which PCE has to compute alternative paths. Such traffic belongs to silver and bronze classes. Traffic to be re-routed increases with OSNR degradation because more services are impacted by the soft failure. All impacted silver traffic is re-routed. Bronze traffic to be re-routed is lower than silver because part of impacted bronze traffic is handled with operations like modulation format adaptation. Half bronze traffic has to be re-routed, successfully completing the recovery.

Table I
FAILURE DETECTED BY OAM AND ACTIONS.

Failure	**Impact**	**Impacted class**	**Actions**
Hard: link or node	Loss of connection	Gold	• switch to protection path • notify the upper layer in the hierarchy
Hard: link or node	Loss of connection	Silver or Bronze	• notify the upper layer in the hierarchy • ABNO controller orders re-routing
Soft: e.g., amplifier, equalizer, fiber aging	BER increase	Gold	• switch to protection path • notify the upper layer in the hierarchy
Soft: e.g., amplifier, equalizer, fiber aging	BER increase	Silver	Possibilities: 1) FEC adaptation is expected to successfully recover from failure A. not implying slot width increase: • FEC is adapted • notify the upper layer in the hierarchy B. implying slot width increase: • notify the upper layer in the hierarchy • ABNO controller orders re-routing 2) FEC adaptation is NOT expected to successfully recover from failure • notify the upper layer in the hierarchy • ABNO controller orders re-routing
Soft: e.g., amplifier, equalizer, fiber aging	BER increase	Bronze	Possibilities: 1) FEC adaptation is expected to successfully recover from failure • FEC is adapted • notify the upper layer: if bit rate is lost ABNO could take actions on lost bit rate 2) Change of modulation format is expected to successfully recover from failure • Change of modulation format • notify the upper layer in the hierarchy: ABNO could take action on lost bit rate 3) Change of modulation format or FEC is NOT expected to successfully recover from failure • notify the upper layer in the hierarchy • ABNO controller orders re-routing
Soft: inter-channel interference	BER increase	Gold	• switch to protection path • notify the upper layer in the hierarchy
Soft: inter-channel interference	BER increase	Silver or Bronze	• notify the upper layer in the hierarchy • Depending on resource availability, ABNO controller orders one of the following actions: - shift in frequency - power re-equalization - re-routing

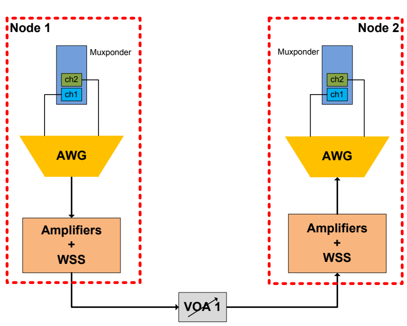

Figure 5. Testbed for alarm generation.

VI. CONCLUSIONS

This paper presented the hierarchical monitoring architecture proposed within the EU ORCHESTRA project. OAM Handler functionalities (collections, correlation, and interpretation of monitoring information, and actions for recovery) are spread into several hierarchical layers, enabling to confine sets of monitored physical parameters within specific levels in the hierarchy. This results in the limitation of the management plane overload. Measurements have been performed to identify the generated alarms in a commercial system. Results have been used to implement simulations in order to evaluate the performance of the hierarchical architecture against a centralized OAM Handler. The proposed hierarchical architecture guarantees high scalability.

REFERENCES

[1] N. Sambo, F. Cugini, A. Sgambelluri, and P. Castoldi, "Hierarchical monitoring architecture and OAM Handler," in *Proc. of ECOC*, 2015.

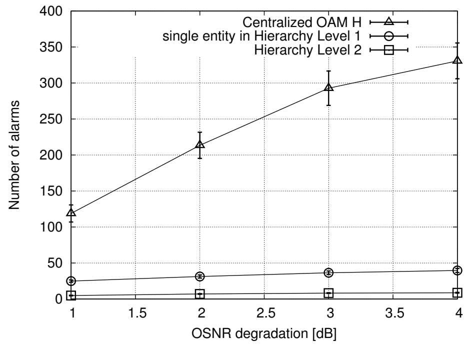

Figure 6. Number of received alarms vs. OSNR degradation in case of link soft failure.

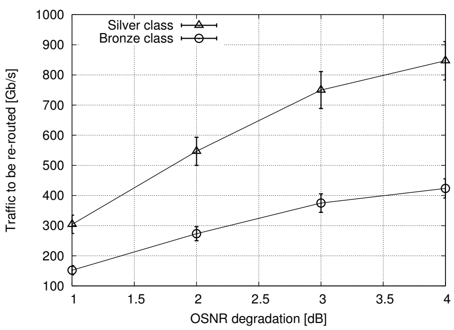

Figure 7. Traffic to be re-routed (i.e., traffic for which ABNO controller invokes PCE for computation upon failure) vs. OSNR degradation in case of link soft failure.

[2] A. Napoli, M. Bohn, D. Rafique, A. Stavdas, N. Sambo, L. Poti, M. Nolle, J. Fischer, E. Riccardi, A. Pagano, A. Di Giglio, M. Moreolo, J. Fabrega, E. Hugues-Salas, G. Zervas, D. Simeonidou, P. Layec, A. D'Errico, T. Rahman, and J.-P. Gimenez, "Next generation elastic optical networks: The vision of the european research project IDEALIST," *Communications Magazine, IEEE*, vol. 53, no. 2, pp. 152–162, Feb 2015.

[3] N. Sambo, P. Castoldi, A. D'Errico, E. Riccardi, A. Pagano, M. Moreolo, J. Fabrega, D. Rafique, A. Napoli, S. Frigerio, E. Salas, G. Zervas, M. Nolle, J. Fischer, A. Lord, and J.-P. Gimenez, "Next generation sliceable bandwidth variable transponders," *Communications Magazine, IEEE*, vol. 53, no. 2, pp. 163–171, Feb 2015.

[4] M. Svaluto Moreolo, J. Fabrega, L. Nadal, F. Vilchez, V. Lopez, and J. Fernandez-Palacios, "Cost-effective data plane solutions based on OFDM technology for flexi-grid metro networks using sliceable bandwidth variable transponders," in *Optical Network Design and Modeling, 2014 International Conference on*, May 2014, pp. 281–286.

[5] F. Guiomar, S. Amado, A. Carena, G. Bosco, A. Nespola, and A. Pinto, "Transmission of PM-64QAM over 1524 km of PSCF using fully-blind equalization and Volterra-based nonlinear mitigation," in *Optical Communication (ECOC), 2014 European Conference on*, Sept 2014.

[6] Y. Loussouarn, E. Pincemin, M. Song, S. Gauthier, Y. Chen, and Z. Shengqian, "400 Gbps real-time coherent Nyquist-WDM DP-16QAM transmission over legacy G.652 or G.655 fibre infrastructure with 2 dB margins," in *Optical Fiber Communications Conference and Exhibition (OFC), 2015*, March 2015, pp. 1–3.

[7] J. Cai, H. Batshon, M. Mazurczyk, H. Zhang, Y. Sun, O. Sinkin, D. Foursa, and A. Pilipetskii, "64QAM based coded modulation transmission over transoceanic distance with > 60 Tb/s capacity," in *Optical Fiber Communications Conference and Exhibition (OFC), 2015*, 2015.

[8] L. Liu, R. Munoz, R. Casellas, T. Tsuritani, R. Martinez, and I. Morita, "Openslice: An OpenFlow-based control plane for spectrum sliced elastic optical path networks," in *Optical Communications (ECOC), 2012 38th European Conference and Exhibition on*, Sept 2012, pp. 1–3.

[9] L. Liu, T. Tsuritani, I. Morita, R. Casellas, R. Martinez, and R. Munoz, "Control plane techniques for elastic optical networks: GMPLS/PCE vs OpenFlow," in *Globecom Workshops (GC Wkshps)*, Dec 2012.

[10] N. Sambo, G. Meloni, F. Paolucci, F. Cugini, M. Secondini, F. Fresi, L. Poti, and P. Castoldi, "Programmable transponder, code and differentiated filter configuration in elastic optical networks," *JLT*, vol. 32, no. 11, June 2014.

[11] Y. Yoshida and et al., "First international SDN-based network orchestration of variable-capacity OPS over programmable flexi-grid EON," in *Proc. of OFC*.

[12] M. Khaddam, L. Paraschis, and J. Finkelstein, "SDN multi-layer transport benefits, deployment opportunities, and requirements," in *Optical Fiber Communications Conference and Exhibition (OFC), 2015*, 2015.

[13] A. Martinez, M. Yannuzzi, V. Lopez, D. Lopez, W. Ramirez, R. Serral-Gracia, X. Masip-Bruin, M. Maciejewski, and J. Altmann, "Network management challenges and trends in multi-layer and multi-vendor settings for carrier-grade networks," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 4, pp. 2207–2230, Fourthquarter 2014.

[14] D. King and A. Farrel, "A PCE-based architecture for application-based network operations," IETF RFC 7491.

[15] L. Gifre, F. Paolucci, L. Velasco, A. Aguado, F. Cugini, P. Castoldi, and V. Lopez, "First experimental assessment of ABNO-driven in-operation flexgrid network re-optimization," *Lightwave Technology, Journal of*, vol. 33, no. 3, pp. 618–624, Feb 2015.

[16] R. Munoz, R. Vilalta, R. Casellas, R. Martinez, F. Francois, M. Channegowda, A. Hammad, S. Peng, R. Nejabati, D. Simeonidou, N. Yoshikane, T. Tsuritani, V. Lopez, and A. Autenrieth, "Experimental assessment of ABNO-based network orchestration of end-to-end multi-layer (OPS/OCS) provisioning across SDN/OpenFlow and GMPLS/PCE control domains," in *Optical Communication (ECOC), 2014 European Conference on*, Sept 2014, pp. 1–3.

[17] A. Aguado, V. Lopez, J. Marhuenda, O. Gonzalez de Dios, and J. Fernandez-Palacios, "ABNO: a feasible SDN approach for multi-vendor IP and optical networks [invited]," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 7, no. 2, pp. A356–A362, February 2015.

[18] A. Mayoral, R. Vilalta, R. Munoz, R. Casellas, and R. Martinez, "Performance analysis of SDN orchestration in the cloud computing platform and transport network of the ADRENALINE testbed," in *Transparent Optical Networks (ICTON), 2015 17th International Conference on*, July 2015, pp. 1–4.

[19] L. Gifre, N. Navarro, A. Asensio, M. Ruiz, and L. Velasco, "iONE: An environment for experimentally assessing in-operation network planning algorithms," in *Transparent Optical Networks (ICTON), 2015 17th International Conference on*, July 2015, pp. 1–4.

[20] J.-L. Auge, "Can we use flexible transponders to reduce margins?" in *Proc. of OFC/NFOEC, 2013*, March 2013.

[21] A. Mitra, A. Lord, S. Kar, and P. Wright, "Effect of link margins and frequency granularity on the performance and modulation format sweet spot of multiple flexgrid optical networks," in *Optical Fiber Communications Conference and Exhibition (OFC), 2014*, 2014.

[22] "D2.1 - ORCHESTRA dynamic optical network, reference scenarios and use cases," in *www.orchestraproject.eu*, Downloads/Public Documents.

[23] "Draft revised G.694.1 version 1.3," Unpublished ITU-T Study Group 15, Question 6.

[24] F. Paolucci, F. Cugini, A. Giorgetti, N. Sambo, and P. Castoldi, "A survey on the path computation element (PCE) architecture," *Communications Surveys Tutorials, IEEE*, vol. 15, no. 4, pp. 1819–1841, Fourth 2013.

[25] K. Christodoulopoulos, P. Kokkinos, A. Di Giglio, A. Pagano, N. Argyris, C. Spatharakis, S. Dris, H. Avramopoulos, J. Antona, C. Delezoide, P. Jenneve, J. Pesic, Y. Pointurier, N. Sambo, F. Cugini, P. Castoldi, G. Bernini, G. Carrozzo, and E. Varvarigos, "ORCHESTRA - optical performance monitoring enabling flexible networking," in *Transparent Optical Networks (ICTON), 2015 17th International Conference on*, July 2015, pp. 1–4.

[26] F. Cugini, F. Fresi, F. Paolucci, G. Meloni, N. Sambo, A. Giorgetti, T. Foggi, L. Poti, and P. Castoldi, "Active stateful PCE with hitless LDPC code adaptation," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 7, no. 2, pp. A268–A276, February 2015.

[27] "Strongest deliverable," in *http://www.ict-strongest.eu/upload-files/deliverables-2/d2-1-15/download*.

[28] N. Sambo, M. Secondini, F. Cugini, G. Bottari, P. Iovanna, F. Cavaliere, and P. Castoldi, "Modeling and distributed provisioning in 10-40-100-Gb/s multirate wavelength switched optical networks," *Lightwave Technology, Journal of*, vol. 29, no. 9, pp. 1248–1257, May 2011.