


Toward efficient, reliable, and autonomous optical networks: the ORCHESTRA solution [Invited]

K. CHRISTODOULOPOULOS,^{1,2,*} C. DELEZOIDE,³ N. SAMBO,⁴ A. KRETSIS,^{5,6} I. SARTZETAKIS,^{5,6}
A. SGAMBELLURI,⁴ N. ARGYRIS,¹ G. KANAKIS,¹ P. GIARDINA,⁷ G. BERNINI,⁷ D. ROCCATO,⁸
A. PERCELSI,⁸ R. MORRO,⁸ H. AVRAMOPOULOS,¹ P. CASTOLDI,⁴ P. LAYEC,³  AND S. BIGO³

¹NTUA, Athens, Greece

²Current address: Nokia Bell Labs, Stuttgart, Germany

³Nokia Bell Labs, Nokia Paris-Saclay, Nozay, France

⁴SSSA, Pisa, Italy

⁵CTI Patras, Greece

⁶Current address: NTUA, Athens, Greece

⁷Nextworks, Pisa, Italy

⁸TIM, Turin, Italy

*Corresponding author: konstantinos.1.christodoulopoulos@nokia-bell-labs.com

Received 19 April 2019; revised 24 June 2019; accepted 24 June 2019; published 15 August 2019 (Doc. ID 365671)

Optical networks have historically been designed to be operated statically. Connections are overprovisioned so that they remain uninterrupted over several (e.g., 10) years, using high physical-layer margins to cover the evolution of the physical conditions and modeling uncertainties. As a first step, we can increase the efficiency without sacrificing network reliability by removing uncertainties and reducing long-term margins, observing and adjusting them at intermediate periods. This requires certain automation steps in monitoring and data processing. Increasing the efficiency further, and thus further reducing the margins, comes at a trade-off in reliability, and should be done according to service classes and the level of network automation. The ORCHESTRA network makes use of coherent optical transponders as software-defined optical performance monitors (soft-OPMs) to improve the optical network observability. ORCHESTRA developed digital signal processing (DSP) OPM algorithms and a hierarchical monitoring plane to carry and process physical-layer monitoring data. ORCHESTRA uses data analytics methods to understand the physical-layer conditions and feed cross-layer optimization algorithms. ORCHESTRA closes the observe–decide–act control loop, automating the mechanisms required to trade efficiency for reliability. © 2019 Optical Society of America

<https://doi.org/10.1364/JOCN.11.000C10>

1. INTRODUCTION

To accommodate continuous traffic growth and dynamics, coherent transmission and elastic optical networks (EONs) are the main solutions currently being deployed [1]. New-generation commercial transponders use coherent receivers (Rxs) and digital signal processing (DSP) for modulation/demodulation and for mitigating several physical-layer impairments. Such transponders provide programmable multiformat transmission features, and combined with flex-grid reconfigurable optical add–drop multiplexers (ROADMs) are typically referred to as EONs. EONs provide higher granularity, flexibility, and efficiency, and enable dynamic network re-optimization.

However, optical networks have historically been designed to be operated statically following the “set and forget” approach. Optical connections are designed so that once

established, they remain uninterrupted over several (e.g., 10) years; that is, their quality of transmission (QoT) is acceptable until their end of life (EoL) [2,3]. To guarantee reliability over long time periods, high physical-layer margins are used to cover the evolution of the physical conditions and related modeling uncertainties. In particular, the margins are used to cover statistical variability of equipment performance due to the manufacturing process, performance fluctuations from polarization effects, equipment aging, future fiber reparations after cuts, increasing interference due to increased network load, modeling errors, etc. Yet, high margins do not fully prevent outages, i.e., *hard failures*. These occur due to fiber cuts or equipment breakdowns and are recovered by the implemented resiliency strategy (protection, restoration) according to the related service-level agreement (SLA)/class of service. Between high margins and resilience strategies, a lot of the efficiency can be harvested.

In previous-generation wavelength-division multiplexing (WDM) optical networks, with fixed transponders and limited-flexibility optical switches, the static operation and high-margins approach were appropriate. However, EONs exhibit much broader optimization dimensions, which are wasted with the current static design approach [2,4]. Telecom operators must face steep traffic increases while maintaining the revenue, and they clearly need to reduce overprovisioning and increase the network efficiency.

A first solution is to remove long-term margins and certain uncertainties and maintain short to medium term margins. The margins are checked, using monitoring information and appropriate models, and adjusted at intermediate periods (i.e., on the order of years) to ensure reliability according to the related SLAs. As a result, connections are allocated enough margin to safely reach the next period—still considerably lower than with the current, static, approach. Moreover, uncertainties in estimation can be reduced. These translate to postponing or avoiding the purchase of regenerators or transponders, which leads to capital expenditure (CAPEX) and operational expenditure (OPEX) savings [5–7].

If we want to find a better trade-off between efficiency and reliability, i.e., if we want to further reduce margins, we need to make further steps in network automation. For example, we can envision operating a connection in a marginless manner, where we regulate its margins and bring its bit error ratio (BER) closer to the limit. Consequently, we might face cases where the connection would run out of margin. We refer to a QoT degradation as a *soft failure*. When margins are low, a soft failure would make the QoT unacceptable and the connection infeasible, and would be equivalent to a hard failure. This would in turn require overprovisioning for protection, defeating the initial purpose of cutting costs. To avoid that, we need a highly automated network following a feedback-based control loop, able to predict, detect, and diagnose health issues and automatically prevent them, either by directly acting at the failure cause or by downgrading capacity, diverting traffic, or ordering manual interventions, according to the service class [8–10].

To achieve an efficient optical network operation by regulating the margins, we first need to observe the network [11]. With the advent of coherent transmission, several physical-layer parameters can be monitored at the Rx via DSP. However, until recently, such information was not exploited for network optimization purposes. The vision of the European project ORCHESTRA (<http://orchestraproject.eu>) was to close and automate an observe–decide–act control loop. ORCHESTRA harvests the information of coherent transponders, using them as software-defined optical performance monitors (soft-OPMs). The feedback, the monitoring information, is transferred over a hierarchical monitoring plane and correlated using machine learning (ML)/data analytics algorithms to obtain accurate knowledge of the physical layer. This, in turn, is used to cross-optimize the network and the physical layers, regulate margins, and reduce overprovisioning. ORCHESTRA uses its observe–decide–act control loop to automate the required mechanisms for trading margin efficiency for reliability.

The rest of the paper is organized as follows. In Section 2 we discuss the margin efficiency and reliability trade-off and network automation. Then in Section 3 we present the ORCHESTRA network solution and continue by presenting its key modules. In Section 4 we present ORCHESTRA's feedback-based QoT estimators, which are key modules in understanding the actual conditions and estimating margins. In Section 5 we present a multiperiod planning study that quantifies the benefits of planning with accurate QoT estimation and reduced margins. In Section 6 we present the ORCHESTRA management and control plane, which implements the observe–decide–act control loop. In Section 7 we present field experiments performed on Telecom Italia's premises deploying ORCHESTRA management and control plane to demonstrate reliable, efficient, and automated network operation. Our conclusions are presented in Section 8.

2. MARGINS, EFFICIENCY, RELIABILITY, AND AUTONOMOUS NETWORKING

Previous works [2,3] classified the margins and discussed their impact on network operation. Margins are defined as the difference between the signal-to-noise ratio (SNR), or another QoT metric such as quality (Q) factor or BER, at the Rx at the network's current conditions and the related performance limit. According to [3], there are three types of margins: *system*, *design*, and *unallocated*. The system margins include the equipment aging and nonlinear interference (NLI) effects, which vary with time. As time passes, equipment such as transponders, amplifiers, filters, etc., age and their performance deteriorates. Fiber cuts are repaired, increasing the fiber attenuation. Moreover, typically as the network evolves, new connections are established to support new and increasing traffic demands, thus increasing the cross-channel NLI. In addition to these components of the system margin, we add the operator's margin and a margin for fast-varying impairments, both of which can be considered constant over time. The design margin is the second type of margin and accounts for QoT modeling inaccuracies. QoT estimation is a key operation performed when planning, upgrading, or reconfiguring the network. The inaccuracies come from two factors: (i) the inaccuracy of the estimation model itself, due to certain simplifications to reduce complexity and make the calculations efficient, and (ii) uncertainties of the input parameters of the model, since some of them cannot be accurately and/or continuously measured. Finally, the third margin type is the unallocated margin, which pertains to the mismatch of the demand and the capabilities of the transponders. This is already reduced with the existing-rate flexible/elastic transponders (with 50–100 Gbit/s granularity) and will diminish with next-generation finer granular transponders based on probabilistic constellation shaping or time-domain hybrid modulation techniques.

As expected, lowering the margins and increasing efficiency reduces the network costs. This has motivated various research studies. The authors of [6,7] studied the planning of an EON over multiple periods to take advantage of the evolution of margins over time. Instead of overprovisioning the lightpaths to ensure an acceptable QoT until the EoL [Fig. 1(a)], the

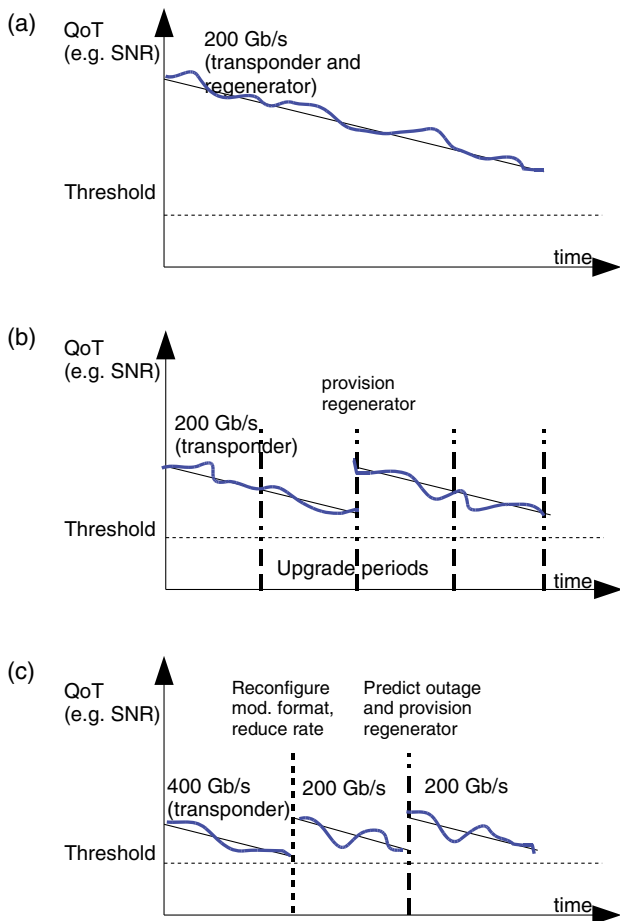


Fig. 1. (a) Traditional planning with end-of-life margins, (b) planning with reduced margins, (c) marginless operation of a 200 Gbit/s connection.

network is operated with just-in-time, or to be more precise, with just-in-period margins [Fig. 1(b)].

Note that the resulting savings come with no deterioration of reliability. The margins in this case are reduced in two ways. The first is the reduction of long-term/aging margins: instead of targeting acceptable performance at EoL (i.e., after 10 years), we target shorter time frames (e.g., 2 years). In this time frame, the network typically has some scheduled maintenance or upgrade window. Before such a scheduled window, we can recheck and re-evaluate the margins and adjust them by adding equipment (e.g., regenerator or a parallel line) or with appropriate reconfiguration actions. This is done to avoid any decrease in the service availability. Note that the long-term margins are not meant to cover hard failures, such as fiber cuts or switch breakdowns, which are covered by the resilience scheme. The long-term margins cover QoT degradations due to aging and increased load, which happen incrementally and slowly.

Second, we also reduce the design margin, by refining the QoT estimation model and reducing its uncertainties. For example, since the transmitter (Tx) laser may become detuned, especially as it ages, a certain part of the design margin accounts for the uncertainty in the alignment of the Tx laser and the filters on the path [12]. By developing a process to align the Tx

and filters, we remove the related uncertainty and safely remove the associated margin. Consider, as another example, the case where the QoT estimator requires knowledge of the dispersion coefficients of all fiber spans. These values would be known only within a certain accuracy, since they are hard to measure, especially in an operating network. Also, even if they were measured, the used equipment would introduce some error. The uncertainty of such a QoT estimator's input parameters is covered by the design margin, which would be reduced if we reduced the uncertainty of these parameters, e.g., with monitoring and data analytics techniques [2,13,14].

The efficiency increase from the aforementioned margin reduction methods requires a small change in operation mode, from the "set and forget" approach to a "check and upgrade" every couple of years, and yields important cost savings. One of the key modules is an accurate QoT estimator that understands the physical-layer conditions at each period and estimates the QoT for the new or reconfigured connections. The actual network condition is the reference on top of which we apply the aging model, so it should be understood to reduce the system aging margin from EoL to a few periods. By looking at the evolution over time, we can also refine the aging model. Moreover, by monitoring the network we can understand its conditions and improve the accuracy of the QoT estimator's input parameters, thereby reducing the design margin.

Thus, improving QoT estimation is a strategic step toward reliable and efficient optical networking. Tremendous research efforts have investigated the refinement of existing QoT estimation models [15,16] to account for lightpath specificities such as channel frequency, dispersion map [17], channel count, filtering impairments [18,19], linear cross talk, and polarization-dependent loss [20]. Efforts have also been undertaken to simplify the models, to reduce computation time, and thus to make them suitable for dynamic network operation. In parallel, QoT estimation can also be improved through a reduction of the uncertainties of the model inputs, as proposed in [13,14] and demonstrated in [21,22]. Sometimes, sources of QoT uncertainty can be removed by introducing some automation processes, e.g., to stabilize parameter values to their nominal values [12,23]. In a different approach called probabilistic design, the network can be designed to minimize the impact of the variations of parameters such as channel power on the QoT [24]. Recently, many works [25–27] have suggested that QoT could be accurately predicted without any form of analytical model, through *blind* ML using monitored parameters as inputs. Since such techniques rely on monitoring, they cannot be used in greenfield network design. Results suggest that blind ML approaches would only outperform model-based QoT estimators in cases where very little is known or when model input parameter uncertainty is extremely high. A discussion on model-based and ML techniques can be found in [14]. Since it has been demonstrated that monitoring can effectively refine model input parameters, it seems fair to state that blind ML approaches can only be useful when combined with a model to cover effects that are currently not sufficiently well modeled, although even this use case poses significant challenges in the training.

An EON can be operated even more dynamically, as it is possible to adapt the transmission parameters of transponders and reconfigure the ROADM switches at shorter time frames. We can reduce margins even further, going closer to the minimum performance threshold. This is known as marginless operation. This would increase the efficiency but could result in certain loss of reliability [8–10]. Considering the expectations of highly dynamic and volatile traffic with the emergence of fifth-generation cellular network technology (5G) and new applications, statistical multiplexing gains currently not present in optical transport networks can arise. In such an era, new use cases with classes of service with a wide variety of resiliency requirements also seem probable.

The benefits of marginless operation can take many forms, e.g., connections can be assigned higher rates than requested [Fig. 1(c)], cost in terms of transponders or spectrum can be saved by adjusting existing connections to meet varying traffic and/or harvest multiplexing gains, or energy can be saved by adjusting the symbol rate and/or the forward error correction (FEC) scheme. For example, assume that a connection carries both high- and low-priority traffic, and also assume that we consider as an option the downgrade of its modulation format, such as from polarization-multiplexing 8 quadrature amplitude modulation (PM-8QAM; 150 Gbits/s) to PM quadrature phase-shift keying (PM-QPSK; 100 Gbits/s) when its QoT degrades (soft failure). We should design the connection so that we have sufficient capacity for the high-priority traffic, i.e., the aggregated high-priority traffic constitutes no more than 100 Gbits/s. The lower-priority traffic would experience a downgrade in its rate until a new interface and transponder are provisioned. Alternatively, if the deterioration is foreseen and the interface and transponder were already deployed, the lower-priority traffic could be dynamically restored with minor disruption. Regarding the dynamic adaptation of client and line optical rates, the recently introduced Flexible Ethernet (FlexE) can provide the needed flexibility. Another solution is reported in [28] and relies on a sliceable transponder with a multiwavelength source, which is capable of adjusting the client and line rates.

We can also envision connections that consist of pure high-priority traffic and require their rate to be maintained. For such a connection, we could either provision it with enough margin to reach the next upgrade period [as discussed above; see Fig. 1(b)], predict and schedule an upgrade far in advance of when it is needed [Fig. 1(c) and [29]], or have a pre-computed reconfiguration option that guarantees the required rate (e.g., increase the bandwidth of the filters and push adjacent channels to reduce the filtering effect and interference), which can be done hitlessly (without traffic interruption). We could also have various types of lower-priority classes, which could accept certain service deterioration levels. To harvest the efficiency and reliability trade-offs of these options, we need further advances in network automation and, in particular, in failure identification, localization, prediction, and recovery mechanisms.

Therefore, control and management play very important roles in marginless network operation [25,30–33]. Such a mode of operation requires a closed control loop that continuously monitors, re-evaluates, regulates the margins, and

re-optimizes the network. The control plane has to handle new types of failures, since soft failures could make connections unacceptable. A critical phase when a problem is detected (e.g., a BER above a threshold) is failure identification and localization [8,10]. First, the control and management plane should be able to identify the type of problem, which could be due to any of several devices (e.g., amplifier, fiber, switches, filter, laser). As an example, the authors of [10] proposed a method of discerning the following causes: signal overlap, tight filtering, and signal drift. Then, with the knowledge of the failure type and its localization, the control plane can react to maintain the service. The problem of localization was addressed in [8,34–37] by employing multiple monitoring locations and correlating monitoring information.

When a traditional hard failure such as a fiber cut occurs, traffic can only be recovered if it is (re)routed over an alternative path. For hard failures arising from soft failures, however, simpler actions such as changing the modulation format can be taken to avoid rerouting. A summary of possible actions in the presence of soft failures such as fiber and amplifier aging and interchannel interference according to classes of service can be found in [30]. Changing the modulation format has also been adopted in space-division-multiplexed optical networks [31]. The works in [8,12,23] proposed techniques to recover from filter/laser alignment soft-failures. Recently, several works implemented recovery or reconfiguration of transmission parameters via the NETCONF protocol, since it has emerged for the control of disaggregated networks [31–33,38,39]. In these works, YANG models are proposed to enable NETCONF in the reconfiguration of the identified parameters.

Another important parameter is the speed, as well as the complexity, of recovery. It has been demonstrated in [30] that centralized solutions can present scalability problems, leading to delays in the processing of alarms and, in turn, in the reaction. Such solutions can exploit (i) on-demand requests for monitoring information to the network device, as in the case of [32], which adopted OpenFlow to obtain monitoring information, or (ii) notifications/alarms, as supported by the NETCONF Notification message [39]. Such protocol messages report to the central controller the value of a parameter monitored at a specific time instant. An alternative solution could be a telemetry stream between the device and the controller [33] to report a flow of real-time monitored values. Note that NETCONF has not been designed for “streaming,” and a widely adopted solution for telemetry is the gRPC protocol [40]. NETCONF and gRPC could also coexist and be used according to the network’s needs and applications. Another solution is to enable some distributed control actions when recovery speed and responsiveness are of utmost importance. Preprogramming [41] is a method that instructs/programs transponders to operate autonomously, e.g., by selecting the proper transmission parameters (modulation format, symbol rate, FEC, etc.) depending on the local monitored values, thus avoiding interacting with the central controller.

Finally, an interesting feature in network automation is failure prediction, which can be done at various time scales [29,42]. Medium-term failure prediction can be used to examine re-optimization/reconfiguration and restoration options

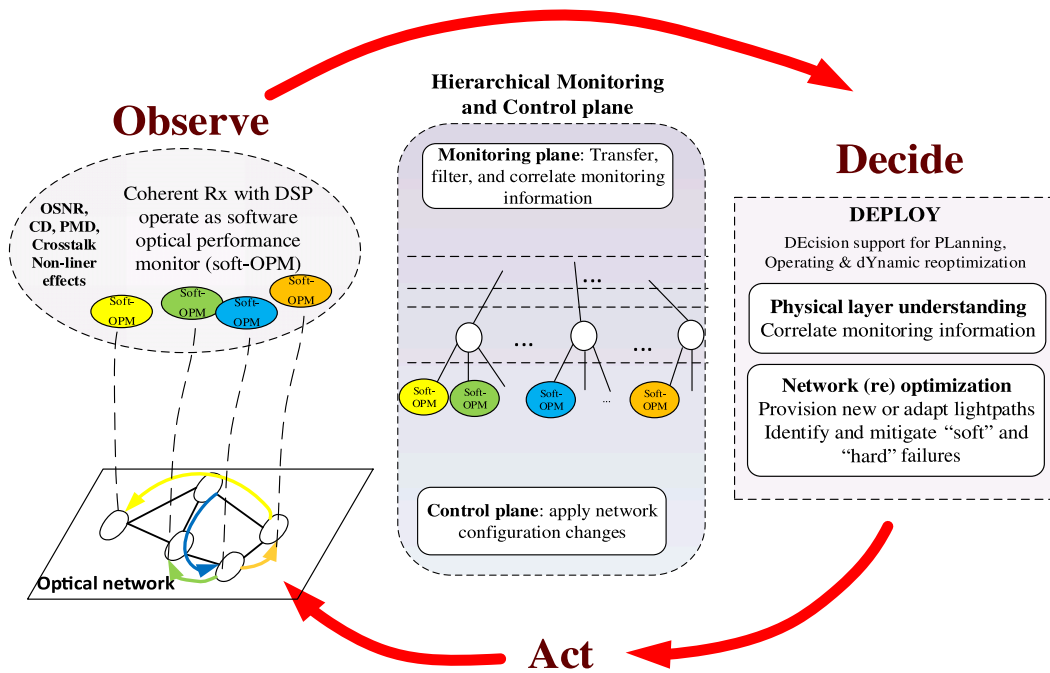


Fig. 2. ORCHESTRA's automated observe–decide–act control cycle. Coherent receivers operate as software optical performance monitors (soft-OPM), ORCHESTRA's hierarchical monitoring architecture carries this information, and the DEPLOY engine decides how to act and (re-)optimize the network.

and schedule a manual intervention (e.g., provision a regenerator) in case re-optimization cannot resolve the problem. Shorter-term prediction can be used for proactive restoration, thus improving the restoration efficiency.

3. ORCHESTRA SOLUTION

The vision of the European project ORCHESTRA was to close and automate the observe–decide–act control loop (Fig. 2), enabling higher network efficiency and dynamicity. The loop starts at the optical coherent transponders that are used as soft-OPMs. State-of-the-art impairment-monitoring DSP algorithms were integrated into a soft-OPM suite and interfaced with the ORCHESTRA monitoring and control plane. ORCHESTRA's hierarchical monitoring infrastructure could program the soft-OPM and select the monitored parameters, the monitoring algorithm, the period, etc., and also efficiently transfer and analyze/correlate data from multiple soft-OPMs. Custom-developed monitoring agents formed a hierarchy rooted at the central controller. The application-based network operations (ABNO) framework [43] was used as the reference central controller, playing the role of the global coordination entity, and also bridging the monitoring and control functionalities. Depending on the use case at hand, certain control actions were examined initially at a local level, i.e., at the monitoring agent of the connection. Then, if the problem could not be resolved, control actions were examined gradually at higher levels of the hierarchy for multiple and an increasing number of connections. The target was to increase scalability by keeping the complexity, intervention, and centralized processing low. The hierarchical monitoring and control architecture is described in more detail in Section 6.

Every transponder in the network can potentially serve as a soft-OPM. But we can do even more: a soft-OPM at a Rx provides aggregate measures over a multilink path. The ORCHESTRA optimization engine analyzes/correlates information from multiple soft-OPMs and other optical monitors (e.g., power monitors at optical amplifiers and switches), opening up a multitude of capabilities, such as accurate QoT estimation under actual network conditions (Section 4), detection, and anticipation and recovery from failures. This in turn is used to cross-optimize the network and physical layers. The advent of EONs vastly increased the optimization dimensions, while introducing new types of problems. ORCHESTRA relies on the feedback from soft-OPMs to feed cross-layer optimization algorithms for offline planning and also dynamic use cases. Lightpaths are provisioned and operated with optimum network parameters (routes, spectrum, modulation format, FEC, power, etc.), regulating their margins according to the related SLAs.

In particular, we developed multiperiod planning algorithms that exploit all of the optimization dimensions available in EONs (Section 5). These algorithms account for the actual physical network state, learned through the feedback and the QoT estimator, to provision lightpaths with reduced/just-in-time margins that are adjusted at intermediate periods to ensure their resilience level. Margins can be reduced even further, increasing the network's efficiency with a trade-off in reliability. To this end, we also developed dynamic algorithms that utilize the closed control loop and operate connections close to their minimum acceptable performance thresholds, regulating their margins and maintaining high efficiency, continuously, over an infinite time horizon (Section 7).

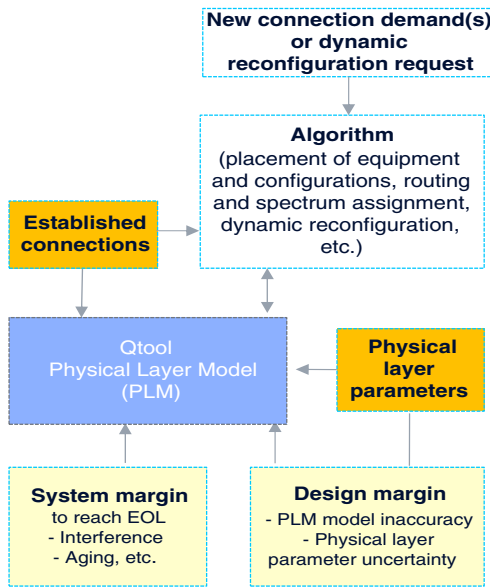


Fig. 3. QoT estimation tool (Qtool) based on a physical layer model (PLM) used by an algorithm to establish or reconfigure lightpaths, inputs, and the design and system margins.

4. ACCURATE QoT ESTIMATION

As discussed, an accurate QoT estimator is key for (i) reducing the margins during planning/upgrades and (ii) realizing more dynamic operation of the network. Figure 3 shows how a QoT estimation tool (or Qtool) is used by an optimization (planning or dynamic) algorithm. The Qtool is typically a physical layer model (PLM), an analytical or semi-analytical model of the physical layer based on certain assumptions that estimates the QoT of new or reconfigured connections with a certain accuracy (e.g., SNR, BER). The Qtool takes as input the parameters of the established connections (in the case of an operating network), and also certain physical-layer parameters, such as spans, fibers, amplifiers, and node parameters. The values of the physical-layer input parameters are not accurately known due to a lack of measuring equipment, or to their limited accuracy, outdated measurements, etc. To cover the model and input parameter inaccuracies, the design margin is used (2 dB in the SNR, as in Refs. [2,3]). Moreover, traditional network design is typically performed such that new connections will have an acceptable QoT at EoL, e.g., after 10 years. Modeling equipment aging, increasing interference, reparations of fiber cuts, etc., contributes to the system margin (3 dB in the SNR as a reference).

In ORCHESTRA we use feedback from the physical layer and ML to improve the accuracy of QoT estimation and reduce the design and system margins (Fig. 4) [13,14]. By understanding the actual network conditions, we improve the accuracy of the parameters used as input in the estimation and thus reduce the design margin. Moreover, this operation is repeatable, so we can target shorter time scales, i.e., on the order of a few years.

In the following we discuss two QoT estimators we developed that complement each other (see Section 4.C). The first is used in network upgrading, where we establish new

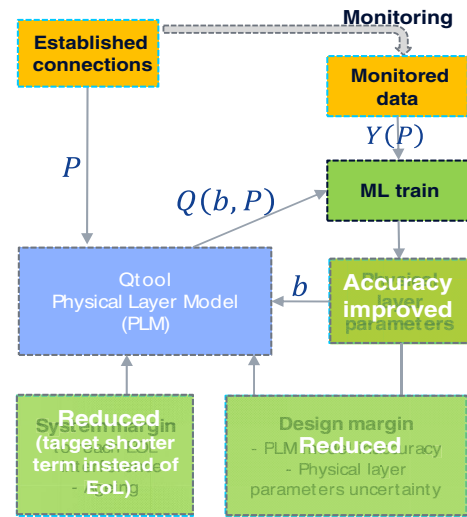


Fig. 4. Machine learning physical layer model (ML-PLM) for reducing the margins.

connections, while the second is used in operation, where we dynamically reconfigure established connections.

A. Offline QoT Estimator for Network Planning

We assume that we obtain monitoring information from the coherent Rxs deployed in the network. In particular, we denote by $Q^*(P)$ the vector that contains the SNR values of all established paths $p \in P$. The mapping $Q^*(P)$ is nonlinear and unknown to us. For any set of established lightpaths P , we denote by $Y(P)$ the vector of their monitored SNR values. The monitoring error consists of a systematic and a random error. The systematic error can be reduced through proper calibration, while the random error can be reduced by averaging over time (on a shorter time scale than a natural change of the monitored value would occur). As monitoring errors are small and can be reduced, we ignore them here for simplicity. Therefore, we will assume we monitor the true QoT for the lightpaths ($Y(P) = Q^*(P)$).

We consider the case where a new lightpath $w \notin P$ is about to be established, and we want to estimate (i) the QoT of that new lightpath w and (ii) the effect of establishing w on the QoT of the existing lightpaths $p \in P$. QoT estimation, our goal is to identify a parametric function $\tilde{Q}(b, P)$ that accurately approximates the actual QoT function $Q^*(P)$. Here b is a set of parameters of the model. The function $\tilde{Q}(b, P)$ does not have to be a closed-form expression; it can be the output of a computation program or a simulation. What is important is that (i) $\tilde{Q}(b, P)$ approximates $Q^*(P)$ well, and (ii) given the vector b , it is computationally easy to obtain $\tilde{Q}(b, P)$.

We consider a PLM as the approximating architecture. In this case, $\tilde{Q}(b, P)$ takes as input b the PLM input parameters, which can include the length, fiber attenuation, dispersion and nonlinear coefficients per span, the noise figure of each erbium-doped fiber amplifier (EDFA), its gain, etc. In our studies, we used the Gaussian noise (GN) model [15] as the PLM, and in particular in our implemented GN model, we

account for the “actual” cross-channel NLI. To be more specific, we calculate the noise power of the NLI (both self and cross-channel) for each span, for the actual channel spectrum position, using Eqs. (128) and (129) of [15]. Then, assuming incoherent noise accumulation and using Eq. (127) of [15], we accumulate the NLI noise power over a link and over the whole path. Although we used the GN model, our ML method is generic and applicable to other PLMs.

The objective of the ML-PLM is to use monitoring information from the established lightpaths $Y(P)$ to train and learn the physical-layer input parameters and thus improve the accuracy of future lightpath QoT estimations. The model also uses the set of lightpath parameters that are included in the network state P , which can be the route, central frequency, symbol rate, modulation format, launch power, etc. The parameters in P are assumed to be known perfectly, as opposed to the parameters in b . Regarding the learning algorithm, if we have closed forms for the partial derivatives $\partial \tilde{Q}/\partial b_j$ for all the physical-layer parameters b_j of b , we could use the gradient descent method to obtain better b estimates, as is done in [13]. Here we assume a generic case where $\tilde{Q}(b, P)$ is unknown. Assuming that QoT depends nonlinearly on some input parameter, we can use a nonlinear fitting method, such as the Levenberg–Marquardt algorithm. If we start far away from the optimum, the method might get stuck in a local minimum, so we can repeat from different starting points.

We start with some default initial physical-layer parameter vector b^0 for each span (e.g., fiber attenuation coefficient, EDFA noise). In each iteration i we obtain the updated parameters b^i that minimize the distance, defined as some function of the error $Q(P) - \tilde{Q}(b^i, P)$ [e.g., the mean squared error (MSE)]. The iterations stop and we obtain the final b^* values when certain criteria are met, e.g., the error function reaches a predefined threshold or a certain number of iterations are executed. Figure 4 shows a schematic of the ML-PLM. When we want to establish new lightpaths, we have a new set P' , and we estimate $\tilde{Q}(b^*, P')$. This estimation is done with more accurate parameters b^* , which, as discussed above, improves the estimation accuracy and reduces the design margin.

In the following we present an evaluation of the accuracy of the ML-assisted QoT estimator, ML-PLM, described above through simulations. We considered the Deutsche Telekom (DT) topology with 12 nodes and 40 bidirectional links. We assumed four traffic loads of 100, 200, 300, and 400 total connections with uniformly chosen source/destination pairs and random symbol rates from the set {32, 43, 56} Gbaud. Regarding the physical layer, we assumed a span length of 80 km, an EDFA noise figure of 5 dB, and standard single-mode fiber (SSMF) with a mean attenuation coefficient of 0.23 dB/km, a mean dispersion coefficient of 16.7 ps/nm/km, and a mean nonlinear coefficient of 1.3 (W km)^{-1} . We set the launch power at 0 dBm. The actual (unknown) values of the fiber coefficients for each span were drawn from uniform distributions ranging by 0%, 10%, or 20% around the above means. The GN model, which accounts for the actual cross-channel interference, was used as the ground truth with these values. For the ML-PLM estimator, the b vector consisted of the attenuation, dispersion, and nonlinear coefficients, initiated with

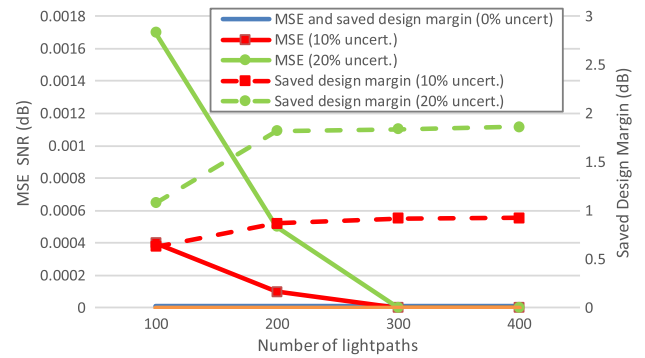


Fig. 5. Accuracy performance of the offline QoT estimator. MSE and saved design margin for 0%, 10%, and 20% uncertainty as a function of the number of connections (lightpaths) established in the network.

their mean values. Thus, the case with a 0% range implies that the ML-PLM estimator has perfect knowledge of the physical-layer parameters. The maximum error made by the ML-PLM before being trained (so using the mean values) was 0 dB, 1 dB, and 1.9 dB for 0%, 10%, and 20% uncertainty, respectively. The training of ML-PLM was done with nonlinear regression, and in particular the Levenberg–Marquardt algorithm.

For each traffic load and uncertainty setting, we executed 500 iterations. For each instance we used 85% of the lightpaths for training and 15% for testing. We did not use cross validation, since it was not deemed necessary to compare and select the ML model or some parameter of the model. We excluded from the testing sets the lightpaths that include links for which we have no QoT information. The training goal was to minimize the MSE, but the max overestimation is also a very useful metric because it defines the design margin. This has to be used to be on the safe side, so that we never overestimate the QoT and establish a lightpath with an unacceptable QoT.

In Fig. 5 we see that ML-PLM achieves a very good MSE. As expected, ML-PLM is affected by the uncertainties. At perfect knowledge (0%), the fitting algorithm starts from the correct values and immediately stops with zero error estimations. As uncertainties increase, the ML-PLM yields a higher MSE and max overestimation error for a specific load/number of established and monitored lightpaths. Viewed differently, it requires a higher number of monitored lightpaths to obtain the same level of accuracy. The ML-PLM’s max overestimation was on the order of 0.1 dB for more than 200 lightpaths (not shown in the graph). The training time was on the order of minutes for the ML-PLM/nonlinear regression. Once the model is trained, estimations are quite fast (<0.1 s). Detailed results are presented in [14].

B. Online QoT Estimator for Network Operation

We now focus on a specific established connection (or lightpath). To determine whether a lightpath reconfiguration is feasible, the typical method consists of estimating its pre-FEC BER and comparing it to the associated FEC limit. This is typically done with the PLM, which, as discussed, takes various inputs (e.g., amplifier noise factor, fiber nonlinearity, channel

power), which are independent of the Rx. So the PLM calculates the generalized optical signal-to-noise ratio (gOSNR) [15] before the Rx, where the gOSNR is an extension of the OSNR that includes, in addition to amplifier noise, all other optical noise [e.g., from the Kerr effect (nonlinearities) and linear cross talk]. Then a typical assumption is a linear function to translate the gOSNR to SNR and closed-form formulas to move from the SNR to BER.

Instead of relying on a PLM to estimate the gOSNR of an established connection, we estimate it *live* using monitored parameters. The accuracy of the gOSNR estimation is therefore independent from the knowledge of the PLM model inputs. In this way we consider a more generic function for the transformation of the gOSNR to BER that can capture the particularities of the transponder at hand such as Tx/Rx imperfections, nonideal performance of DSP for different modulation formats, and some advanced impairments that are typically covered by margins, e.g., filtering effects.

To formulate this, assume an established connection p , where p represents the state (specific modulation format, symbol rate, routed over a path that crosses a set of nodes/filters), and monitored parameters $Y(p)$ (BER and others such as the 10 dB signal bandwidth) that we want to reconfigure. In this process we want to estimate the BER at another state p' , which includes one or more adaptations such as a change in its modulation format or an increase in the bandwidth of the filters (to reduce filtering and interference effects). To do so, we use a function f ,

$$\text{gOSNR} = f(Y(p), p),$$

that considers the current state p and monitored values $Y(p)$ and estimates the current gOSNR. Then we apply the inverse function $f^{-1}(\text{gOSNR}, p')$ to estimate BER' at state p' . Function f can be *learned* with ML by fitting data captured in transponder calibration in back-to-back and basic filter configurations, to cover estimates for various modulations formats, symbol rates, and filter configurations.

Figure 6 shows the fitted BER–gOSNR curves for three modulation formats supported by the ORCHESTRA prototype transponder. Figure 6 also shows an example of moving from a monitored BER in PM-QPSK to an estimated BER' in PM-8QAM and 16QAM. For this fitting we used only the monitored BER as $Y(p)$. We also extended the model to capture filtering penalties on symbol rate and allocated spectrum (filter bandwidth) reconfigurations. For that fitting we used the monitored BER and the 10 dB bandwidth provided by the coherent Rx. Further details are presented in [8].

C. QoT Model Synergy

The online QoT estimator described in Section 4.B complements the offline QoT estimator described in Section 4.A. The offline estimator correlates information about established connections to determine the actual network conditions and improve the accuracy of gOSNR/SNR estimation for new connections. It is agnostic to the specifications of the transponder to be used and considers the filtering effect and other impairments with a certain accuracy. Thus it reduces the margins compared to previous approaches but still allocates some to




Fig. 6. Fitted BER–gOSNR for three supported modulation formats of the ORCHESTRA prototype transponder.

capture such inaccuracies. The online estimator targets the reconfiguration of an established connection that uses a specific transponder, or that crosses specific filters, etc. Therefore, that model works on top of the gOSNR and captures the particularities that are left out of the offline model, further reducing the already reduced margins.

5. PLANNING WITH REDUCED MARGINS

If the network is planned with high margins, which is the traditional approach, the QoT estimator uses a high design margin and EoL system margin. Thus, the efficiency is low and more transponders/regenerators are placed, but lightpaths are certain to be uninterrupted, and to have an acceptable QoT, until EoL. Instead, if we leverage an accurate QoT estimator, we can use an appropriate algorithm to plan the network with reduced margins. In the first period, we will serve the new demands with a high design margin (no feedback/no established connections to monitor, train, and refine the QoT estimator) and with a reduced system margin to reach the next upgrade period (several months to a few years, but sooner than EoL). At an intermediate period, we train the QoT estimator, and we reduce the design margin. Therefore, when planning for the next period, the incremental planning algorithm performs two tasks: (i) it checks the remaining margins of previously established connections and reconfigures/adds transponders or regenerators to restore those that have run out of margin (will have unacceptable QoT performance), and (ii) it serves the new demands by placing transponders/regenerators. In both cases it chooses the configuration of the transponders/regenerators, and decides the routes and the spectrum allocation, by interacting with the QoT estimator to check the physical-layer performance. An algorithm for planning the network over multiple periods with reduced margins is presented in [7].

To quantify the benefits of the ORCHESTRA approach, the accurate offline QoT estimator developed (Section 4.A), and the incremental planning algorithm with low margins, we dimensioned a network over multiple periods and calculated the CAPEX at each period. We compare that with planning with EoL margins.

The network topology was the 12-node DT topology, as above. We planned the network over 11 periods (the initial/greenfield and 10 incremental/brownfield periods); a period would roughly correspond to a year. The initial traffic (period τ_0) consisted of 200 connections with a uniformly

chosen source–destination and uniformly demanded traffic between 100 and 200 Gbits/s. Thus, the initial traffic was ~ 30 Tbps, which was increased by 20% each period by creating new demands in the same way. We assumed two types of elastic transponders (ETs): (i) 32 Gbaud, modulating with DP-QPSK, DP-8QAM, or DP-16QAM and supporting a capacity of 100, 150, or 200 Gbits/s, respectively, and (ii) 64 Gbaud, modulating from DP-QPSK to DP-32QAM and supporting a capacity of 200 to 500 Gbits/s, respectively. The first ET of 32 Gbaud was assumed to be available at the initial period τ_0 with a price of 1 cost unit (CU), and the second ET of 64 Gbaud was assumed to be available at period τ_5 with a price again of 1 CU (at that period). Prices were assumed to fall by 10% per period (so when the second ET is introduced, the first ET costs 0.59 CU).

We again used the GN model to model the physical layer. For the first period τ_0 we initialized the model with heterogeneous span parameters, similar to those in Section 4. The actual (unknown) values of the fiber coefficients for each span were drawn from uniform distributions ranging by 10% around the aforementioned means. We studied 10 problem instances with different initial traffic and span parameters and averaged the results. To model the aging of the network, we considered the following: increase of fiber attenuation (e.g., due to cuts) and aging of ETs, EDFAs, and nodes [optical cross-connects (OXC)]. The interference was modeled according to the network load. Table 1 shows the increase in the parameters' values per period. Note that the increase was uniform for all spans, but since we started from a heterogeneous network model, the network remained heterogeneous for all subsequent periods.

When planning the network with EoL margins, the Qtool uses a system margin based on the parameters of Table 1, assuming 10 periods and a full network load (each link having 60×32 Gbaud connections), which in total is about 3 dB. The design margin was set to 2 dB (1 dB for model inaccuracy and 1 dB for input parameter inaccuracy). When planning the network with reduced margins, the system aging margin was based on the parameters in Table 1, assuming two periods (~ 0.6 dB) and also actual interference (increasing from ~ 0.5 to 1.5 dB according to load). Then we assumed that at each period we monitor the network before the upgrade and obtain the SNR values of the established connections (the monitored values were calculated by the GN model, with the random created initial parameters and aging according to Table 1, which was unknown to the QoT estimator). We used the monitored SNR values to train the ML estimator. The design margin was set equal to 1 dB for the model inaccuracy, plus 0.2 dB

Table 1. Parameters to Model the Network Aging

Physical-layer parameters evolution			Increase per period
System margins	Aging	Transponder margin (dB)	0.05
		Attenuation (dB/km)	0.0015
		EDFA noise figure (dB)	0.1
		OXC loss (dB)	0.3
		Interference	According to load

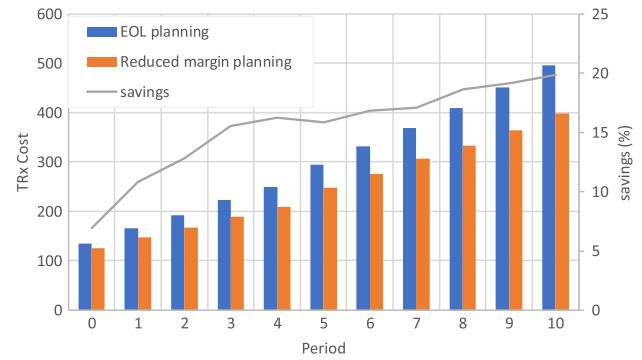


Fig. 7. Total cost (in cost units) and savings of deployed elastic transponders (ETs) per period when planning with accurate QoT estimation/reduced margins and with high margins.

(the max overestimation after the extensive simulations in the previous section), plus the max training error. Note that the choice of the system margin for two periods and the design margin greater than 1.2 dB are conservative. In total we harvest ~ 3 dB compared to planning with EoL margins over all periods (noting that interference increases but the design margin falls as the load increases).

Figure 7 presents the total cost of the deployed ETs per period for the two provisioning approaches. As expected, reducing the margins yields lower costs. Reducing the system margins postpones the purchase of ETs, and we obtain savings from the 10% depreciation. Reducing the design margin avoids the purchase of equipment. Figure 7 also presents the relative savings, found to be about 20% at the end of the examined periods. Note that the savings would flatten or even decrease at later periods if we assumed a single type of ET for all periods. Also note that in all simulations (10 instances \times 11 periods), we never observed a QoT problem; all lightpaths had an adequate QoT to reach the next period. The closest we came to QoT blocking was at period τ_5 , where the new ET (and thus a new symbol rate) was introduced, and some parameters were not learned accurately enough, having only the first ET (32 Gbaud) in the network. A factor not included above is the time value of money; money saved at intermediate periods can be invested (or loans can be avoided), resulting in extra savings. Additional savings can be obtained by power optimization [7].

6. ORCHESTRA MONITORING AND CONTROL PLANE: DESIGN AND PROTOTYPE

The ORCHESTRA project designed and implemented a control and management plane reflecting the main functionalities summarized by the Internet Engineering Task Force (IETF) ABNO architecture [43]. It included optical connection provisioning, path computation exploiting databases storing traffic engineering (TE-DB) and circuit information [label switched path (LSP-DB)], and operation administration and maintenance (OAM) procedures. OAM is one of the core functionalities of ORCHESTRA, since it involves the reception of monitoring information, the reception of alarms, their correlation, and the triggering of actions to preserve the service. According to ABNO, an entity named OAM Handler is in

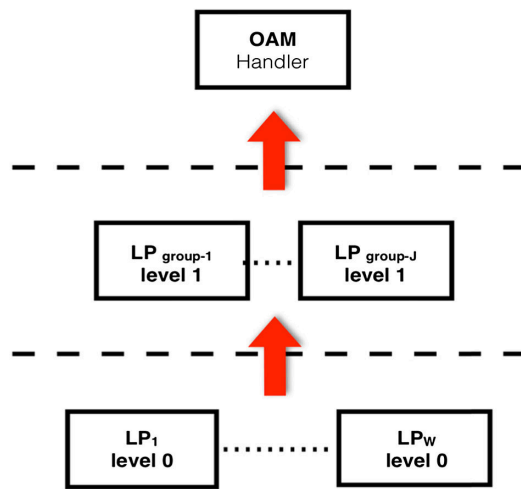


Fig. 8. ORCHESTRA's hierarchical management plane for the exchange and processing of monitoring information.

charge of OAM functionalities. To overcome scalability issues in the management of monitoring information and alarms, as discussed in Section 2, ORCHESTRA proposed a hierarchical architecture for the management plane [30], as shown in Fig. 8. In this architecture, the OAM Handler functionalities are spread among monitoring entities in the hierarchy that are assigned limited responsibilities.

Each Level 0 entity is responsible for a single lightpath. It can be programmed (i) to send monitoring information on demand and (ii) to create and send an alarm when a specified threshold is exceeded (e.g., a BER above threshold). Monitoring data and alarms are sent to a specific Level 1 entity. A Level 0 entity can also make local decisions (e.g., a change in FEC) to recover from failures that do not affect other entities. Then, each Level 1 entity has the visibility of a group of Level 0 entities/lightpaths (we assumed grouping per ingress node). It processes, correlates monitoring data and alarms from those lightpaths, and takes appropriate actions; forwards messages to a specific Level 2 entity; or reconfigures one or more lightpaths for which it is responsible, without affecting lightpaths not belonging to its responsibility set. Higher levels have similar functionalities for larger sets of lightpaths. This type of approach has been demonstrated to increase the scalability of network management [30] and reduce the delay in recovery [41]. Finally, the OAM Handler at the root of the hierarchy is responsible for the whole network and can take actions if lower levels are not able to recover from a failure or re-optimize the network accordingly. The hierarchical architecture can be expanded to account for node and link monitors (e.g., power monitors), which would follow the same concept of limiting responsibilities [30]. In the following we focus on end-to-end/lightpath monitors.

Network (re-)configuration and monitoring information dissemination in the hierarchical architecture is performed with the NETCONF protocol. Within the ORCHESTRA project, several YANG models for NETCONF were proposed to control and manage network devices. One of the most relevant devices is the flexible/ET, whose implemented YANG model is shown in Fig. 9. The parameters that can




Fig. 9. Transponder YANG model: (a) configuration parameters, (b) monitoring parameters.

be configured such as the rate, modulation format, central frequency, and output power are shown in Fig. 9(a). Monitored parameters associated with the coherent Rx and physical-layer performance, such as pre-FEC BER, polarization mode dispersion (PMD), and chromatic dispersion (CD), are shown in Fig. 9(b).

Regarding NETCONF, we implemented the following messages: <edit-config> to (re-)configure devices such as the transponder, <get> to request on-demand monitoring information (e.g., to Level 0 entities), <rpc-reply> to respond with the monitored parameters' values, <subscription> to enable an entity (e.g., at Level 1) in the hierarchical architecture to be notified when a problem occurs on a lightpath under its responsibility, and <notification> to implement the alarm based on the aforementioned subscription. Note that periodic monitoring can be programmed on top of these messages, at the OAM Handler level. So the OAM Handler could be programmed to periodically poll specific elements for specific parameters with appropriate <get> messages.

An integrated prototype of ORCHESTRA monitoring and control plane was developed and experimentally validated in lab experiments and in a field trial to be discussed in the following section. This prototype includes an ABNO controller extended to support ORCHESTRA use cases, a custom hierarchical monitoring architecture and the OAM Handler implementation (as discussed above), extended databases to store physical-layer monitoring information (PL-DB), the DEPLOY module acting as the path computation element (PCE) and QoT estimator, and a provisioning manager to enforce control actions through a software-defined networking (SDN) controller [in particular, OpenDaylight (ODL)].

The monitoring plane is implemented in software as an N -level infrastructure, composed of a minimum of $N=3$ levels: Level 0 maps to lightpaths/end-to-end soft-OPM monitors, Level 1 maps to ingress nodes (grouping all Level 1 monitors of the lightpaths starting at that node), Levels 2 to $N-1$ map to a specific group of nodes (e.g., network geographical regions, and ultimately a network domain), and Level N maps to the OAM Handler, as the root of the hierarchy. Monitor Level 0 interfaces with the Rx DSP as the lowest management point in the hierarchy. It can provide monitoring information on demand and also be configured to generate alarms when a DSP-monitored parameter violates a threshold (e.g., an excessive BER). Monitors from Level 1 to Level $N-1$ follow a generic architecture composed of (i) a NETCONF client to interact with lower monitoring entities, (ii) a correlation engine to process received alarms and produce (if possible) a cumulative notification for higher entities, and (iii) a NETCONF server that sends notifications to higher

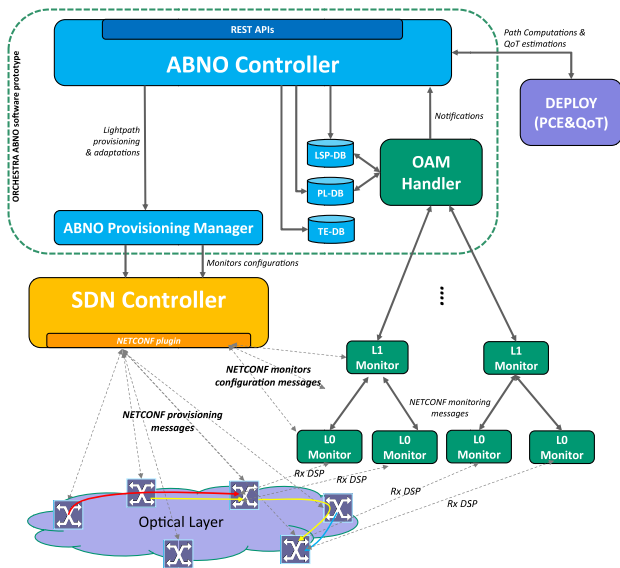


Fig. 10. ORCHESTRA ABNO prototype and external control and monitoring entities.

entities. The OAM Handler was interfaced with the hierarchical infrastructure and plays the role of its root. It opens the NETCONF sessions toward lower monitoring entities, subscribes to streams, processes and correlates notifications, and forwards aggregated notifications to the ABNO controller.

Through ABNO's northbound representational state transfer application program interfaces (REST APIs), the network operator can set up, tear down, and reconfigure lightpaths. Such requests are translated into NETCONF messages to configure (i) the transponders and the switches to set up the lightpath and (ii) the monitoring entities to create/adjust the hierarchical monitoring infrastructure. Failure handling use cases are triggered through alarms, which are properly aggregated and escalated to the OAM Handler (when needed) and then managed by related ABNO controller workflows to apply control reactions as instructed by the DEPLOY tool. Thus, the ORCHESTRA observe–decide–act control loop achieves automated recovery from various failure scenarios, yielding a network with self-healing capabilities. Figure 10 presents the schematic of the ABNO prototype that was used in the lab and field experiments.

An innovative control method proposed in the framework of ORCHESTRA is preprogramming [41]. Such a method can work both in a hierarchical and a classical management

plane architecture. Preprogramming was proposed to further increase the scalable management of monitoring information and alarms. According to this scheme, transponders are instructed/programmed to autonomously select the proper transmission parameters such as the modulation format, symbol rate, and FEC depending on the monitored parameters' values. This way, there is no need to send an alarm to the central controller and wait for computation because devices have been programmed to react. Such an approach was demonstrated in [41] to increase alarm scalability management and to speed up lightpath recovery. In the next section we also present a demonstration of this approach.

7. DYNAMIC CONFIGURATION EXPERIMENTS

The ORCHESTRA dynamic network operation was demonstrated in lab experiments and also in a field trial on Telecom Italia (TIM) premises in Torino. The cable used in the field trial is deployed between the Torino (Stampalia) and Chivasso network exchanges and is composed of eight 76-km-long G.652 fiber spans (Fig. 11). We used five spans to create two links of three and two spans, denoted as I1 (228 km) and I2 (152 km), respectively. Two lab-hosted ROADMs switched traffic from I1 to I2. The Tx supported QPSK, 8QAM, and 16QAM modulation formats and 28 and 32 Gbaud baud rates to achieve net capacities of 100, 150, and 200 Gbits/s with 12% [low-FEC (LF)] or 28% [high-FEC (HF)] coding rates. The Rx consisted of a 40 Gsample/s oscilloscope and offline DSP processing to handle both signal demodulation and monitoring of multiple transmission parameters.

The Rx DSP was interfaced with the ORCHESTRA Level 0 monitoring agent to report the monitored values and raise alarms. The various aging scenarios were emulated through a combination of a tunable optical filter (TOF) and a variable optical attenuator (VOA). The Level 0 agent was connected to a three-level hierarchical infrastructure (rudimentary in this experiment, since only one lightpath was used) and the OAM Handler. The ABNO controller included the OAM Handler, the databases, the DEPLOY tool as an online QoT estimator and PCE, and the ODL controller as the provisioning manager (see Section 6 and Fig. 10).

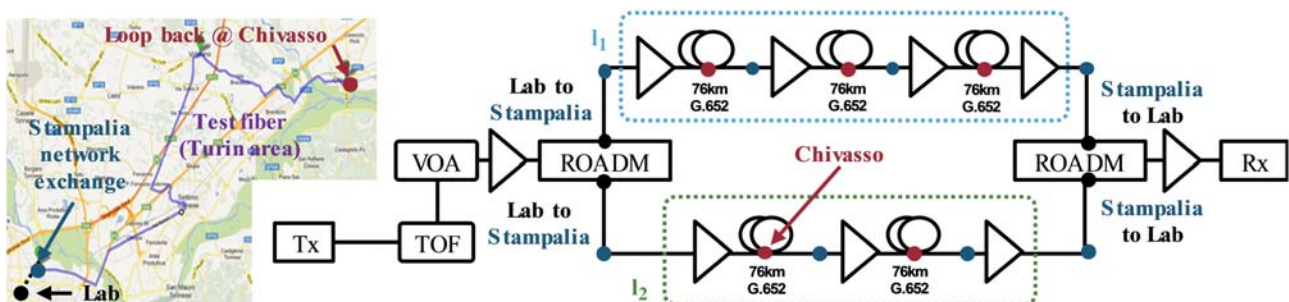


Fig. 11. Field trial setup at the TIM premises in the Turin region, Italy.

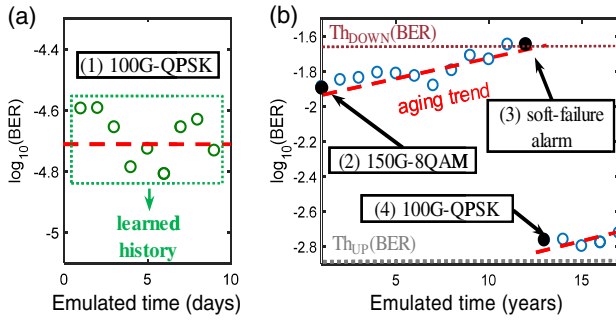


Fig. 12. Automatic rate adaptation in the presence of emulated fiber/amplifier aging: (a) learning of short-term performance fluctuations and (b) evolution over a long time.

A. Rate Adaptation under Fiber/Amplifier Aging

As the first use case, we considered the rate adaptation of a connection. This could correspond to a best-effort connection that operates close to its FEC limit and can reduce its net rate as time evolves, or a connection with a mix of high- and low-priority traffic that guarantees 100 Gbits/s and can make use of a higher rate for best-effort traffic if available. Fiber/amplifier aging is emulated by increasing the optical attenuation using the VOA (Fig. 11). As a consequence, the OSNR falls, where on top of the aging, long-term-trend and short-term-performance fluctuations naturally occur in the field.

Prior to the rate adaptation experiment, we measured nine points of BER and OSNR back-to-back for each modulation format (Fig. 6). This was done to train the online QoT estimator residing in DEPLOY to convert BER into gOSNR and also do the opposite, as described in Section 4.B. Thanks to these BER–gOSNR conversions, we could define appropriate thresholds for reconfiguring the modulations formats.

At its beginning of life (BoL), the lightpath operated at 100G-QPSK, indicated by state “(1)” in Fig. 12(a). During a short period, i.e., a few days or weeks of operation (minutes in the lab), DEPLOY learned the BER average μ and standard deviation σ in state (1) [Fig. 12(a)]. Based on those and the soft FEC threshold $TH_{SF} = 0.02$, DEPLOY calculated the threshold $TH_{down} = TH_{SF} - k \cdot \sigma$ to reconfigure from 8QAM to QPSK, where σ is the learned standard deviation and $k = 4$. It also calculated the threshold $TH_{up} = f^{-1}(f(TH_{SF}, p_{QPSK}), p_{8QAM}) + k \cdot \sigma$, to reconfigure from QPSK to 8QAM, where $f^{-1}(f(TH_{SF}, p_{QPSK}), p_{8QAM})$ makes use of the fitted modulation curves f discussed in Section 4.B. The reason for using k is to avoid a ping-pong effect; the QoT performance experiences certain natural short-term fluctuations, which could make the connection fluctuate between the two modulation formats. To avoid this, we introduced a hysteresis: we increased/decreased the thresholds accordingly to cover some performance variations. In this context, parameter k balances the modulation format stability and the minimal margin operation. The calculated thresholds were $TH_{up} = 5.8 \cdot 10^{-5}$ to switch from QPSK to 8QAM and $TH_{down} = 0.01998$ for the opposite. Based on the observation of $BER = 2.5 \cdot 10^{-5}$ and the related thresholds, the lightpath was reconfigured to 150G-8QAM, state “(2)” in Fig. 12(b),

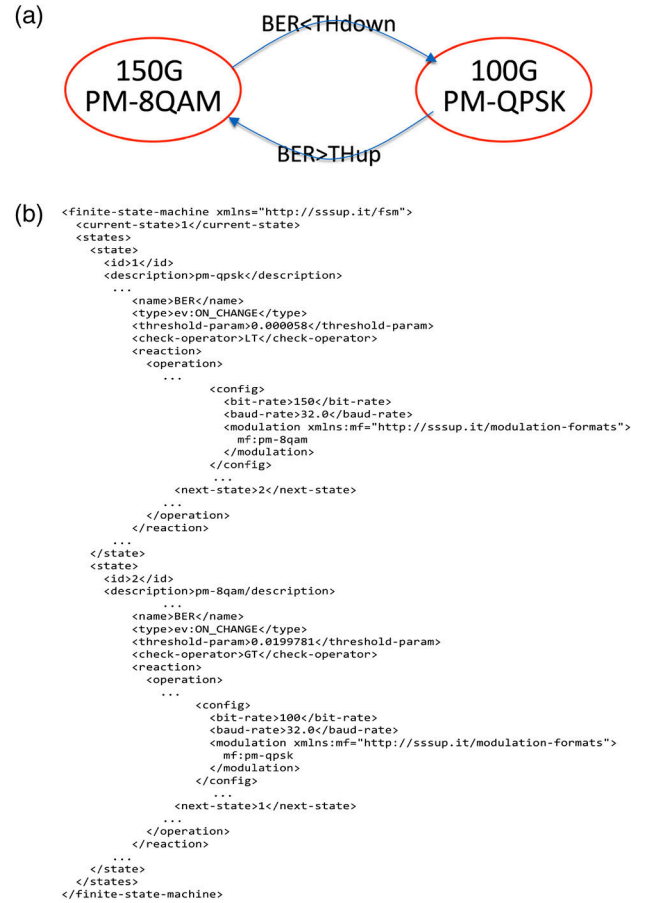


Fig. 13. FSM for automatic rate adaptation: (a) schematic and (b) configuration message snippet.

and reached a BER of $1.3 \cdot 10^{-2}$ [cf. Fig. 12(b)], within 0.5 dB of the assumed design margin.

Then the BER increased progressively due to emulated aging. Once the BER crossed the related threshold in state “(3)” of Fig. 12(b), we switched back to 100G-QPSK, corresponding to state “(4)” in Fig. 12(b). At state (4), the hysteresis scheme prevented the system from reverting back to state (3), despite natural BER fluctuations that might temporarily validate such a reconfiguration. The connection remained in QPSK until the network’s EoL.

If we look at the control plane side, the above experiment was performed with the use of preprogramming. Initially, the lightpath was established with a northbound interface (NBI) call to the ABNO controller. Then through the NBI we preprogrammed the automatic rate adaptation of the lightpath. The ABNO controller obtained a certain number of monitored values to calculate the mean and standard deviation. Then it calculated the thresholds TH_{up} and TH_{down} , created the finite-state machine (FSM) for the dynamic rate adaptation, and installed that to the Rx agent. The Rx agent monitored the BER, compared it to the current state’s threshold, and automatically switched to the next state when the condition was satisfied. The central controller (ABNO) was informed about the modulation format adaptations but did not participate in those decisions/actions. A schematic of the FSM is shown in




Fig. 14. Connection's BER and BW as a function of time, under increasing filter penalty and reconfiguration.

Fig. 13(a), and the snippet of a configuration message is shown in Fig. 13(b).

B. Rate Maintenance under ROADMs Aging

In this second use case, we considered a gold-class 100G-QPSK connection that needs to maintain its net capacity. This can be achieved through FEC (hence the symbol rate) and/or slot/bandwidth size adaptation. Here we considered the increased penalty due to the progressive detuning of crossed filters (in the ROADMs) and the Tx laser frequency over the years. In this context, we exploited the capacity of the online QoT estimator (DEPLOY) to accurately predict the BER based on monitored bandwidth (BW) and BER, as outlined in Section 4.B and [12].

We emulated 10 cascaded filters and their aging through progressive reduction of the bandwidth of a TOF (Fig. 11). At BoL, the gold-class 100G-QPSK connection was established with high soft FEC (HF = 28%) and three flex-grid slots, that is, 37.5 GHz (Fig. 14). When, because of filter aging, the HF BER threshold was reached, the online QoT estimator

```

2017-12-15 12:23:30,800 AbnoController:502 - Received message from
queue
{ "msgType": "NOTIFICATION",
  "failures": [{"id": "3", "status": "DOWN", "failure-type": "SOFT"}],
  "sender": "OAM_HANDLER", "msgType": "NOTIFICATION"
}

2017-12-15 12:23:32,082 DeployManager:337 - DEPLOY response is:
{ "Recovery_Status": [1],
  "Recovery_Actions": [{"reconfigured_lightpath_id": 3,
    "action":{"action_type":"transmission_param_adaptation",
    "action_description": {"tsp_id": 80, "configuration_id": 81,
    "grid": 3,"n": 80, "m": 4, }, "expected_BER": 0.000366 }]}]
}

2017-12-15 12:23:32,568 OrchestraProvisioningManager:1437 - Connection
configuration message to be sent to node 3 is:
<filter xmlns="sss:sup:filter" xmlns:nc="urn:ietf:params:xml:ns:netconf:
base:1.0">
<connections nc:operation="replace">
<connection-id>3</connection-id>
<n>80</n>
<m>4</m>
<output-port-id>1</output-port-id>
<input-port-id>0</input-port-id>
</connections>
</filter>

```

Fig. 15. Top: Alarm (notification) for the soft failure. Middle: Soft failure restoration computation request to DEPLOY and DEPLOY response to increase the bandwidth of filters. Bottom: Configuration message to the filter.

(DEPLOY) estimated the BER for the available combinations of FEC and slot sizes (Fig. 14). DEPLOY then informed ABNO to maintain the same high FEC and increase the slot size to 50 GHz. This drastically reduced the filter penalty and successfully recovered the performance with a $2 \cdot 10^{-4}$ BER, within 0.5 dB of the assumed design margin, as seen in Fig. 14.

This experiment was performed under centralized control. The connection was initially established with an NBI call to the ABNO controller, which also installed the threshold at the Rx Level 0 agent. Then as the network aged, the Rx agent monitored the BER. Once the BER exceeded the threshold, an alarm was created at the Rx, propagated through the hierarchical monitoring plane, and reached the OAM Handler and the ABNO controller (Fig. 15, top). Then DEPLOY estimated the BER for the reconfiguration options and decided to keep the same FEC and increase the bandwidth of the filters (Fig. 15, middle). This information was sent from the ABNO provisioning manager to the ROADMs filters (Fig. 15, bottom).

C. Filter: Central Frequency Correction

In the previous experiment, we considered a gold connection that suffered from an increasing filter penalty due to the detuning of the Tx laser and the filters. The previous approach showed a reactive solution to this problem through reconfiguration. In the following we present an alternative solution, that of automated Tx-filter alignment. In this experiment we established a connection with a central frequency of 193.6 THz. The central frequency of the tunable filter was (de)tuned to +2 GHz to emulate the misalignment of the Tx/Rx and the path filter cascade. Figure 16 shows the power spectral density (PSD) before the central frequency alignment process, where we can clearly observe a distortion in the right edge of the signal.

Figure 17 shows the monitored parameters (SNR, BER, and signal bandwidth) during the central frequency alignment process. We followed a trial-and-error approach that involved several steps to optimize/align the Tx frequency with the filter. Initially, we searched around the starting frequency (by ± 1 GHz) to decide on the correct direction to move. Then we chose the direction (+) given the better performance of +1 GHz compared to -1 GHz (i.e., lower BER and higher SNR). Then at each state we monitored and obtained the average and standard deviation, and when we had enough

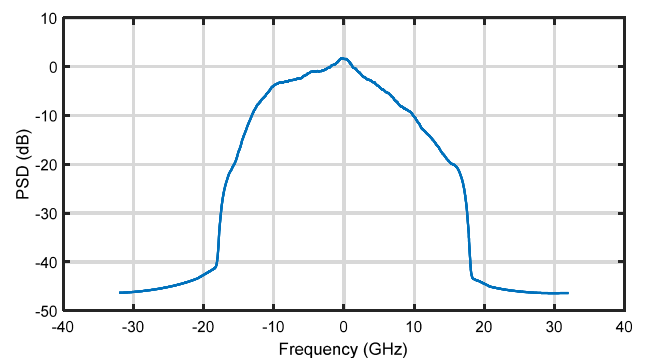


Fig. 16. Power spectral density before central frequency optimization.




Fig. 17. Monitored SNR, $\log(\text{BER})$, and signal bandwidth during the optimization procedure.

accuracy, we compared them to the values of the previous states to decide on the next move. Until we reached 193.602 THz, the performance improved. When we moved to 193.603 THz (step 4), the central frequency was shifted beyond the optimum and the performance was degraded (higher BER and lower SNR). Thus, we made a step back and finalized the alignment at a central frequency of 193.602 THz. In total our alignment process improved the SNR by 0.5 dB. Note that a 2 GHz misalignment of the Tx/Rx and filters is considered as a reference for the related effect. A sophisticated extension of this automatic filter alignment process can be found in [12].

The three field trial experiments discussed above are also presented in videos at <http://orchestraproject.eu>.

8. CONCLUSION

Optical networks' efficiency is directly connected to their transmission margins. We can increase the efficiency with no reliability deterioration by removing estimation uncertainties and by reducing long-term margins and adjusting them at intermediate periods. Further reduction of margins trades off higher efficiency for resiliency and should be done according to service classes. We presented the work carried out in the ORCHESTRA EU project, which developed a closed control loop that monitors, understands, predicts, and regulates the margins. ORCHESTRA developed the automation mechanisms to move toward low-margin, efficient, and reliable optical networks.

REFERENCES

- V. Lopez and L. Velasco, eds., *Elastic Optical Networks: Architectures, Technologies, and Control* (Springer, 2016).
- Y. Pointurier, "Design of low-margin optical networks," *J. Opt. Commun. Netw.* **9**, A9–A17 (2017).
- J. L. Auge, "Can we use flexible transponders to reduce margins?" in *Optical Fiber Communication Conference and Exhibition and the National Fiber Optic Engineers Conference* (IEEE, 2013).
- "Transforming Margin into Capacity with Liquid Spectrum," Ciena White Paper, https://media.ciena.com/documents/Transforming_Margin_into_Capacity_with_Liquid_Spectrum_WP.pdf.
- D. Ives, P. Bayel, and S. J. Savory, "Assessment of options for utilizing SNR margin to increase network data throughput," in *Optical Fiber Communication Conference* (Optical Society of America, 2015), paper M21.3.
- J. Pesic, T. Zami, P. Ramantanis, and S. Bigo, "Faster return of investment in WDM networks when elastic transponders dynamically fit ageing of link margins," in *Optical Fiber Communication Conference and Exhibition* (2016).
- P. Soumplis, K. Christodouloupoulos, M. Quagliotti, A. Pagano, and E. Varvarigos, "Network planning with actual margins," *J. Lightwave Technol.* **35**, 5105–5120 (2017).
- C. Delezoide, K. Christodouloupoulos, A. Kretsis, N. Argyris, G. Kanakis, A. Sgambelluri, N. Sambo, P. Giardina, G. Bernini, D. Roccatò, and A. Percelsi, "Marginless operation of optical networks," *J. Lightwave Technol.* **37**, 1698–1705 (2019).
- I. Sartzetakis, K. Christodouloupoulos, and E. Varvarigos, "Cross-layer adaptive elastic optical networks," *J. Opt. Commun. Netw.* **10**, A154–A164 (2018).
- A. P. Vela, M. Ruiz, F. Fresi, N. Sambo, F. Cugini, G. Meloni, L. Potì, L. Velasco, and P. Castoldi, "BER degradation detection and failure identification in elastic optical networks," *J. Lightwave Technol.* **35**, 4595–4604 (2017).
- Z. Dong, F. N. Khan, Q. Sui, K. Zhong, C. Lu, and A. P. Lau, "Optical performance monitoring: a review of current and future technologies," *J. Lightwave Technol.* **34**, 525–543 (2016).
- C. Delezoide, P. Layec, and S. Bigo, "Automated alignment between channel and filter cascade," in *Optical Fiber Communication Conference and Exhibition* (2019), paper Th2A.48.
- E. Sève, J. Pesic, C. Delezoide, S. Bigo, and Y. Pointurier, "Learning process for reducing uncertainties on network parameters and design margins," *J. Opt. Commun. Netw.* **10**, A298–A306 (2018).
- I. Sartzetakis, K. Christodouloupoulos, and E. Varvarigos, "Accurate quality of transmission estimation with machine learning," *J. Opt. Commun. Netw.* **11**, 140–150 (2019).
- P. Poggiolini, G. Bosco, A. Carena, V. Curri, Y. Jiang, and F. Forghieri, "A detailed analytical derivation of the GN model of non-linear interference in coherent optical transmission systems," arXiv:1209.0394 (2012).
- E. Torrenço, R. Cigliutti, G. Bosco, A. Carena, V. Curri, P. Poggiolini, A. Nespola, D. Zeolla, and F. Forghieri, "Experimental validation of analytical model for nonlinear propagation in uncompensated optical links," *Opt. Express* **19**, B790–B798 (2011).
- E. Sève, P. Ramantanis, J. C. Antona, E. Grellier, O. Rival, F. Vacondio, and S. Bigo, "Semi-analytical model for the performance estimation of 100 Gb/s PDM-QPSK optical transmission systems without inline dispersion compensation and mixed fiber types," in *European Conference and Exhibition on Optical Communication* (2013), paper Th.1D2.
- C. Delezoide, P. Ramantanis, and P. Layec, "On the performance prediction of optical transmission systems in presence of filtering," in *19th International Conference on Transparent Optical Networks* (2017).
- C. Delezoide, P. Ramantanis, and P. Layec, "Weighted filter penalty prediction for QoT estimation," in *Optical Fiber Communication Conference and Exposition* (2018).
- P. Serena, S. Musetti, S. Almonacil, S. Bigo, A. Benoni, P. Jenneve, N. Rossi, and P. Ramantanis, "The Gaussian noise model extended to polarization dependent loss and its application to outage probability estimation," in *European Conference on Optical Communication* (2018).
- E. Sève, J. Pesic, C. Delezoide, A. Giorgetti, A. Sgambelluri, N. Sambo, S. Bigo, and Y. Pointurier, "Automated fiber type identification in SDN-enabled optical networks," *J. Lightwave Technol.* **37**, 1724–1731 (2019).

22. M. Bouda, S. Oda, O. Vassilieva, M. Miyabe, S. Yoshida, T. Katagiri, Y. Aoki, T. Hoshida, and T. Ikeuchi, "Accurate prediction of quality of transmission based on a dynamically configurable optical impairment model," *J. Opt. Commun. Netw.* **10**, A102–A109 (2018).
23. M. Dallaglio, Q. P. Van, F. Boitier, C. Delezoide, D. Verchere, P. Layec, A. Dupas, N. Sambo, S. Bigo, and P. Castoldi, "Demonstration of SDN-based spectrum monitoring of elastic optical networks," in *Optical Fiber Communication Conference and Exhibition* (2017).
24. H. M. Chin, D. Charlton, A. Borowiec, M. Reimer, C. Laperle, M. O'Sullivan, and S. J. Savory, "Probabilistic design of optical transmission systems," *J. Lightwave Technol.* **35**, 931–940 (2016).
25. S. Yan, F. N. Khan, A. Mavromatis, Q. Fan, H. Frank, R. Nejabati, A. P. Lau, and D. Simeonidou, "Field trial of machine-learning-assisted and SDN-based optical network planning with network-scale monitoring database," in *Optical Fiber Communication Conference* (2017), paper Th.PDP.B.4.
26. G. Liu, K. Zhang, X. Chen, H. Lu, J. Guo, J. Yin, R. Proietti, Z. Zhu, and S. B. Yoo, "The first testbed demonstration of cognitive end-to-end optical service provisioning with hierarchical learning across multiple autonomous systems," in *Optical Fiber Communication Conference* (2018), paper Th4D.7.
27. R. Morais and J. Pedro, "Machine learning models for estimating quality of transmission in DWDM networks," *J. Opt. Commun. Netw.* **10**, D84–D99 (2018).
28. N. Sambo, D. Antonio, C. Porzi, V. Vercesi, M. Imran, F. Cugini, A. Bogoni, L. Potì, and P. Castoldi, "Sliceable transponder architecture including multiwavelength source," *J. Opt. Commun. Netw.* **6**, 590–600 (2014).
29. D. Rafique, T. Szyrkowicz, H. Grießer, A. Autenrieth, and J. P. Elbers, "Cognitive assurance architecture for optical network fault management," *J. Lightwave Technol.* **36**, 1443–1450 (2018).
30. N. Sambo, F. Cugini, A. Sgambelluri, and P. Castoldi, "Monitoring plane architecture and OAM handler," *J. Lightwave Technol.* **34**, 1939–1945 (2016).
31. R. Muñoz, N. Yoshikane, J. M. Fàbrega, R. Vilalta, L. Rodríguez, M. S. Moreolo, R. Casellas, R. Martínez, S. Beppu, D. Soma, and T. Tsuritani, "SDN control and monitoring system for soft-failure detection and optical restoration of spectral/spatial superchannels," in *European Conference on Optical Communication* (2018).
32. L. Velasco, A. Sgambelluri, R. Casellas, L. Gifre, J. L. Izquierdo-Zaragoza, F. Fresi, F. Paolucci, R. Martínez, and E. Riccardi, "Building autonomic optical whitebox-based networks," *J. Lightwave Technol.* **36**, 3097–3104 (2018).
33. F. Paolucci, A. Sgambelluri, F. Cugini, and P. Castoldi, "Network telemetry streaming services in SDN-based disaggregated optical networks," *J. Lightwave Technol.* **36**, 3142–3149 (2018).
34. S. S. Ahuja, S. Ramasubramanian, and M. Krunz, "Single-link failure detection in all-optical networks using monitoring cycles and paths," *IEEE/ACM Trans. Netw.* **17**, 1080–1093 (2009).
35. J. Tapolcai, P.-H. Ho, L. Ronyai, and B. Wu, "Network-wide local unambiguous failure localization (NWL-UFL) via monitoring trails," *IEEE/ACM Trans. Netw.* **20**, 1762–1773 (2012).
36. K. Christodoulopoulos, N. Sambo, and E. Varvarigos, "Exploiting network kriging for fault localization," in *Optical Fiber Communication Conference and Exhibition* (2016).
37. T. Panayiotou, S. P. Chatziz, and G. Ellinas, "Leveraging statistical machine learning to address failure localization in optical networks," *J. Opt. Commun. Netw.* **10**, 162–173 (2018).
38. N. Sambo, P. Giardina, I. Sartzetakis, A. Sgambelluri, F. Fresi, M. Dallaglio, G. Meloni, G. Bernini, K. Christodoulopoulos, P. Castoldi, and E. Varvarigos, "Experimental demonstration of network automation based on QoT estimation and monitoring in both single- and multi-domains," in *European Conference on Optical Communication* (2017).
39. S. Chisholm and H. Trevino, "NETCONF event notifications," IETF RFC 5277 (2008).
40. gRPC, "A high performance, open-source universal RPC framework," <https://grpc.io/>.
41. N. Sambo, A. Giorgetti, F. Cugini, and P. Castoldi, "Sliceable transponders: pre-programmed OAM, control, and management," *J. Lightwave Technol.* **36**, 1403–1410 (2018).
42. F. Boitier, V. Lemaire, J. Pesic, L. Chavarría, P. Layec, S. Bigo, and E. Dutisseuil, "Proactive fiber damage detection in real-time coherent receiver," in *European Conference on Optical Communication* (2017).
43. D. King and A. Farrel, "A PCE-based architecture for application-based network operations," IETF RFC 7491, 2015.