

Marginless Operation of Optical Networks

Camille Delezoide , Kostas Christodoulopoulos, Aristotelis Kretsis , Nikos Argyris , Giannis Kanakis , Andrea Sgambelluri , Nicola Sambo , Pietro Giardina, Giacomo Bernini, Diego Roccatto, Alessandro Percelsi, Roberto Morro , Hercules Avramopoulos, E. Varvarigos, Piero Castoldi , Patricia Layec , and Sebastien Bigo 

(Invited Paper)

Abstract—Considering flexible technologies available nowadays, operating optical networks much closer to their physical capacities is very tempting but necessarily requires efficient network automation. To achieve this, the two main challenges are handling failures, and accurately predicting performance in dynamic environments. We experimentally demonstrate the ability of the ORCHESTRA solution for early detection and localization of failures, to preventively mitigate their impact, and thus guarantee smooth network operation. Then, leveraging machine learning for live performance estimation and closed-loop software-defined network control, we demonstrate a fully automated reconfiguration of marginless connections undergoing critical performance variations over 228 km of field-deployed fiber.

Index Terms—Cross-layer optimization, failure localization, marginless operations, optical networks, soft-failures.

I. INTRODUCTION

A TYPICAL optical network is designed with high margins [1] to be resilient to any form of uncertainty related to its deployment and operation. Typical uncertainty sources are the statistical variability of network equipment performance, ageing, fluctuations from polarizations effects, increasing network load, fiber repairs, etc. Yet, high margins do not fully prevent outages, i.e., *hard failures*. These commonly occur due to fiber cuts [2] or critical network equipment breakdowns. To recover

traffic after hard failures, various restoration strategies exist, entailing network capacity overprovisioning from 25 to 100% [3]. With both high margins and restoration strategies, a lot of deployed physical capacity currently remains available.

Simultaneously, the industry must face steep traffic increase with constant revenue [4]. *Marginless operation* appears as a potent solution to overcome this challenge. It consists in meeting traffic demands at minimal cost through optimized use of resources. Marginless operation can take many forms, e.g., existing connections can be upgraded to meet increased demands or to postpone equipment purchase, optical bandwidth can be densified by closely matching slot-size and symbol rate, energy can be saved by downgrading forward error correction (FEC) scheme, etc.

With multi-rate transponders and flexgrid reconfigurable optical add-drop multiplexers (ROADMs), modern network elements already support marginless operation. At the lightpath level, marginless operation translates into a bit error ratio (BER) continuously brought close to FEC limit. As consequence, marginless operation would significantly increase hard failure probability with current practices in network maintenance. Marginless networks would thus require more overprovisioning for restoration, defeating the initial purpose of cutting costs. For optimal results, networks should first be able to detect and diagnose network health issues, i.e., *soft failures*, then to automatically prevent hard failures, either by directly acting at the cause [5], by downgrading capacity, by diverting traffic or by ordering manual intervention. In this last instance, failure classification and localization will reduce the mean-time-to-repair and thus operational expenditures.

In the ORCHESTRA solution [6], hard failures are prevented through the early detection, classification and localization of uncommon events and the automatic selection of the most appropriate recovery action, e.g., a change of modulation format, baudrate, FEC, slot-size, or a lightpath reroute. In marginless operation, the reversed actions are periodically envisaged for optimization purposes. For both applications, ORCHESTRA relies on its capacity to analyze all potential configurations to find an optimized one.

In this paper, we extend [7] and focus on failure recovery and network optimization at the lightpath level. Independently from what triggers the search for an optimized configuration, the same steps apply. First, the bit error ratio (BER) is estimated for each candidate configuration to filter out those with BER above FEC

Manuscript received October 11, 2018; revised November 9, 2018; accepted November 13, 2018. Date of publication November 16, 2018; date of current version April 2, 2019. This work was supported by the European Commission through H2020 ORCHESTRA under Grant 645360. (Corresponding author: Camille Delezoide.)

C. Delezoide, P. Layec, and S. Bigo are with Nokia Bell Labs, Nokia Paris-Saclay, Nozay 91620, France (e-mail: camille.delezoide@nokia-bell-labs.com; patricia.layec@nokia-bell-labs.com; sebastien.bigo@nokia-bell-labs.com).

K. Christodoulopoulos, A. Kretsis, and E. Varvarigos are with the Computer Technology Institute & Press, Athens 105 63, Greece (e-mail: kchristo@mail.ntua.gr; akretsis@ceid.upatras.gr; manos@ceid.upatras.gr).

N. Argyris, G. Kanakis, and H. Avramopoulos are with the National Technical University of Athens, Athens 106 82, Greece (e-mail: nikosa@mellanox.com; giankan@mail.ntua.gr; hav@mail.ntua.gr).

A. Sgambelluri, N. Sambo, and P. Castoldi are with the Sant'Anna School of Advanced Studies, Pisa 56127, Italy (e-mail: a.sgambelluri@santannapisa.it; n.sambo@sssup.it; castoldi@sssup.it).

P. Giardina and G. Bernini are with Nextworks, Pisa 56122, Italy (e-mail: p.giardina@nextworks.it; g.bernini@nextworks.it).

D. Roccatto, A. Percelsi, and R. Morro are with TIM, Turin 10122, Italy (e-mail: diego.roccato@telecomitalia.it; alessandro.percelsi@telecomitalia.it; roberto.morro@telecomitalia.it).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JLT.2018.2881840

limit. Then, an optimized configuration is found through the application of selection rules. Beyond technical aspects, these rules can be customized to fit network operator strategy, e.g., a lightpath reconfiguration is always preferred to a reroute, a lightpath should always operate at the maximal capacity achievable, a lightpath should always operate with minimal cost per bit, outside capacity upgrades lightpaths should operate with minimal power consumption, etc.

Regardless of the selection rules, a significant part of existing margins can only be converted into resource savings if the BER estimations made to filter out configurations are more accurate than those made at initial network design [1]. In this context, neural networks were recently proposed to predict BER-related metrics for spectral efficiency adaptation [8] and impairment aware service provisioning [9]. At current state-of-the-art however, the prediction accuracy achieved with neural networks remains below model-based BER-prediction tools routinely used in optical network design [10]. It also remains unclear how such high-complexity machine learning techniques could be pragmatically trained to consistently achieve accurate BER predictions in full-size deployed optical networks. In the ORCHESTRA solution, we propose the first BER estimation tool that leverages receiver-based monitoring and low-complexity machine learning to improve the accuracy of state-of-the-art model-based BER prediction through a refined and regularly updated knowledge of model inputs.

In this paper, we validate through simulations, and both laboratory and field trial experiments the capacity of the ORCHESTRA solution to meet the fundamental requirements for marginless operation of optical networks. We first experimentally demonstrate the early detection and localization of soft failures before they can lead to hard failures [11]. Then, after validating the accuracy of our new BER prediction tool through simulations, we demonstrate the first fully automated operation of marginless connections undergoing major performance variations over 228km of field-deployed fiber [7]. In a first use case, we present the capacity adaptation of a best-effort connection with emulated amplifier and fiber ageing. Finally, we present the net capacity maintenance of a gold class connection facing BER variations from emulated laser and wavelength selective switch (WSS) ageing.

II. THE ORCHESTRA ARCHITECTURE

We use ORCHESTRA's hierarchical and programmable management infrastructure [12] illustrated in Fig. 1. The two key building blocks are the Operations, Administration and Maintenance (OAM) Handler, and the Application-Based Network Operations (ABNO) Controller. The ABNO controller implements workflows for several ORCHESTRA use cases and uses the DEPLOY [13] software module as a Path Computing Element (PCE) extended with Quality of Transmission (QoT) estimation ability. DEPLOY performs monitoring-based BER predictions and determines appropriate recovery actions. Finally, the ABNO controller (re)configures lightpaths - Transmitters (Tx), Receivers (Rx) and ROADMs nodes - through a provisioning manager and Software Defined Network (SDN) controllers.

The control logic is applied via NETCONF following YANG model and REST. This versatile control plane architecture can notably be programmed to detect, identify, localize and react to various types of soft and hard failures. At lowest control and monitoring level, agents placed in receivers (Rx) can check any monitored parameter against one or more thresholds, and generate parameter-specific alarms when corresponding thresholds are exceeded. An alarm can be handled at all levels of the hierarchical architecture. Additionally, the way an alarm is handled by each agent is fully programmable: the agent can either apply correlation or suppression functions, try to solve the problem locally, or forward the alarm to a higher level.

III. DETECTION AND LOCALIZATION OF SOFT-FAILURES

A. Detection Method

Setting safety thresholds close to the FEC limit is appropriate to handle slow variations of performance, e.g., as expected from equipment ageing. In contrast, a sudden and relatively large BER increase - compared to natural variations - could be due to an unexpected event such as a fiber degradation or equipment malfunction. Such BER variations should be distinguished from natural variations and treated as threats and systematically investigated. Accordingly, the soft-failure detection method we implemented relies on a dynamic threshold Th_{SF} individually set for each lightpath according to the current BER average $\langle BER \rangle$ and standard deviation σ , as $Th_{SF} = \langle BER \rangle + k\sigma$ where $\langle BER \rangle$ and σ are periodically evaluated from monitored BER values. When a monitored BER value is above Th_{SF} , a soft-failure alarm is raised.

By simulation, we test the impact of parameter k assuming gaussian-distributed BER variations outside soft-failures (Fig. 2(a)–(b)). For simplicity, we also assume that soft-failures do not modify the BER standard deviation (Fig. 2(b)). Since the pre and post-failure BER probability distributions overlap, there is a non-zero probability that similar BER values are measured before and after the soft-failure because of natural BER fluctuations, as can be observed in Fig. 2(a). In such realistic conditions, we understand that a simpler soft-failure detection method, e.g., based on a fixed threshold, would generate many false alarms. Although eradicating false alarms is statistically impossible, the method we propose allow to control the false alarm probability through the parameter k . We also define the soft-failure *Detection Limit* (DL) as the smallest relative variation of BER identified as a soft-failure. As can be expected, we observe in Fig. 2(c)–(d) that k trades-off soft-failure detectability with false alarm probability. Experimentally, we choose $k = 4$ to achieve a false alarm probability below 0.01%.

B. Localization Algorithm

Once a soft-failure is detected at the lightpath level, the monitoring agent in the corresponding Rx (Fig. 1) generates an alarm that reaches the OAM handler. The OAM handler introduces a *correlation delay* to capture similar alarms from other lightpaths. This delay is set according to the BER polling rate. Once the correlation delay is over, the OAM Handler forwards all

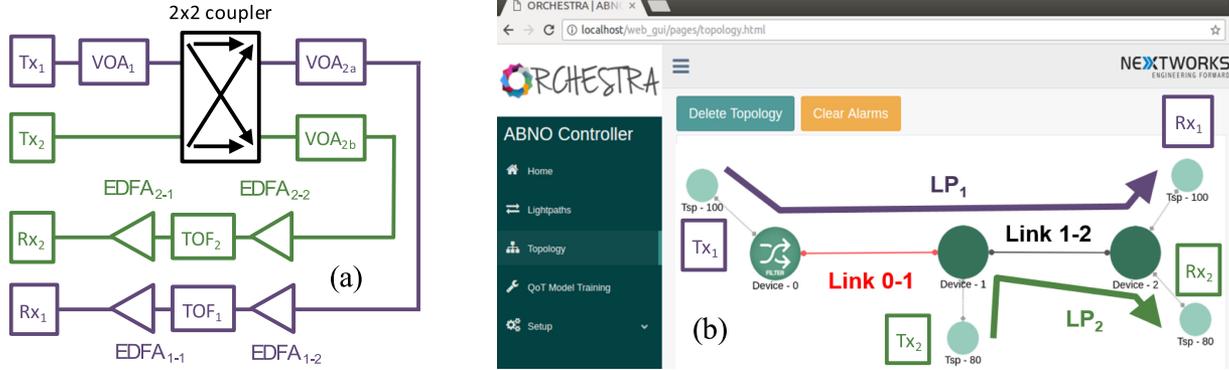


Fig. 3. Data plane and control plane view (ABNO user interface) of the soft-failure detection experiments.

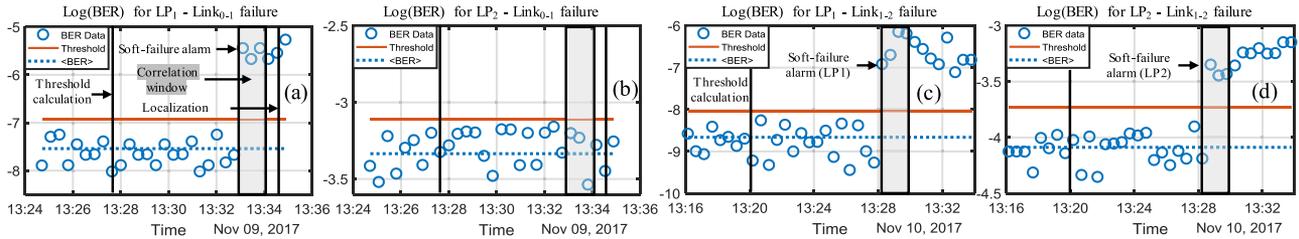


Fig. 4. Monitored BER in the soft-failure detection and localization experiments.

issue is mitigated by minimizing the correlation window through faster BER polling rates.

IV. PRE-FEC BER ESTIMATION FRAMEWORK

To determine the feasibility of a lightpath configuration, we estimate its pre-FEC BER and compare it to the relevant FEC limit. This first relies on the estimation of the generalized optical signal to noise ratio (gOSNR) [14], related to BER by:

$$BER = f \left[\sum_{k=0}^2 a_k (R/B)^k gOSNR^{-k} \right] \quad (1)$$

where $a_k = 0, 1, 2$ are constants depending on transponder and filter cascade, R is the baudrate, B the 0.1nm reference bandwidth and f is mathematically derived and depends on modulation format and coding scheme [15]. The gOSNR extends OSNR definition by including, in addition to amplifier noise, all other optical noises e.g., from Kerr effect (nonlinearities) and linear crosstalk.

For unestablished lightpaths, the current best practice is to calculate gOSNR as a known function of system parameters e.g., amplifier noise factor and fiber nonlinearity, and use a_k values obtained through back-to-back transmissions [14]–[16] to finally convert gOSNR into pre-FEC BER using Eq. 1. In this estimation scheme, imperfect knowledge of system parameters values leads to prediction errors on gOSNR and thus on BER.

For marginless operations, the same method could be applied to determine the feasibility of a reconfiguration for active lightpaths, with the same limitations in terms of accuracy. Instead, the estimation framework we propose leverages monitoring of active lightpaths and low complexity machine learning to improve BER estimation accuracy for reconfigurations.

A. Monitoring and Machine-Learning Based BER Estimation

The goal of the method we propose is to estimate, for a given active lightpath, the impact on BER of a configuration change, i.e., a change of baudrate, modulation format, FEC or slot-size.

Before deployment, BER and gOSNR are simultaneously measured in back-to-back for all transponder configurations. This constitutes the training set, used to train a stochastic gradient descent polynomial regression model and obtain the a_k constants for the different transponder configurations offline. During operation, we feed the monitored BER and use the trained parameters of the current configuration to invert Eq. 1 and obtain the gOSNR value. Then, we feed obtained gOSNR and trained parameters values for the target configuration into Eq. 1 to estimate BER in case of reconfiguration.

With such scheme, the BER monitored in any current configuration can be used to estimate the BER for any other configuration. Compared to the typical estimation method used for unestablished lightpaths, the BER estimation described here relies on a direct measurement of gOSNR rather than a calculation based on assumed input parameters values. The accuracy of our method is thus independent to model input uncertainties responsible for design margins.

B. Validation Through Numerical Simulations

We simulate an optical connection at 32 Gbd and RRC0.1 pulse shape crossing a cascade of six 37.5 GHz-wide filters, equivalent to three route-and-select ROADMs nodes. We test the monitoring and machine-learning based estimation method by calculating the BER estimation errors going from QPSK to 16QAM, and reversely from 16QAM to QPSK. In the simulations, we first perform, for both modulation formats, best-case

and worst-case calibrations as described in [16] by simultaneously measuring both BER and gOSNR for various quantities of noise, loaded after Tx (best-case) or before Rx (worst-case). Then, we simulate the optical connection for both modulation formats and measure the BER for various loaded noise levels. Note that these simulations differ from back-to-back calibrations essentially because the loaded noise is evenly distributed between the six cascaded filters.

We test the impact of the number N of calibration points used for interpolation on the accuracy of the BER prediction. In Fig. 5(a-c), we plot measured and estimated BER for 16QAM (Fig. 5(a)) and QPSK (Fig. 5(c)), and the corresponding relative BER prediction errors in Fig. 5(b)-(d). We first observe that the number N of calibration points has very little impact on the BER prediction accuracy. In Fig. 5(a)-(c), estimated and measured BER are a close match for all gOSNR values with only 3 calibration points. From Fig. 5(b)-(d), we further observe that relative prediction error increases when the BER values decreases. This is because a given absolute error on BER prediction leads to larger relative errors for small BER values. Therefore, the relative errors increase with the gOSNR value, and errors are generally larger for QPSK than for 16QAM. However, since large relative errors are obtained for BER values far from FEC limit, they have no impact on lightpath operation. Near FEC limit, for BER = 10^{-2} , BER prediction error is within 5% for 16QAM, and within 10% for QPSK.

C. Method Extension to Account for Actual Filter Cascade

In the previous sections, we assumed for simplicity that the filter cascade used to establish the training set is the same as the field deployed one. For the current standard transmissions with 32 GBd baudrate and 50 GHz allocated slot-size, this assumption has no significant impact on BER estimation accuracy due to the negligible impact of ROADMs filtering on a_k parameters. However, for more bandwidth-efficient transmissions e.g., 32 GBd in 37.5 GHz slot-size, using the method described in Section IV-A in a full-size network would require individual training sets for all possible filter cascades. Moreover, residual BER estimation errors would stem from the unavoidable discrepancy between expected and actual filter cascade.

In the ORCHESTRA solution, we thus extended the BER estimation method described in Section IV-A to adjust to the actual, field-deployed filter cascade through the monitored bandwidth (MBW). The MBW is defined as the width at -10 dB of the signal spectrum measured at the output of the analog-to-digital converter (ADC) inside the coherent receiver. Following experimental observations [15], we assume that the impact of the filter cascade can be treated as a constant penalty on the signal to noise ratio SNR. The SNR is related to gOSNR by:

$$SNR^{-1} = \sum_{k=0}^2 a_k (R/B)^k gOSNR^{-k} = f^{-1}(BER) \quad (2)$$

Before deployment, MBW and filter penalty D_{SNR} are simultaneously measured in a back-to-back transmission with a tunable optical filter inserted before noise loading, for various bandwidth settings. This is repeated for every transponder con-

figuration. As previously, this data is used as training set for the BER estimation. In operation, the MBW provided by the coherent receiver is converted into a filter penalty D_{SNR} based on relevant training data. Then, D_{SNR}^* , the filter penalty for the candidate configuration is estimated. If the change of configuration does not entail a variation of MBW , e.g., a change of modulation format, then the MBW value is directly converted into the proper D_{SNR}^* value based on training data relevant to the candidate configuration. When the change of configuration modifies the MBW , we first need to estimate the monitored bandwidth MBW^* in the new configuration. To achieve this, we rely on the reference bandwidth (RBW) as a model-based estimation of the monitored bandwidth. This estimation is made through numerical evaluation of the filter function equivalent to the filter cascade, based on general knowledge e.g., baudrate, pulse shape, slot-size, number of crossed ROADMs, and filter characteristics. Note that we expect MBW to differ from RBW due to device variability and ageing, and we assume that the ratio between MBW and RBW is independent of the configuration, so that the estimation of the new monitored bandwidth is $MBW^* = MBW \cdot RBW^* / RBW$. Then, MBW^* is converted into D_{SNR}^* using the training data. Finally, the new bit error ratio $BER^* = f(SNR^{*-1})$ is calculated using:

$$SNR^* = SNR + D_{SNR} - D_{SNR}^* + \Delta_R + \varepsilon \quad (3)$$

where $SNR^{-1} = f^{-1}(BER)$, Δ_R is the calibrated SNR penalty due to configuration change (outside filter effects), and $\varepsilon = 0.5$ dB an added margin to account for model error. We illustrate the application of this model extension in Section V-C.

V. FIELD TRIAL OF MARGINLESS NETWORK OPERATION

A. Field Trial Setup

The fiber cable is deployed between Turin (Stampalia) and Chivasso network exchanges and is composed of eight 76 km-long G.652 fiber spans. We use 5 spans to create two lightpaths of 3 and 2 spans, denoted as LP1 (228 km) and LP2 (152 km) respectively (cf. Fig. 6). Two lab-hosted reconfigurable optical add-drop multiplexers (ROADMs) switch traffic from LP1 to LP2. The transmitter (TX) supports QPSK, 8QAM, 16QAM modulation formats, and 28 and 32 GBd baudrates to achieve net capacities of 100, 150 and 200 Gb/s with 12% (low-FEC) or 28% (high-FEC) coding rates. The receiver (RX) integrates a 40Gsamples/s oscilloscope, offline signal processing to handle both signal demodulation and monitoring of multiple transmission parameters, and a software agent to report monitored values and raise alarms. The various ageing scenarios are emulated through a combination of tunable optical filter (TOF) and variable optical attenuator (VOA).

B. Best Effort Rate Adaptation Under Fiber/Amplifier Ageing

Here, we consider a best-effort connection that leverages end-of-life (EOL) margins [1] to first increase its capacity and operate as close as possible to FEC limit during network lifetime. Fiber/amplifiers ageing is emulated by increasing the optical attenuation using a VOA (Fig. 6). As consequence, the OSNR

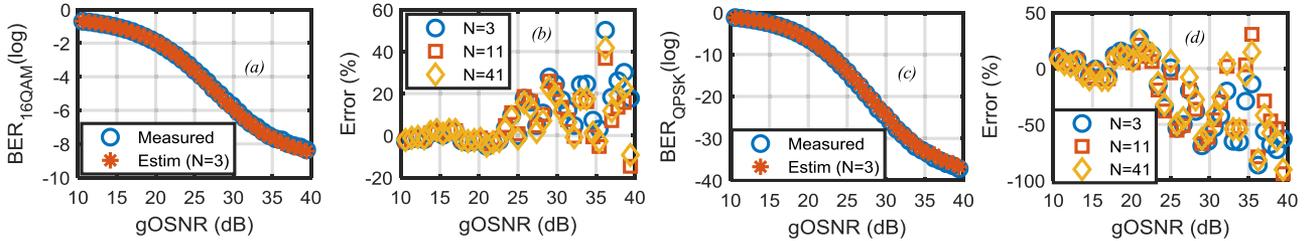
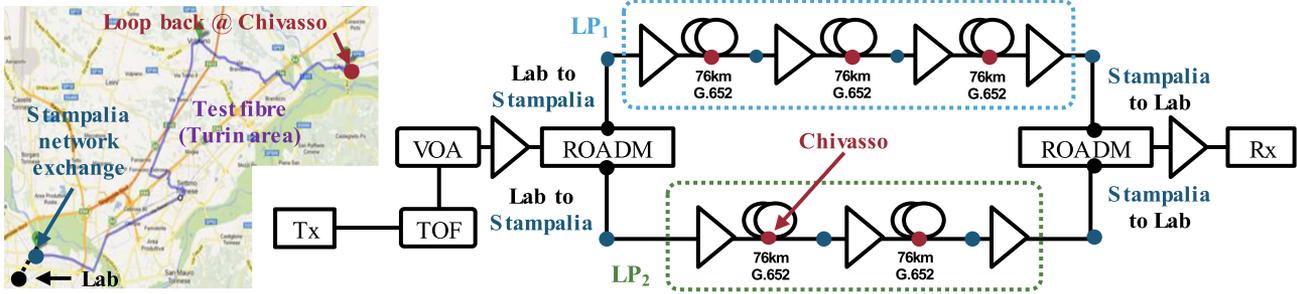

 Fig. 5. Impact of the number N of calibration points on the BER prediction accuracy.


Fig. 6. Field trial setup at TIM premises in Turin region, Italy.

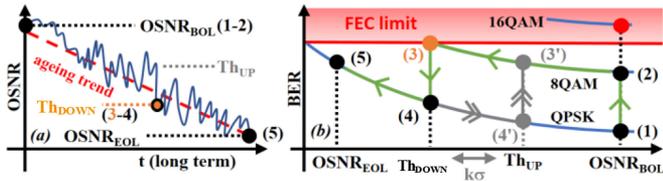


Fig. 7. OSNR variations and format adaptations.

evolves as illustrated in Fig. 7(a), where on top of the ageing trend fluctuations naturally occur in the field. The system reacts to OSNR variations as illustrated in Fig. 7(b).

Prior to the rate adaptation experiment, we measure seven points of BER and OSNR in back-to-back for each modulation format (cf. Fig. 8(a)). This is to train DEPLOY to convert BER into gOSNR and reversely, as described in Section IV-B. Thanks to these BER–gOSNR conversions, the performance thresholds related to alarms can be defined for either metric. At begin-of-life, the lightpath operates at 100G–QPSK. During a short period, i.e., a few minutes during the trial emulating days or weeks of operation, DEPLOY learns (cf. Fig. 8(b)) the BER average and standard deviation σ in state (1). Based on the monitored BER value at $2.5 \cdot 10^{-5}$ and training data (Fig. 8(a)), DEPLOY estimates the BER to be 10^{-2} in 150G–8QAM and $3.5 \cdot 10^{-2}$ in 200G–16QAM. With a FEC threshold at $2.10 \cdot 10^{-2}$, the lightpath is naturally reconfigured to 150G–8QAM (2) and reaches a BER of $1.3 \cdot 10^{-2}$ (cf. Fig. 8(c)), within the 0.5 dB of prediction error margin. Then, the BER progressively increases due to emulated ageing. Rate downgrade is based on an gOSNR threshold Th_{DOWN} (Th_{SF} in Fig. 7(a)–(b)) defined according to FEC scheme. Once the gOSNR crosses Th_{DOWN} in state (3), the various possible reconfiguration options are explored. Con-

sequently, DEPLOY requests to switch back to 100G–QPSK (Fig. 8(c)–(d)) corresponding to state (4).

At this point, DEPLOY could still impose a rate upgrade if judged possible. Because of gOSNR fluctuations, this reversibility could lead to a ping-pong effect between modulation formats. To mitigate this, we implemented a history-based hysteresis (Fig. 7(b)) so that once a downgrade occurred, DEPLOY only allows capacity upgrade when gOSNR exceeds $Th_{UP} = Th_{DOWN} + k\sigma$, where σ is the learned standard deviation of the OSNR and $k = 4$. In this context, the parameter k balances modulation format stability with minimal margin operation.

From state (4), the history-based hysteresis prevents the system from reverting back to state (3) despite natural BER fluctuations. The connection remains in QPSK until network’s end of life.

C. Gold-Class Rate Maintenance Under ROADM Ageing

In this second use case, we consider a gold-class 100G–QPSK connection automatically maintaining its net capacity through FEC (hence baudrate) and/or slot-size adaptation. For marginless operation, we aim to convert EOL margins into a 33% gain in spectral efficiency by allocating the connection 37.5 GHz of bandwidth instead of the typical 50 GHz. Thus, we consume EOL margins through the increase of filter penalty. As side effect, the bandwidth reduction makes the connection more vulnerable to filter ageing. This effect typically stems from the progressive detuning of crossed ROADMs and transmitter laser frequency over the years, further increasing the filter penalty. In this context, we exploit DEPLOY’s capacity to accurately predict BER based on monitored bandwidth (MBW) inside the coherent receiver, as described in Section IV-C.

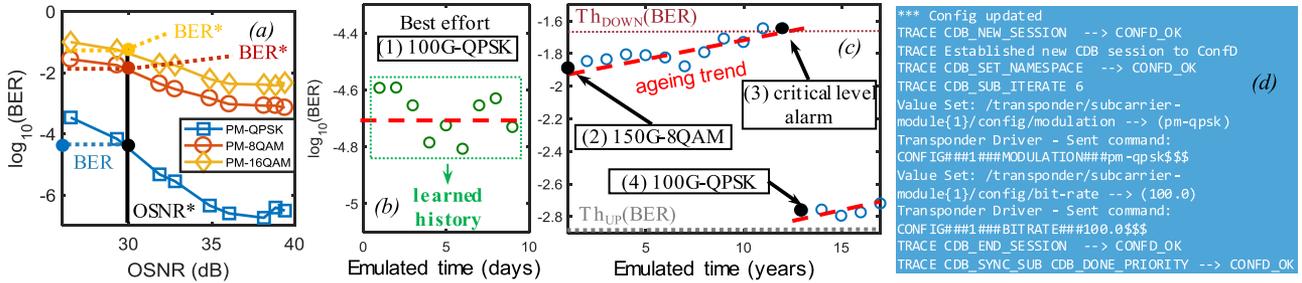


Fig. 8. Best-effort automatic rate adaptation in presence of emulated fiber/amplifier ageing.

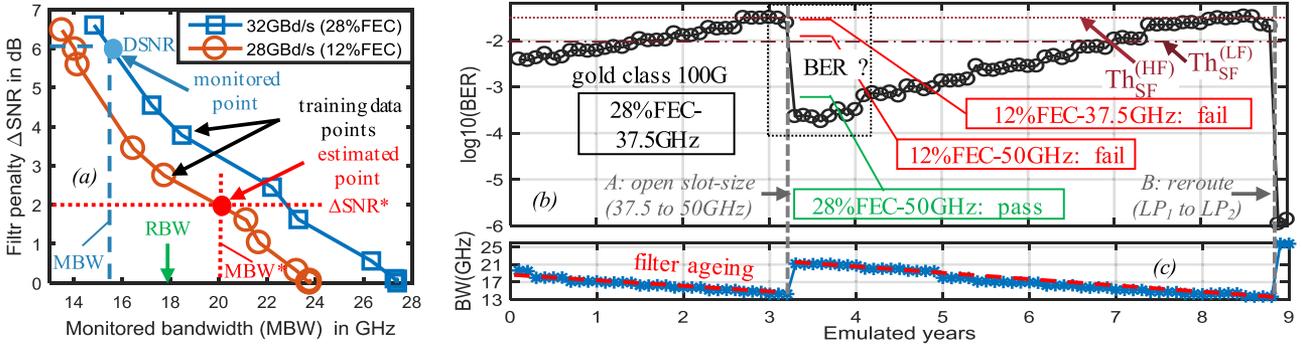


Fig. 9. Gold class automatic FEC and slot-size adaptation for net capacity maintenance under filter ageing.

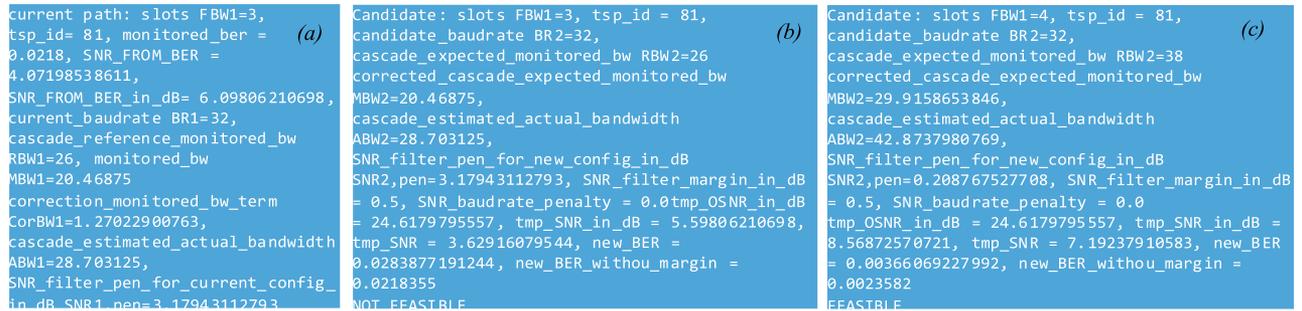


Fig. 10. Control plane logs corresponding to the exploration of various candidate solutions for lightpath recovery.

We emulate ten cascaded filters and filter ageing through initial setting and progressive reduction of bandwidth of a tunable optical filter (TOF in Fig. 6). Prior to the net capacity maintenance experiment, we measure ten points of filter penalty D_{SNR} and MBW in back-to-back configuration for both baudrates, 28 GBd and 32 GBd. Experimental data is plotted in Fig. 9(a) along with an illustration of the various conversions occurring to get from an MBW value to a BER prediction through Eq. 3 in an arbitrary case.

At begin-of-life, the gold class 100G-QPSK connection is first established with high-FEC (HF = 28%) and 37.5 GHz slot-size (Fig. 9(b)). When, because of filter ageing (Fig. 9(c)), the HF BER threshold is reached, DEPLOY computes BER estimates for current and available combinations of FEC and slot-size (Fig. 10). DEPLOY then informs ABNO to maintain high-FEC and increase slot-size to 50 GHz. As can be seen in Fig. 9(b),

this drastically reduces filter penalty and successfully recovers the LP with a 2.10^{-4} BER, within 0.5 dB of the estimation margin. After this first adaptation, the emulated filter cascade keeps ageing and BER reaches the FEC threshold once again. This time, the only option left to DEPLOY is to order a reroute to a new, shorter lightpath (LP2) with reduced filter penalty.

VI. CONCLUSION

We validated through simulations, and both laboratory and field trial experiments the capacity of the ORCHESTRA solution to meet the fundamental requirements for marginless operation of optical networks. We first demonstrated its ability to detect and localize soft failures as unexpected BER variations, requiring further analysis and potentially reconfiguration actions. We then introduced a novel BER estimation

method leveraging analytical models, machine learning and monitoring to achieve accurate BER predictions in the search for an optimized lightpath configuration, and demonstrated its accuracy through numerical simulations. Finally, we integrated this novel BER estimation with the ORCHESTRA solution to demonstrate in a field trial a full cross-layer network solution able to steadily operate with minimal margins despite ageing and short-term performance fluctuations. Under fiber/amplifier ageing, the SDN-based optical testbed dynamically adapted modulation format to upgrade and downgrade capacity. Under filter ageing, we leveraged our new filter-aware BER estimation for FEC and slot-size adaptation.

ACKNOWLEDGMENT

The authors would like to thank Philippe Jennev  and Fabien Boitier for lab support.

REFERENCES

- [1] Y. Pointurier, "Design of low-margin optical networks", *IEEE/OSA J. Opt. Commun. Netw.*, vol. 9, no. 1, pp. A9–A17, Jan. 2017.
- [2] M. Hoffman, [Online]. Available: <http://all.net/CID/Attack/papers/CableCuts.html>
- [3] R. Eisenach, "Advances in optical layer restoration," Nokia White Paper, 2016.
- [4] M. Weldon, *The Future X Network: A Bell Labs Perspective*. Boca Raton, FL, USA: CRC Press, 2016.
- [5] M. Dallaglio *et al.*, "Demonstration of SDN-based spectrum monitoring of elastic optical networks," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2017, Paper Tu3L.5.
- [6] [Online]. Available: <http://www.orchestraproject.eu>
- [7] C. Delezoide *et al.*, "Field trial of marginless operations of an optical network facing ageing and performance fluctuations," in *Proc. Eur. Conf. Opt. Commun.*, 2018, Paper Tu1D.3.
- [8] S. Yan *et al.*, "Field trial of machine-learning-assisted and SDN-based optical network planning with network-scale monitoring database," in *Proc. Euro. Conf. Opt. Commun.*, 2017, Paper Th.PDP.B.4.
- [9] G. Liu *et al.*, "The first testbed demonstration of cognitive end-to-end optical service provisioning with hierarchical learning across multiple autonomous systems," in *Proc. Opt. Fiber Commun. Conf. Expo.*, 2018, Paper Th4D.7.
- [10] P. Poggiolini *et al.*, "The GN-model of fiber non-linear propagation and its applications," *J. Lightw. Technol.*, vol. 32, no. 4, pp. 694–721, Feb. 2014.
- [11] C. Delezoide *et al.*, "Pre-emptive detection and localization of failures towards marginless operations of optical networks," in *Proc. 20th Int. Conf. Transparent Opt. Netw.*, 2018, Paper We.D2.3.
- [12] N. Sambo *et al.*, "Experimental demonstration of fully disaggregated white bow including different types of transponders and monitors, controlled by NETCONF and YANG," in *Proc. Opt. Fiber Commun. Conf. Expo.*, 2018, Paper M4A.3.
- [13] K. Christodoulopoulos *et al.*, "Observe-decide-act: experimental demonstration of self-healing network," in *Proc. Opt. Fiber Commun. Conf. Expo.*, 2018, Paper M3A.7.
- [14] E. Torrenge *et al.*, "Experimental validation of analytical model for non-linear propagation in uncompensated optical links," *Opt. Express*, vol. 19, no. 26, pp. B790–B798, 2011.
- [15] C. Delezoide *et al.*, "On the performance prediction of optical transmission systems in presence of filtering," in *Proc. 19th Int. Conf. Transparent Opt. Netw.*, 2017, Paper WE.B1.1.
- [16] C. Delezoide *et al.*, "Weighted filter penalty prediction for QoT estimation," in *Proc. Opt. Fiber Commun. Conf. Expo.*, 2018, Paper W2A.56.

Authors' biographies not available at the time of publication.