Experimental Validation of Lightpath Operation based on Continuous Digital-Twin Model Tuning

P. González[(1)], A. Moawad[(2)], S. Ghasrizadeh[(1)], R. Casellas[(3)], M. Ruiz[(1)], H. Zaid[(2)], and L. Velasco[(1),*]

[(1)] Optical Communications Group (GCO), Universitat Politecnica de Catalunya (UPC), Spain [*]luis.velasco@upc.edu
[(2)] Fraunhofer Heinrich-Hertz-Institute (HHI), Germany - [(3)] Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Spain

Abstract: A complete digital twin-assisted lightpath operation workflow for multiband optical networks is proposed for assured performance and reliable surveillance. Accurate performance estimation and prompt degradation detection through continuous model tuning using telemetry are experimentally shown. © 2025 The Author(s)

1. Introduction

Optical Network Digital Twins (NDT) have become essential control plane systems enabling intelligent and predictive network operation. Such tools provide Quality of Transmission (QoT) estimation during lightpath provisioning and facilitate the detection of performance degradation coming from soft-failures, while distinguishing from degradation coming from network operation [1]. Note that in the case of Multiband (MB) optical transmission, lightpaths are affected by the inter-channel stimulated Raman scattering non-linear effects, and thus the QoT of established lightpaths changes when other lightpaths are set up or torn down [2].

In our recent works in [2],[3], we extended the OCATA optical time domain NDT for MB optical transmission and demonstrated its accuracy for QoT estimation, specifically the pre-Forward Error Correction (FEC) Bit Error Rate (BER), which enables its application during the lightpath provisioning process. OCATA lightpath models are based on the concatenation of span and node models for the elements along the route of the lightpaths. In practice, however, it might happen that real spans in the network do not match the characteristics of the pre-trained models, e.g., the length. In such cases, lightpath models would provide QoT estimation that might differ from the real ones after the lightpath is established on the network. For that very reason, we proposed to fine tune the initial lightpath models with telemetry data [4], which are available only after the lightpaths are set up.

In this paper, we propose and experimentally validate a complete workflow for lightpath operation that includes lightpath provisioning using an NDT, as well as the continuous model tuning procedure based on telemetry data collected from optical devices to improve degradation detection.

2. Intelligent Control Plane Architecture

Fig. 1a presents the architecture of the control plane for MB optical network operation. A Software-Defined Networking (SDN) controller is in charge of network programmability, including the configuration of the network devices for lightpath operation (i.e., provisioning and reconfiguration). Node agents are in charge of actually configuring the optical devices following the detailed directions received from the SDN controller. Node agents play also a key role for telemetry collection, as they gather detailed measurements from the optical devices with very small periodicity and process them to reduce their dimensionality [5]. Processed telemetry is stored in a centralized telemetry database (DB) and made available for other systems in the control plane. A message bus in publish-subscribe mode is used to convey telemetry data from the node agents to the centralized telemetry DB. A centralized NDT (in this paper we assume OCATA) is used to assist the SDN controller. Focusing on lightpath operation, the OCATA NDT provides accurate QoT estimation used during lightpath provisioning to help on the selection of the route and channel across bands, as well as for QoT degradation detection based on telemetry measurements after the lightpath is set up. Finally, an inventory stores the characteristics of optical devices.

3. Distributed Telemetry Processing and NDT-assisted Operation

The OCATA NDT predicts optical signal propagation in the time domain. Models are trained for individual optical components, e.g., optical spans and switches, and for specific characteristics, e.g., the length of the optical span. When the propagation through an end-to-end lightpath needs to be modelled, component models for the specific characteristics of those in the route of the lightpath and for the selected channel, are concatenated. Once the lightpath model is obtained, it can be used to estimate the pre-FEC BER, which is paramount, not only during lightpath provisioning, but also for performance surveillance.

Because models assume fully loaded systems, QoT estimation is for the worst case, which guarantees lightpath performance under any load scenario. However, using such models for QoT surveillance would result in delayed degradation detection since the margin of QoT variation between telemetry and estimation would be larger in the common case of not fully loaded systems. Therefore, lightpaths' models need to be tuned to the actually observed values, so degradation detection algorithms can analyze QoT variations, e.g., whenever they exceed some tight threshold.

Fig. 1b presents the complete operation workflow, which includes: *i*) lightpath set up; *ii*) telemetry collection; *iii*)

The research leading to these results has received funding from the European Union's Horizon Europe research and innovation programme SEASON (G.A. 101096120), project PID2024-157824OB-I00 funded by MICIU /AEI /10.13039/501100011033 / FEDER, UE, and from the ICREA Institution.

This Article © 2026 The Author(s)

model tuning; and *iv*) performance surveillance. Note that model tuning might occur also after performance surveillance.

During lightpath set up, the SDN controller follows a routing and spectral allocation procedure assisted by the OCATA NDT (see [3] for details) and it requests QoT estimation of a selected route and channel (message 1 in Fig. 1b). OCATA then queries the Inventory system to obtain detailed descriptions of the optical components specified in the route (2). Based on the retrieved characteristics, the most appropriate pretrained component models are selected from its internal Model DB. The end-to-end lightpath model is created together with associated metadata, e.g., the id of the lightpath, and stored in the Model DB. QoT is then estimated by propagating a pseudorandom sequence of IQ symbols (3). OCATA then replies the SDN controller with the estimated pre-FEC BER (4). If the estimated performance is acceptable, the lightpath is established in the network (not shown in the figure) and a notification is sent to OCATA (5). Such notification triggers the configuration of the telemetry collector in the node agent, including periodicity T_1 (e.g., few sec) and the modulation order of the signal (m) for the Feature Extraction (FeX) algorithm, as well as timers for model tuning (T_2) (e.g., hours) and surveillance (T_3) (e.g., few min).

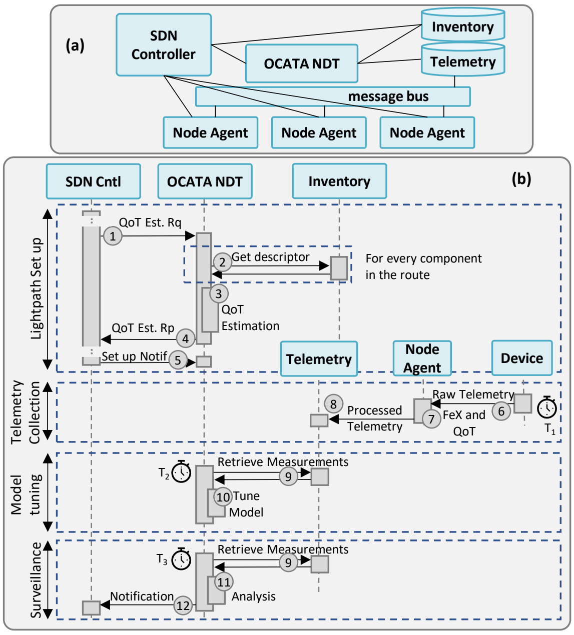

Fig. 1 Control plane (a) and operation workflow (b)

The node agent starts collecting IQ constellation samples just after the lightpath is set up (6), which are locally processed. Specifically, the FeX procedure uses Gaussian Mixture Model (GMM) fitting, so each constellation point (CP) i of a m-QAM modulated optical signal is modelled as a bi-variate Gaussian distribution defined by set of features $Y^i=[\mu^I,\mu^Q,\sigma^I,\sigma^Q,\sigma^{IQ}]$, representing the mean, variance and covariance of such distribution, respectively. Once features $Y=[Y^i]_{i:1..m}$ are obtained, the pre-FEC BER is calculated by computing the probability (Φ^i_{out}) that a symbol originally belonging to CP i is received outside of the detection area defined for that CP (see [3] for details). Features Y and the pre-FEC BER estimated from the collected IQ samples are published in the message bus to be tagged and stored in the telemetry DB (8).

With periodicity T_2, the lightpath model is checked and tuned, if needed, using the latest measurements. To that end, pre-FEC BER measurements collected during the last period are retrieved from the telemetry DB (9) and a subset of them are randomly selected. The average pre-FEC BER that the lightpath has experienced during the period is compared to the one estimated by the model to check whether the relative error exceeds an acceptable limit. In that case, the concatenated model is sent to a Sandbox to be re-trained with the features Y from telemetry and the tuned, adjusted model is replaced in the model DB (10).

Finally, the surveillance process is executed with periodicity T_3. As for model tuning, OCATA retrieves the most recent pre-FEC BER measurements from the telemetry DB (9) and runs an algorithm that analyzes its evolution with time, compares that to the value predicted by the lightpath model and detects any degradation of the performance of the lightpath (11). If a degradation is detected, a notification is sent to the SDN controller (12), so it can initiate any subsequent action.

4. Testbed Details and Experimental Results

The experiments were carried out on a distributed testbed connecting premises in HHI, UPC and CTTC, where control and data plane components were deployed. The data plane (Fig. 2a) consists of a 16-QAM MB transmitter (Tx) and a receiver (Rx) connected through an Optical Line System (OLS) [6]. The Tx includes a 4-ch 64 GBd DAC and a commercial C-band optical multi-format transmitter comprising driver amplifiers and a LiNbO3-based dual-polarization IQ-modulator. Two tunable External Cavity Laser (ECL) sources are used to cover the tested wavelengths in the C- (1527.5 – 1565 nm) and L- (1570 – 1600 nm) bands. The ECLs were set to a constant output power of approximately 16 dBm. The OLS consists of two 80-km and two 60-km SMF spans. Variable Optical Attenuators (VOA), placed in the link, allow to introduce degradation.

The architecture defined in Fig. 1a was setup for the control plane. A node agent, co-located with the Rx, collects and processes telemetry data, which is then published in the Apache Kafka message bus to be stored in the telemetry DB. OCATA exposes a REST API to the SDN controller and consumes available telemetry data.

The lightpath operation workflow in Fig. 1b has been experimentally assessed and the most relevant messages have been captured. In particular, Fig. 3 presents the content of messages 1, 6, and 8. Message 1 is the request sent by the SDN controller to the OCATA NDT for QoT estimation, as part of the lightpath setup process. The request includes the lightpath UUID, a list with the components along the selected route, the selected channel, the characteristics of each component, the modulation format, and the symbol rate. The data model used in this request

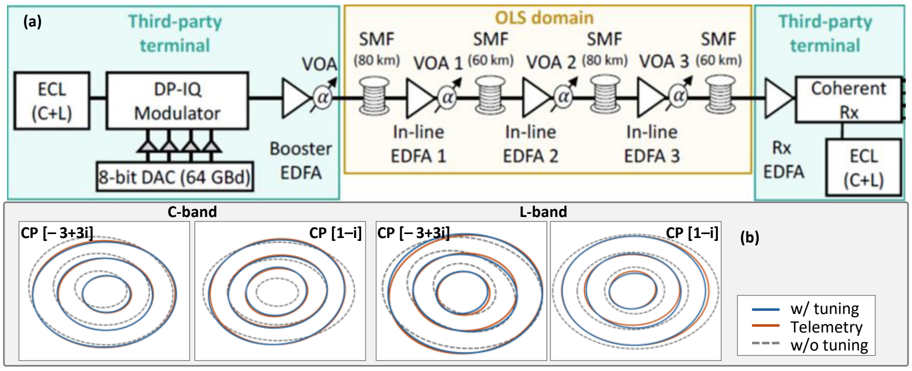

Fig. 2 Experimental Setup (a). Contours of the fitted GMM: Models vs Telemetry (b).

"id": "221cae86-94b3-4b80-b4c9-097f153bf559",
"route": ["A", "B", "C"],
"channel": 1,
"components":{
"A": {"type": "ROADM", "name": "ROADM_1"},
"B": {"type": "link", "name": "link_1"},
"C": {"type": "ROADM", "name": "ROADM_2"}},
"modulation_format": "16-QAM",
"symbol_rate": "32G"
(1)
"raw_samples": [["-3.08616", "2.98467i"],
:
["2.96258", "2.97984i"]]
(6)
"Y": [["2.96543", "3.09510",..,"0.00455"],
:
["1.03531", "1.01812",..,"0.00124"]],
"BER": "3.75e-3"
(8)

Fig. 3 Captures of selected messages

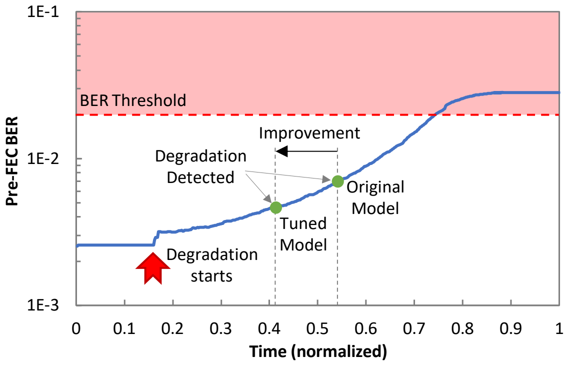

Fig. 4 Example of degradation detection

is consistent with that used in the Inventory, ensuring seamless interoperability between the SDN controller and the OCATA NDT. For illustrative purposes, we intentionally assume that only models for 80km spans are available, so lightpath models will overestimate the QoT of the lightpath. Message 6 contains the raw telemetry data collected from the Tx, which consists of a list with 10,000 IQ symbols. The FeX algorithm in the node agent processes the collected samples and summarizes each constellation point into features Y^i, computes the pre-FEC BER and publishes the processed samples in the Kafka bus (message 8) to be stored in the telemetry DB.

Fig. 2b illustrates the accuracy of model tuning by showing the contours of the fitted GMM for two CPs: one external, [-3+3i], and one internal, [1-i], across two channels, one in the C-band and the other in the L-band. The grey dashed contours represent the initially concatenated lightpath models, used during set up procedures. Although the QoT estimation allowed the lightpath to be established, we observe differences with the contours derived from the samples collected from telemetry. Such differences, and those between the measured and predicted pre-FEC BER values are minimized after the lightpath models are tuned. This allows accurate QoT evolution analysis and degradation detection process.

To demonstrate the effectiveness of the proposed NDT-assisted operation, we introduced a gradual degradation on the signal that increased the pre-FEC BER at the RX. Fig. 4 shows the evolution of the degradation as a function of the normalized time. The surveillance process periodically analyzes the collected telemetry data and compares the measured pre-FEC BER against the expected value from the lightpath model. For the sake of simplicity, any measured value that exceeds 2 times the expected one is assumed to be a performance degradation. For this experiment, we compared the measurements against the prediction from the original concatenated model and from the tuned model. Since the tuned model is more accurate than the original one, the surveillance process was able to detect the degradation at its earliest stage, which gives more time to the maintenance team to schedule repair activities. Note that this greatly reduces the operational costs of the network.

5. Conclusions

A complete lightpath operation workflow for MB optical networks has been presented and experimentally validated on a distributed testbed. The OCATA optical time-domain NDT is used for pre-FEC BER estimation during lightpath provisioning. Once the lightpath is set up, telemetry data is collected and locally processed to reduce its dimensionality and then stored in a centralized DB. Telemetry data is used to tune the NDT lightpath model, which resulted into large improvement for degradation detection.

References

[1] M. Devigili*et al.*, "Applications of the OCATA Time Domain Digital Twin: from QoT Estimation to Failure Management," JOCN 2024.
[2] P. Khare *et al.*, "Simulation and Modeling of C+L+S Multiband Optical Transmission for the OCATA Digital Twin," Sensors, 2025.
[3] S. Ghasrizadeh *et al.*, "Digital Twin -Assisted Lightpath Provisioning and Nonlinear Mitigation in MB Optical Networks," Sensors, 2024.
[4] S. Ghasrizadeh *et al.*, "Creating and Tuning Multiband Optical Transmission Digital Twin Lightpath Models," ONDM, 2025.
[5] L. Velasco *et al.*, "Distributed Intelligence for Pervasive Optical Network Telemetry," JOCN, 2023.
[6] M. Sena *et al.*, "DSP-Based Link Tomography for Amplifier Gain Estimation and Anomaly Detection in C+L-Band Systems," JLT, 2022.