

Monitoring and Data Analytics for Optical Networking: Benefits, Architectures, and Use Cases

L. Velasco, A. Chiadò Piat, O. Gonzalez, A. Lord, A. Napoli, P. Layec, D. Rafique, A. D'Errico, D. King, M. Ruiz, F. Cugini, and R. Casellas

ABSTRACT

Operators' network management continuously measures network health by collecting data from the deployed network devices; data is used mainly for performance reporting and diagnosing network problems after failures, as well as by human capacity planners to predict future traffic growth. Typically, these network management tools are generally reactive and require significant human effort and skills to operate effectively. As optical networks evolve to fulfil highly flexible connectivity and dynamicity requirements, and supporting ultra-low latency services, they must also provide reliable connectivity and increased network resource efficiency. Therefore, reactive human-based network measurement and management will be a limiting factor in the size and scale of these new networks. Future optical networks must support fully automated management, providing dynamic resource re-optimization to rapidly adapt network resources based on predicted conditions and events; identify service degradation conditions that will eventually impact connectivity and highlight critical devices and links for further inspection; and augment rapid protection schemes if a failure is predicted or detected, and facilitate resource optimization after restoration events. Applying automation techniques to network management requires both the collection of data from a variety of sources at various time frequencies, but it must also support the capability to extract knowledge and derive insight for performance monitoring, troubleshooting, and maintain network service continuity. Innovative analytics algorithms must be developed to derive meaningful input to the entities that orchestrate and control network resources; these control elements must also be capable of proactively programming the underlying optical infrastructure. In this article, we review the emerging requirements for optical network management automation, the capabilities of current optical systems, and the development and standardization status of data models and protocols to facilitate automated network monitoring. Finally, we propose an architecture to provide Monitoring and Data Analytics (MDA) capabilities, we present illustrative control loops for advanced network monitoring use cases, and the findings that validate the usefulness of MDA to provide automated optical network management.

INTRODUCTION AND MOTIVATION

After years of research and development, the Elastic Optical Networking (EON) technology is currently being deployed in optical transport networks. This technology enables among others: the capacity and/or increase the reach and reliability of optical connections (hereafter, lightpaths), and a finer and dynamic spectrum allocation. The first is enabled by the joint usage of coherent detection, advanced Digital Signal Processing (DSP) techniques, novel modulation formats and soft-decision Forward Error Correction (FEC) codes to recover Bit Error Rate (BER) within the optical transponders (TP). The second is possible thanks to programmable wavelength selective switches (WSS) and reconfigurable optical add-drop multiplexers (ROADM).

From a network control perspective, an enormous amount of research and standardization effort has been carried out in recent years to implement the Software-Defined Networking (SDN) concept in optical networks [1]. SDN relies on the separation of the data plane and control plane, and leverages programmability and the usage of open interfaces. However, little to no attention has been paid to the operational loop (including monitoring, intelligence and management functionalities), relegating some of them to the network management systems (NMS), and with limited practical operation capabilities.

Although EON and SDN technologies can fulfil current capacity and dynamicity requirements, transport networks are expected to support the deployment of upcoming 5G mobile infrastructures in the near future; 5G mobile will extend far beyond previous generations and require an enhanced quality of experience for the final users with new services and improved network performance. To meet the goals of 5G, network infrastructures should provide increased levels of flexibility and automation, together with higher priority given to network optimization, security, energy consumption, and cost efficiency. In fact, disaggregation at the optical layer has been conceived to enrich the offer of available solutions and to enable the deployment of optical nodes that better fit optical network operators' needs.

As future network complexity increases, the main challenge for operators will be to promptly respond to variable network conditions while

Use case	Description	Expected benefits	Modeling and parameters involved
Network planning and provisioning with reduced margins.	Application of <i>just enough</i> margin in the network design and in lightpath provisioning.	CAPEX saving opportunity by avoiding or postponing unnecessary investments at a given time.	Attenuation, dispersion and other fiber parameters, the noise figure of amplifiers, WSS passband, the sensitivity of TPs, etc. Those parameters can be used together with an analytical model to estimate the QoT of lightpaths accurately. ML-based methods to predict the probability that the QoT of a candidate lightpath will not exceed a defined threshold.
Dynamic network adaptation.	Leveraging on configurable TPs the allocation of <i>just enough</i> data rate for any connection at any time to cope with traffic dynamics at a minutes or hours scale.	Better exploitation of network resources and potential savings by reducing the typical overprovisioning of static allocation.	Use of models to evaluate the expected QoT of a lightpath at any new TP configuration. Use of models for traffic analysis to evaluate traffic trends and periodicity.
Lightpath degradation and failure localization.	QoT reduces over time due to network and device degradation (e.g., fiber cuts and repairs), ageing, or load increasing.	Degradation anticipation allows appropriate tuning of systems' parameters before alarm triggering. Localizing the element responsible for a failure facilitates network maintenance by planning a human intervention.	Predictive analysis based on QoT evolution. Localization based on the per-system analysis. Algorithms that find the potential cause of the failure.

TABLE 1. Target use cases.

ensuring full availability and optimization of network resources. Nonetheless, current optical networks are incorporating a complex ecosystem of devices and sensors, which will produce a large amount of data that can be exploited to optimize a network in real-time. To cope with complex and time-variable 5G service scenarios, Machine Learning (ML)-based algorithms [2] are being proposed to facilitate network operation and predictive maintenance. ML algorithms, fed with real measurements, are able to accurately estimate the Quality of Transmission (QoT) of new lightpaths, to anticipate capacity exhaustion and degradations, or to predict and localize failures, among others (see, e.g., [3–8]).

Based on the above facts, network operators are looking with high interest to the opportunities that Monitoring and Data Analytics (MDA) can offer to their optical transport networks, as it emerges from applied research and standardization bodies. In fact, such solutions can be made available only after monitoring and telemetry protocols, together with data models, are standardized. There are multiple ongoing standardization efforts within several technology areas, where most of the proposals are based on three main principles: data modeling language that provides structured data models for technology and function specific data points; management protocol for encoding and carrying the data model information; and the operational process governing how the protocol interface is used and connections are managed. In addition to research activity working on network telemetry (see, e.g., [9]), practical industrial projects exist, including: OpenConfig (see openconfig.net) and the OpenROADM (see openroadm.org) efforts.

In this article, we review the operators' vision, as well as the capabilities of current optical systems and present three wide-scope use cases that require MDA-based solutions and whose application will bring clear benefits: i) network planning and provisioning with reduced margins, ii) dynamic network adaptation, and iii) lightpath degradation detection and failure localization. Next, the state-of-the-art of data models and monitoring and

telemetry protocols is reviewed as well. With this in mind, several MDA architectures are proposed, and the pros and cons of each of them are highlighted. Finally, illustrative control loops for the considered use cases bring a clear and complete vision of the validity and feasibility of MDA in the context of optical transport networks.

OPERATORS' VISION IN NEAR-TERM AND DATA AVAILABILITY THE NETWORK OPERATORS' VISION

The vision of network operators, regarding the deployment of MDA in their optical networks, mostly concentrates on three wide-scope use cases, as summarized in Table 1.

The *first use case* focuses on minimizing the system operation margins, for example, linear optical signal-to-noise-ratio (OSNR), that are widely used in optical systems to ensure worst-case end-of-life QoT of the lightpaths. Before entering operation, all available combinations of modulation formats, fiber types, FEC codes, and so on, are considered, and exhaustive numerical simulations and lab experiments are conducted to extract engineering rules to be used. This time-consuming analysis can be simplified by utilizing approximate analytical tools such as the Gaussian Noise (GN) model [10]. Both approaches lead to the estimation of QoT for the existing lightpaths. Nonetheless, these solutions are static by nature and based on conservative design principles, which lead to resource underutilization.

To reduce margins, analytical methods or ML-based algorithms can utilize the knowledge of the current network status, that is, the configuration of optical devices (e.g., TPs, WSSs, ROADMs, and so on) and the characteristics of the optical fibers to estimate the QoT of new lightpaths to be established [3]. During operation, the SDN controller is in charge not only of the provisioning process, but also of adapting the network to traffic changes (it is quite common that packet traffic varies from day to night not only in intensity but also in directionality due to, for example,

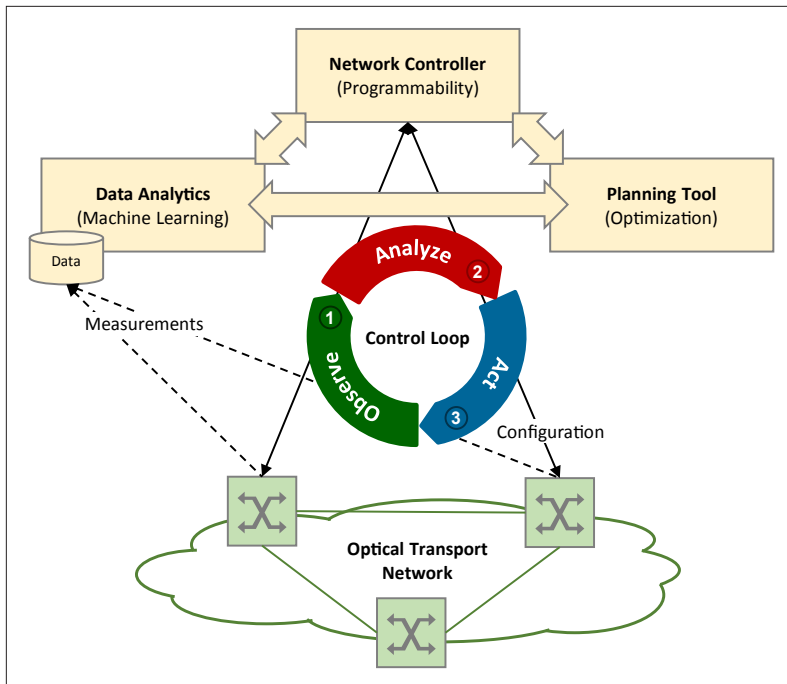


FIGURE 1. Monitoring and Data Analytics enable observe-analyze-act control loop implementation.

data-center activity) aiming at minimizing overprovisioning. In the *second use case*, the role of MDA is to derive models to accurately predict the traffic volume for the short term, in detecting whether the capacity of the lightpaths will be soon exhausted, and so on [4]. With such knowledge, the SDN controller can re-configure the network leveraging the configurability of TPs, that is, adapting the rate and spectrum of already established lightpaths and creating new lightpaths in real time with significant CAPEX and OPEX savings.

The *last use case* concentrates on degradations and failures. All components deployed into an optical network suffer ageing over their lifetime, for example, the amplifiers might decrease their gain, the filters might insert additional losses, the fiber might present several splices, and so on. This leads to a slow, but continuous, decrease of the lightpaths' QoT. Early detection of lightpaths' degradation would allow tuning parameters within the TPs, for example, by increasing the FEC overhead or by switching to a more robust modulation format [5]. When the severity of the degradation increases, localizing its root cause is of paramount importance for maintenance purposes [6, 7]. It is also possible to predict failures and proactively re-route the traffic [8], which allows a high resiliency of the optical network at the just-enough cost. To this end, dedicated optical protection is replaced with just-in-time optical restoration.

Four aspects are particularly important and must be implemented to support the three use cases described above:

- Which data may be obtained, derived or provided by the network devices and collected by the operators.
- Which are the key parameters to be estimated and the accuracy required.
- Identification of technologies that can be used to elaborate the information.

- Definition of the main limitations in terms of data availability, veracity, and frequency that exist and what is needed to overcome them.

DATA AVAILABILITY

Considering the use cases defined in Table 1, optical devices need to be capable of performing measurements on selected points of the networks, named observation points (OP). For example, measurements could be obtained from DSP units within the TPs, as well as from specific monitoring devices installed within the network. Specifically, DSPs can provide measurements or estimations of power levels, fiber channel characteristics (e.g., accumulated dispersion, fiber nonlinear coefficient, polarization mode dispersion) and QoT-related parameters (e.g., linear OSNR and BER). Furthermore, monitoring devices, like cost-effective optical spectrum analyzers (OSA) and optical time-domain reflectometers placed at predefined locations of the network, can provide specific measurements of optical signals and fiber segments.

Among all available and derived data, the most relevant is the OSNR measured at the receiver, which is used to define the system margin of every lightpath. While the estimation in the linear regime is straightforward, the GN model can provide a worst-case accuracy as low as ± 0.75 dB at the optimal power level or in the nonlinear regime. An accurate enough value of the system OSNR would enable strategies that can lead to optimal usage of the optical spectrum. ML-based algorithms can also contribute to estimating this and other parameters, like laser characteristics or amplifier noise figures.

It is clear that operational data (i.e., the network topology, the route of lightpaths, and so on) is of paramount importance to realize all the above use cases, as they allow to correlate measurements and events; such operational data can be collected from the SDN controller. In addition, lightpath provisioning activity can also be collected from the SDN controller and used for traffic modeling. Other parameters can be available as well, like traffic forecasts that can be used to further optimize network operations or to predict failures.

Finally, by deploying low-cost monitoring devices, environmental parameters could also be exploited and eventually correlated for optimal network operation.

CONSIDERATIONS ABOUT MDA-BASED SYSTEMS

Besides data availability, it is also important to consider their accuracy to define the sufficient amount of data to be collected and stored, as the accuracy depends on the amount of data that is considered in the MDA system for decision making. For example, if a system would operate in a pure linear regime, the pre-FEC BER could be enough to estimate the actual OSNR and then the relative system margin. However, real networks do not always operate in a full linear regime, and therefore the pre-FEC BER may result in being unsuitable to always provide an accurate prediction of the instantaneous OSNR margin. Consequently, enough data need to be stored to achieve a pre-defined accuracy, especially under low or zero-margin network operation.

Protocol	Description	Data types	Scalability and performance
IPFIX	<ul style="list-style-type: none"> • IPFIX was developed in IETF for typical IP networks applications. • It works in push mode and supports a many-to-many relationship between OPs and MDA collectors. 	<ul style="list-style-type: none"> • The structure of the IPFIX protocol messages is based on templates that enable the export of any type of data. 	<ul style="list-style-type: none"> • Scalability is considered in the design of the protocol, with a requirement of hundreds of different exporting processes to be supported.
gRPC	<ul style="list-style-type: none"> • gRPC uses by default Protocol Buffers, a mechanism for serializing structured data. • It supports data streaming based on a request/response model. 	<ul style="list-style-type: none"> • Specific data structures can be defined; a compiler can generate source code in various programming languages, representing the data, and methods to serialize them. 	<ul style="list-style-type: none"> • gRPC is carried over HTTP/2 and leverages effective binary framing and header compression that improve data transfer efficiency.
Thrift	<ul style="list-style-type: none"> • Apache Thrift is an open source software library and set of code-generation tools. • Thrift is stream-oriented by design. 	<ul style="list-style-type: none"> • It allows the definition of datatypes and generates all the necessary code in different programming languages. 	<ul style="list-style-type: none"> • Data transfer efficiency is comparable to that of gRPC.

TABLE 2. Monitoring and telemetry protocols.

Another key factor is the *update frequency*; an instantaneous collection of monitoring data could produce negative effects, so it is important to determine the right frequency for data collection. For instance, once a lighthouse is established, and until there are no substantial changes in the network, there is no need to update the fiber channel values. Contrarily, parameters such as amplifier power levels require a higher update frequency, although old values could be discarded if the individual amplifier works properly. Overall, all these data will be ultimately used by the MDA system, which might also result in saturation or in drawing sub-optimal decisions in the case of overwhelming or contradictory data.

Different strategies can be envisioned to solve this issue:

- Using thresholds, which are simple but inaccurate.
- Experience and physical knowledge, which could lead to evaluation errors in the case of not predicted scenarios.
- By designing an intelligent MDA system that can decide based on physical conditions which data should be analyzed and consider possible dependencies.

Finally, it is worth pointing out that the main challenge (and limitations) occur in multi-vendor scenarios. In this context, a proactive MDA system could anticipate issues before they happen and issue the proper recommendations provided that the MDA system is aware of the configuration of all involved nodes at any time.

In conclusion, the opportunities that MDA opens go far beyond a monitoring data collector and storage platform. The analysis of the collected data can discover knowledge and use it to proactively self-configure and self-tune the network in a cost-effective (near) real-time manner by adapting resources to future conditions. Therefore, thanks to the application of data analytics to monitored data, observe-analyze-act control loops can be enabled, where outcomes of such analysis can be used for event notifications together with recommended actions to the SDN controller (Fig. 1). Last but not least, useful models can be estimated from monitoring data to feed planning tools in order to compute optimal solutions for the expected future conditions.

YANG DATA MODELS AND PROTOCOLS

YANG is a data modeling language standardized by the Internet Engineering Task Force (IETF) and designed to operate with the NETCONF protocol for network configuration and management (see IETF RFC 6020). YANG enables:

- Human readability and simplified troubleshooting operations compared to protocols relying on bit encoding.
- Hierarchical structures of data models.
- Extensibility and modularity through augmentation mechanisms and sub-modules.

A YANG data model is represented by a tree structure where nodes are defined by names, data types, data values, or a set of child nodes and lists. In recent years, in the context of optical networks, YANG/NETCONF has emerged as a candidate solution to provide automated control of network elements having common and vendor-neutral standardized models [11]. Several standardization bodies, like the IETF, and working groups, for example, the aforementioned OpenConfig and OpenROADM, have released vendor-neutral YANG models for devices such as X-ponders, optical amplifiers, and ROADMs. However, the related YANG models are significantly different, with relevant incompatibility issues. Although efforts are ongoing to converge toward commonly adopted models, multiple versions of drivers, software implementations and SDN controllers and monitoring customizations are expected in the near future, potentially delaying the adoption in heterogeneous and multi-vendor networking scenarios.

For monitoring purposes, YANG relies on state (read-only) types providing the actual values of the considered system parameters. The SDN controller is able to retrieve YANG-defined parameter values by exploiting NETCONF messages either periodically (e.g., every 15 minutes) or asynchronously (e.g., in case of events) through notification messages. However, NETCONF messages are not particularly efficient for monitoring (especially when the data collection period is short, for example, one minute) let alone for telemetry (e.g., when a continuous stream of data is provided). Thus, other protocols have been proposed for monitoring and telemetry purposes; the most relevant are IP Flow Information eXport (IPFIX) (see IETF RFC 3917); gRPC (see gpc.io), and Apache

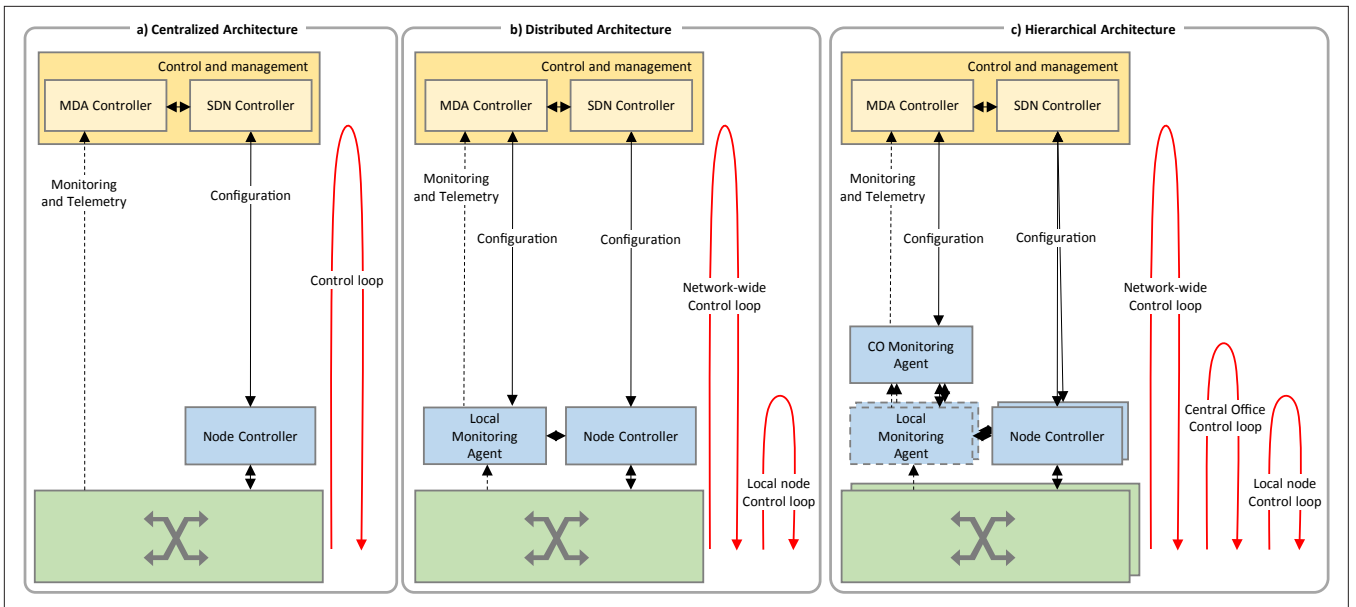


FIGURE 2. Overview of the MDA architectures: (a) centralized, (b) distributed, and (c) hierarchical.

Thrift (see thrift.apache.org). For a brief description refer to Table 2.

With these protocols available, the selection of the collection period is no longer limited to 15 minutes, and it can be reduced to, for example, one second [9]. Note that the shorter the collection period, the shorter the event that can be detected, as well as the shorter the time to detect degradations. However, reducing the measurement period increases the amount of data that has to be collected, stored, and analyzed. Then, an approach to reduce the amount of data is to rely on monitoring using a collection period of minutes and activate telemetry on demand to get insight, by analyzing a continuous stream of measurements, when and where needed.

MDA ARCHITECTURES

In this section, we present and analyze several architectural approaches to bring real MDA capabilities to the network (Fig. 2). Specifically, three architectures are considered depending on where data analytic capabilities are enabled, namely centralized, distributed, and hierarchical.

The *centralized architecture* (Fig. 2a) consists of detaching the monitoring repository and the data analytics system, if any, from the NMS to create a separate specific centralized MDA controller that can interface the SDN controller and other systems within the control plane (see, e.g., Ciena Blue Planet). To keep the MDA architecture simple, let us consider that its only mission is to expose an interface to collect monitoring and telemetry data from the network devices. Measurements are stored in a (big data) repository, and data analytics algorithms can be devised to discover knowledge to be used to predict and/or to detect anomalies and degradations before they negatively impact the network performance. Such predicted events can be notified to the SDN controller together and include a recommended action to guide the SDN controller; the recommended action is a suggestion that the SDN controller can follow or ignore and apply its own policies. As an example, in some cases

BER degradation can be predicted ahead of time in a lightpath before any threshold is exceeded by analyzing the BER evolution as measured at the receiver. This is notified to the SDN controller together with a *recommended action* after analyzing several alternatives, including changing the modulation format (also via probabilistic shaping), re-route of the lightpath (e.g., to avoid some links); or also increasing, if possible, the amount of overhead used by the FEC. The notification to the SDN controller might trigger a re-configuration, hence closing the loop and adapting the network to the new conditions.

The centralized MDA architecture presents some limitations; for instance, the time to detect an anomaly or degradation is related to the update frequency. Therefore, to reduce the detection times, the amount of data to be conveyed to the MDA controller needs to be increased accordingly. Another issue is related to the control of monitoring; specifically, to activate telemetry on-demand once an event has been detected.

To overcome these problems, the *distributed architecture* (Fig. 2b) includes *MDA agents* in charge of collecting measurements from a single node, while keeping the MDA controller centralized [12, 13]. The MDA agent exposes two unified interfaces toward the MDA controller for collecting data and monitoring the configuration. In addition, specific interfaces for data collection and monitoring control allow the MDA agent to connect with the network device. The data analytics capabilities deployed close to the network nodes enable local control loop implementation; measurements can be analyzed locally, and configuration can be tuned and adapted to changing conditions. However, the co-existence of two controllers, the SDN and the MDA, in charge of configuring network devices, might create conflicts, so it would be desirable to clearly separate responsibilities among them.

The distributed architecture includes a dedicated MDA agent for every node in the network, which might present some limitations when disaggregated optical network nodes (e.g., TPs and

Architecture	Features	Strengths	Weaknesses
Centralized	<ul style="list-style-type: none"> Includes a centralized MDA system with a data repository for monitoring/telemetry data where data analytics can be applied. Monitoring and telemetry activation and deactivation is managed by an external system, e.g., the NMS. 	<ul style="list-style-type: none"> Data analytics results can be used for network self-adaptation to changing conditions. Interfaces with the SDN controller and NFV orchestrator can be easily standardized. 	<ul style="list-style-type: none"> Different monitoring/telemetry protocols need to be available at the MDA controller. The amount of data to be collated from the nodes increases exponentially to maintain low reaction times against anomalies or degradations. Configuration tuning is not supported. Network slicing is difficult to be supported.
Distributed	<ul style="list-style-type: none"> Allows data analytics to be applied within the MDA agents, close to the network nodes. Control loops can be implemented locally at the node level. Monitoring and telemetry activation/deactivation is managed by the MDA controller. 	<ul style="list-style-type: none"> Supports configuration tuning [14]. It reduces data to be conveyed to the MDA controller since pattern recognition can be done in the MDA agents. MDA agents expose one single monitoring and telemetry interface to the MDA controller. Supports network slicing [12]. 	<ul style="list-style-type: none"> A configuration interface needs to be defined between the MDA controller and the agents. More complex MDA controller as more features are added, like monitoring and telemetry control, and configuration tuning. CO control loops are not supported
Hierarchical	<ul style="list-style-type: none"> It includes a per-CO MDA agent that connects to all the nodes in the CO. 	<ul style="list-style-type: none"> Control loops can be implemented locally at the node, as well as at the CO level involving more than one node. Appropriate for node disaggregation scenarios, where monitoring devices can be installed in one node, but configuration tuning needs to be done in a different node [14]. It reduces the total number of agents and the number of interfaces toward the MDA controller. 	<ul style="list-style-type: none"> Requires more complex MDA agents to consider complex relations among nodes.

TABLE 3. Strengths and weaknesses of several monitoring and data analytics architectures.

ROADMs) and monitoring devices are deployed within the same central office (CO) [13]. For this reason, the *hierarchical architecture* (Fig. 2c) includes a per-CO MDA agent that collects measurements from every network device in the CO and exposes a single set of interfaces toward the MDA controller. In this case, measurements from one device can be analyzed in the CO MDA agent and configuration can be tuned to another device within the same CO, thus minimizing the intervention of the MDA controller. The per-CO MDA agent could (or not) replace every node MDA agent, thus reducing the systems count.

The strengths and weaknesses of the analyzed MDA architectures are summarized in Table 3, where the features of each architecture include those of the previous.

ILLUSTRATIVE CONTROL LOOP IMPLEMENTATION

This section illustrates how the use cases introduced earlier can be implemented. To this end, let us assume a disaggregated scenario, where COs are equipped with TP nodes and ROADMs and the hierarchical MDA architecture is selected. Apart from the MDA, the control plane includes an SDN controller in charge of configuring the optical network, a planning tool running optimization algorithms for provisioning and in-operation network planning purposes [15], and an NMS for human operators to manage the network.

Additionally, it is worth highlighting that if external systems, such as planning tools, may require access to data stored in the MDA controller upon request, it is necessary to define additional interfaces. The data that is then available, as part of the MDA, is not simply the raw measurements being streamed from the network devices, but also estimated data and derived knowledge generated by ML algorithms.

LIGHTPATH PROVISIONING WITH A REDUCED MARGIN

In this first use case, we focus on the provisioning of lightpaths minimizing the system margin, which can be derived from the OSNR and/or the TP's pre-FEC BER threshold according to the transmission scenario. OSNR estimation at optimal launch power and in a nonlinear regime requires data from monitoring the optical channel and the devices' configuration, which we assume are already available in the MDA controller (labeled 0 in Fig. 3a). Besides this information, parameters related to network infrastructures (such as fiber types and link lengths) are also required. These might be collected, for example, from the SDN controller, the NMS and inventory systems.

When a lightpath setup request arrives at the SDN controller (1), the latter relies on the planning tool for the computation of the route, spectrum allocation, modulation format and other parameters that contribute to minimize the system margin while guaranteeing its QoT (2). In order to compute an optimal solution and meet the objective function criterion, the planning tool needs to access data from the MDA controller (3); once a solution has been found, it is sent back to the SDN controller (4). Here, (at least) three possibilities might exist:

- The lightpath can be established, and an optimal configuration has been found.
- The lightpath can be established provided that the configuration of other lightpaths is first changed.
- No solution has been found.

In the second case, the planning tool returns the optimal configuration found for the requested lightpath, together with a (*reactive*) *recommended action* for the SDN controller to modify the

configuration of a subset of already established lightpaths. In this case, the SDN controller might request the human operator to confirm the re-configuration through the NMS (5). Finally, in the case that the lightpath can be established with or without network re-configuration, the SDN

controller configures the network devices accordingly.

Further optimization can then be achieved by observing the QoT of each established lightpath, aiming at identifying possible transmission adaptations (e.g., FEC, modulation format) leading to margin reductions closer to the predefined target values.

DYNAMIC NETWORK ADAPTATION

In the previous control loop, the planning tool issued a recommended action for re-configuration because of a previous request from the SDN controller (we named them reactive). In this and the next use cases, the MDA controller will issue recommendations to the SDN controller as a result of observing what is happening in the network and aiming at anticipating the most relevant events. In this context, we refer them to as proactive recommendations.

As for the case before, we assume that data are already available within the MDA controller (labeled 0 in Fig. 3b). ML algorithms running in the MDA controller can use the measured packet traffic volume to determine a traffic model for the traffic between every origin and destination CO. Such traffic models can be used to compare the expected traffic against the provisioned capacity and therefore, when the measured or the expected traffic for the near future is close to the allocated capacity, the MDA controller issues a notification to the SDN controller including a recommended action to reconfigure the allocated capacity (1).

In these dynamic cases, the SDN controller might inform the human operator through the NMS (2), and then request the planning tool to compute the optimal capacity configuration for the detected event (3). For such computation, the planning tool needs data from the MDA controller, for example, the expected traffic matrix, for example, for the next hours (4) [4]. Such a traffic matrix can be computed assuming the maximum or the 95th percentile traffic volume expected for every origin-destination pair. With such a traffic matrix, an optimization algorithm running in the planning tool can compute the optical capacity allocation and respond to the SDN controller (5). Finally, the SDN controller implements the re-configuration in the network (6).

LIGHTPATH DEGRADATION DETECTION AND MODULATION FORMAT ADAPTATION

For this use case, let us consider the lightpath is established and being monitored, where BER measurements are collected by the MDA agents connected to the end TPs (labeled 0 in Fig. 3c). A data analytics algorithm running within the MDA agents can be in charge of detecting BER trends to anticipate QoT degradation [5]. In the case of QoT degradation detection, a decision can be locally made without the intervention of the MDA controller. For instance, modern TPs are capable of identifying the modulation format of the received signal by means of DSP. Therefore, a change in the modulation format employed for a lightpath can be initiated in one of the transmitters and the end TPs will automatically realize such change and carry out the same in the opposite direction (such a local control

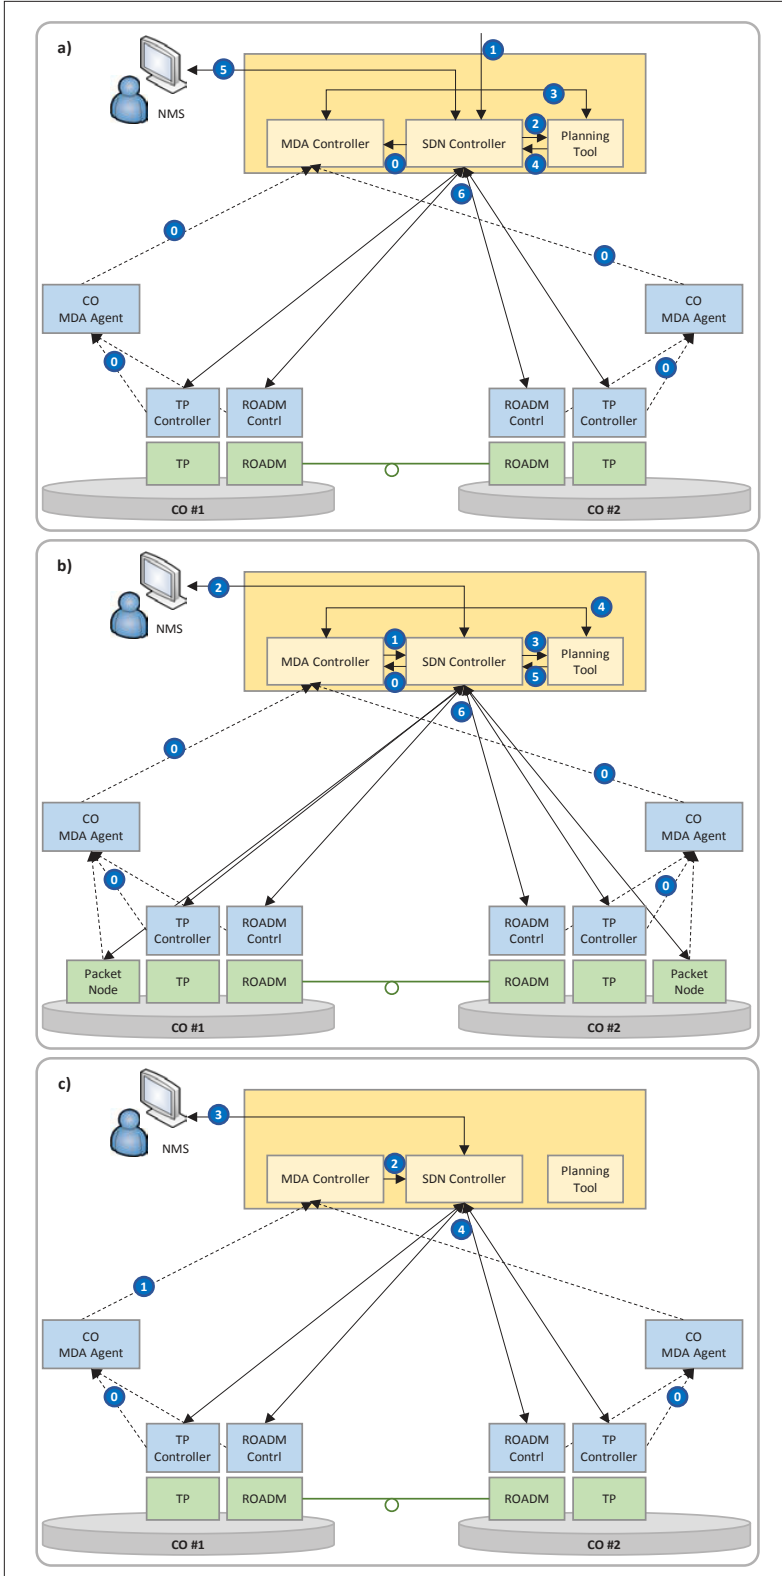


FIGURE 3. Control loops implementation: (a) lightpath provisioning with reduced margin; (b) dynamic network adaptation; and (c) lightpath degradation detection and modulation format adaptation.

loop is not shown in Fig. 3c). However, in disaggregated multi-vendor scenarios, both ends could need to be simultaneously re-configured. To that end, the MDA agent sends a notification to the MDA controller (1) that evaluates the capabilities of both TPs and evaluates the possibilities. The degradation detection together with a recommendation (e.g., change the modulation format to a more robust one) is sent to the SDN controller (3) that implements it in the devices, maybe after checking it with the operator in the NMS (3-4).

SUMMARY

We have provided the network operator vision for automating management of an advanced optical network infrastructure, key requirements and current enabling optical technologies. The role of MDA in optical networking has been studied through three wide-scope use cases covering the main network operations: network planning and provisioning, dynamic network adaptation, and degradation detection and failure localization, where clear benefits have been shown. Interestingly, current networking devices are already capable of performing measurements that support those use cases. Additional data can be collected by installing specific monitoring devices at predefined locations.

A review of the currently ongoing standardization activities revealed that different initiatives are working toward modeling optical components and adopting different solutions. In addition, several protocols can be used for monitoring and telemetry purposes. From the control plane perspective, it is not clear the support of SDN controllers to the MDA functions more than just collecting monitoring data. In view of that, a specific MDA system has been proposed and three different architectures, from centralized to distributed, were analyzed, where an MDA controller is defined in the control plane working together with the SDN controller. Finally, illustrative control loops supporting examples of the selected use cases have been shown.

As a final remark, although the technologies supporting MDA in optical networks are ready, there is still a significant amount of discussion required within the relevant standardization forums and industrial OpenSource projects, to leverage this work fully.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Community through the METRO-HAUL project (G.A. n 761727), from the TWINS project (TEC2017-90097 AEI/FEDER, UE), from the Catalan Institution for Research and Advanced Studies (ICREA), and from the CELTIC project SENDATE-FICUS funded by the German BMBF, contract no. 16KIS0487K.

REFERENCES

- [1] R. Casellas *et al.*, "Control, Management, and Orchestration of Optical Networks: Evolution, Trends, and Challenges," *IEEE/OSA J. Lightwave Technology*, vol. 36, 2018, pp. 1390–1402.
- [2] D. Rafique and L. Velasco, "Machine Learning for Optical Network Automation: Overview, Architecture and Applications [Invited Tutorial]," *IEEE/OSA J. Optical Communications and Networking*, vol. 10, 2018, pp. D126–43.

Although the technologies supporting MDA in optical networks are ready, there is still a significant amount of discussion required within the relevant standardization forums and industrial OpenSource projects, to leverage this work fully.

- [3] L. Barletta *et al.*, "QoT Estimation for Unestablished Lightpaths Using Machine Learning," *Proc. OSA Optical Fiber Commun. Conf.*, 2017.
- [4] F. Morales *et al.*, "Virtual Network Topology Adaptability Based on Data Analytics for Traffic Prediction," *IEEE/OSA J. Optical Commun. and Networking*, vol. 9, 2017, pp. A35–45.
- [5] A. P. Vela *et al.*, "BER Degradation Detection and Failure Identification in Elastic Optical Networks," *IEEE/OSA J. Lightwave Technology*, vol. 35, 2017, pp. 4595–4604.
- [6] A. P. Vela *et al.*, "Soft Failure Localization During Commissioning Testing and Lightpath Operation," *IEEE/OSA J. Optical Commun. and Networking*, vol. 10, 2018, pp. A27–36.
- [7] D. Rafique *et al.*, "Cognitive Assurance Architecture for Optical Network Fault Management," *IEEE/OSA J. Lightwave Technology*, vol. 36, 2018, pp. 1443–50.
- [8] F. Boitier *et al.*, "Proactive Fiber Damage Detection in Real-Time Coherent Receiver," *Proc. European Conf. Optical Commun.*, 2017.
- [9] F. Paolucci *et al.*, "Network Telemetry Streaming Services in SDN-based Disaggregated Optical Networks," *IEEE/OSA J. Lightwave Technology*, vol. 36, 2018, pp. 3142–49.
- [10] P. Poggiolini, "The GN Model of Non-Linear Propagation in Uncompensated Coherent Optical Systems," *IEEE/OSA J. Lightwave Technology*, vol. 30, 2012, pp. 3857–79.
- [11] M. Dallaglio *et al.*, "YANG Models for Vendor-Neutral Optical Networks, Reconfigurable through State Machine," *IEEE Commun. Mag.*, vol. 55, 2017, pp. 170–78.
- [12] L. Velasco *et al.*, "An Architecture to Support Autonomic Slice Networking [Invited]," *IEEE/OSA J. Lightwave Technology*, vol. 36, 2018, pp. 135–41.
- [13] L. Velasco *et al.*, "Building Autonomic Optical White-box-based Networks," *IEEE/OSA J. Lightwave Technology*, vol. 36, 2018, pp. 3097–3104.
- [14] L. Gifre *et al.*, "Autonomic Disaggregated Multilayer Networking," *IEEE/OSA J. Optical Communications and Networking*, vol. 10, 2018, pp. 482–92.
- [15] L. Velasco and M. Ruiz, *Provisioning, Recovery and In-Operation Planning in Elastic Optical Networks*, Wiley, 2017.

BIOGRAPHIES

L. VELASCO (lvelasco@ac.upc.edu) received the M.Sc. degree in physics from the Universidad Complutense de Madrid (UCM) in 1993 and the Ph.D. degree from the Universitat Politècnica de Catalunya (UPC) in 2009. In 2004 he joined UPC, where currently he is a full professor with the Department of Computer Architecture (DAC). His interests include monitoring and data analytics aspects for the service and network layers.

A. CHIADÒ PIAT after graduation in electronic engineering at Politecnico di Torino in 1991, joined Telecom Italia in 1992 where she is currently working in the IP & transport innovation unit. She was involved in the study of transport network technologies, formerly SDH, and then packet transport, with a particular interest in frequency and phase synchronization delivery. More recently she has been involved in the study of optical networks, in particular elastic optical networks.

O. GONZÁLEZ received his M.Sc. in telecommunications engineering, as well as his Ph.D. with honors from the University of Valladolid. He has 16 years of experience in Telefonica I+D, where he has been involved in numerous R&D European projects (recently STRONGEST, ONE, IDEALIST), where he has acted as WP leader.

A. LORD joined BT in 1985 after a B.A. in physics from Oxford University. He has worked on a wide range of optical network systems and technologies, including long haul subsea and terrestrial DWDM networks. He currently heads BT's optical core and access research including SDN and quantum communications.

A. NAPOLI received the master (2002) and Ph.D. (2006) degrees in electronics engineering from Politecnico di Torino, Italy. In 2006, he joined Siemens COM (now Infinera) in Munich, where he worked on EDFA transient suppression and DSP for long-haul optical networks and optical transmission up to 400 Gb/s. His research interests include DSP for optical receivers, modulation techniques for long-haul systems, and system design.

P. LAYEC is a Bell Labs researcher in France. She received her Ph.D. degree from the University of ParisSud XI, France, in 2009. In 2005, she worked at Orange on MIMO transceiver design and multi-user information theory. In 2009, she joined Alcatel-Lucent (now Nokia), first working on wireless physical layer dimensioning (HSPA+/LTE) in the business division, then joining Nokia Bell Labs. Since 2016, she heads the Physics of Optical Networks research team. Her research interests include information theory, digital signal processing and elastic optical networks.

D. RAFIQUE is a senior innovation manager and principal engineer in the CTO Office at ADVA Optical Networking, where he is working on technology roadmap definitions for cognitive network design, control and management and several EU/national projects. He holds B.Sc., M.Sc., and Ph.D. degrees, all in electrical engineering.

A. D'ERRICO is a member of the Ericsson Research branch in Pisa. He received the Ph.D. degree with honors in telecommunication engineering from Scuola Superiore Sant'Anna in Pisa, Italy. He has 15 years of experience in photonics technologies. His main activities have been focused on advanced technological solutions for the optical transport networks. He is now working on photonic enabling technologies toward 5G Mobile Networks.

D. KING holds an MBA from Lancaster University where he is currently a Ph.D. student specializing in SDN and NFV. He has worked extensively in industry for established Internet infrastructure companies (Cisco, Bell Labs, Redback, Marconi) and for

optical start-up Movaz Networks. He was a co-founder of Aria Networks where he served as COO for four years. For the last seven years, he has been an independent consultant working with Old Dog Consulting.

M. RUIZ received the M.Sc. degree in statistics and operations research in 2009 and the Ph.D. degree in 2012, both from UPC. In 2012 he joined UPC, where currently he is an assistant professor in the DAC and a post-doc researcher in the CCABA. His research interests include optimization and data analytics for telecom cloud applications.

F. CUGINI is Head of Research at CNIT, Pisa, Italy. His main research interests include theoretical and experimental studies in the field of packet and optical communications. Specific research activities have been carried out in the context of SDN control plane, management and monitoring architecture for optical networks, segment routing, cloud networking, and service chaining.

R. CASELLAS graduated in telecommunications engineering in 1999 both from UPC (Barcelona) and ENST (Paris), where he later completed a Ph.D. in 2002, funded by France Telecom R&D. He has worked as an undergraduate researcher at France Telecom R&D and British Telecom Labs and as an associate professor at the ENST. He joined the CTTC in 2006, where he currently holds a senior researcher position. His research interests include traffic engineering, GMPLS/PCE architectures, Software Defined Networking (SDN) and Network Function Virtualization (NFV).