

Demonstration of ML-assisted Soft-Failure Localization Based on Network Digital Twins

Kayol S. Mayer, Rossano P. Pinto, Jonathan A. Soares, Dalton S. Arantes, Christian E. Rothenberg, Vinicius Cavalcante, Leonardo L. Santos, Filipe D. Moraes, and Darli A. A. Mello

Abstract—In optical transport networks, failure localization is usually triggered as a response to alarms and significant anomalous behaviors. However, the recent evolution of network control and management leveraging software-defined networking (SDN) and streaming-based telemetry opens up new possibilities for automated methods that can localize even subtle anomalies, the so-called soft failures. This paper reports the experimental demonstration of a machine-learning-based soft-failure localization framework in a small-scale laboratory setup. The SDN telemetry setup includes ONOS-controlled transponders using NETCONF and an optical line system (OLS) providing telemetry via an OLS domain controller. A shallow artificial neural network (ANN) accomplishes ML-based failure localization with principal component analysis to reduce non-essential information. The ANN is trained by combining field data and synthetic data generated in a digital network twin. The field data trains the ANN to tolerate statistical variations in the network telemetry without failures, while the synthetic data generates artificial single-failure scenarios. We show that the soft-failure localization scheme successfully pinpoints the faulty element in all single failures generated in transponders, fibers, and amplifiers. We also demonstrate the system's ability to deal with double-failure scenarios.

Index Terms—Software-defined optical networks, soft-failure, failure localization, machine learning, neural networks.

I. INTRODUCTION

EFFECTIVE failure localization in optical transport networks is essential for proper network operation and service downtime mitigation [1]. Typically, alarm correlation techniques isolate the failure and set off maintenance actions [2], [3]. However, recent advances in software-defined

Manuscript received month day, year; revised month day, year and month day, year; accepted month day, year. Date of publication month day, year; date of current version month day, year. This work was supported by Padtec S/A – Project Number 29-P-19854/201. Kayol S. Mayer is supported in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. (Corresponding author: Kayol S. Mayer). This article was presented in part at the Optical Fiber Communications Conference, San Francisco, CA, USA, Mar. 2021. (Corresponding author: Rossano P. Pinto).

Kayol S. Mayer, Rossano P. Pinto, Jonathan A. Soares, Dalton S. Arantes, Christian E. Rothenberg, and Darli A. A. Mello are with the School of Electrical and Computer Engineering – FEEC, University of Campinas – Unicamp, Campinas, SP, Brazil (e-mail: kayol@decom.fee.unicamp.br; rossano@dca.fee.unicamp.br; jonathan@decom.fee.unicamp.br; dalton@unicamp.br; chesteve@unicamp.br; darli@unicamp.br).

Vinicius Cavalcante, Leonardo L. Santos, and Filipe D. Moraes are with Padtec S/A, Campinas, SP, Brazil (e-mail: vinicius.cavalcante@padtec.com.br; leonardo.santos@padtec.com.br; filipe.moraes@padtec.com.br).

This is an extended version of [1].

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

Digital Object Identifier xxxxxxxxxx

networking (SDN) [4] and intent-based networking (IBN) [5], [6] have contributed to novel control and management capabilities [7]–[9], including *soft-failure management* [10]–[19].

Unlike hard failures, which disrupt the service, soft failures are not severe enough to activate alarms. Eventually, the early repair of a soft failure can avoid the progressive degradation to a hard failure. Soft-failure management can be divided into the processes of *detection*, *localization*, and *identification*. The detection process notices anomalous behaviors without pinpointing the faulty element. The localization process pinpoints the faulty device. Finally, the identification process finds the cause of the failure. As failures in certain network elements affect network parameters distributed all over the network, soft-failure localization is a network-wide process. If-else rules can implement failure localization based on dependence trees applied to the network telemetry dataset. Nonetheless, telemetry parameters may be unavailable or not implemented in some components, particularly in disaggregated scenarios, requiring more sophisticated if-else rules. On the other hand, by re-fitting hyperparameters, machine learning (ML) techniques can automatically learn complex rules and even interpolate missing telemetry data [16], [20].

In recent years, several approaches have been proposed for soft-failure management in optical networks [20], [22]–[24]. Zhang et al. [25] use the extreme gradient boosting (XGBoost) algorithm and the Shapley additive explanations (SHAP) to find high-relevance features related to equipment failures for soft-failure detection. Tanaka et al. [26] detect fiber bends with a deep-neural-network-based diagnoses workflow. Liu et al. [27] detect failures via an autoencoder-based anomaly detection scheme. Aiming at failure identification, Vela et al. [12] use spectrum analyzers and optical test channels during commissioning testing and operation, and Shahkarami et al. [10] monitor the bit error rate (BER) in an experimental setup. Lun et al. [13], [28], Varughese et al. [29], and Sun et al. [30] identify failures using machine learning algorithms applied to receiver digital signal processing (DSP) features. In Shu et al. [31], soft-failure detection and identification are carried out by analyzing the digital spectrum of received signals. Musumeci et al. [32] use domain adaptation and transfer learning for failure detection and cause identification.

Soft failures eventually trigger anomalies in several network parameters, and localizing the original failure is a network-wide [33] process. Barzegar et al. [14], [17] accomplish soft-failure localization by monitoring the end-to-end performance of active lightpaths and looking for correlations. Date et al. [34] localize soft failures in wavelength-selective

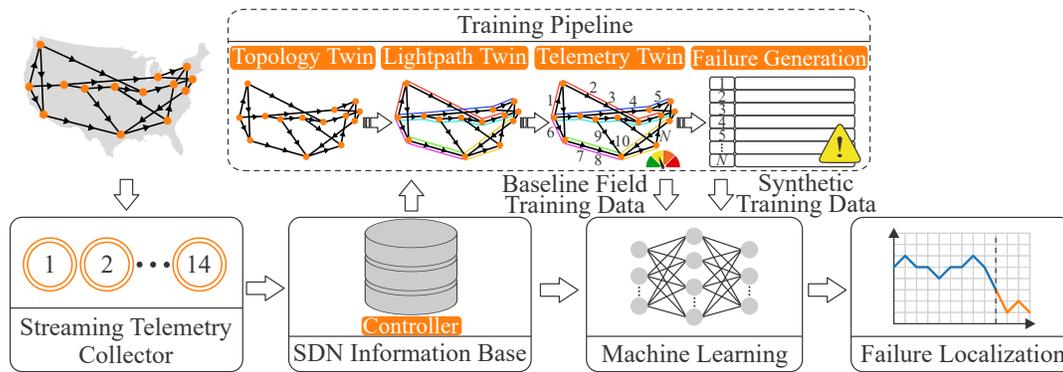


Fig. 1. Machine learning-based soft-failure localization framework. In contrast to the original framework proposed in [16], [21], the testbed demonstration requires the new arrow indicating “baseline field training data.” Telemetry data is retrieved by the streaming telemetry collector and fed into the SDN information base. An ANN localizes the failure based on training carried out by synthetic training data generated in a network digital twin. Training is also carried out by baseline field data to prepare the ANN for statistical variations and avoid false positives.

switches (WSSs) using correlations of DSP anomalies found in transponders. In [21], we develop a framework to localize soft failures using an artificial neural network (ANN) applied to network-wide parameters, following an SDN streaming telemetry service architecture. Unlike [14], [17], [34], which use explicitly programmed correlation rules to localize failures, our approach uses ML to automatically learn the relationship between failures and telemetry. The training of an ANN is carried out using synthetic telemetry data generated in a digital twin and actual readings from a live network. The ANN learns complex interdependencies and enables interesting features, such as missing data interpolation [16] and double-failure localization. The proposed technique was evaluated through simulations and experiments in an emulated scenario. In [16], we extend the solution proposed in [21] for scenarios of partial telemetry, improving the ANN results with principal component analysis (PCA). Furthermore, in [1], we validate our solution in a small-scale laboratory testbed.

In this paper, we extend [1] by providing an extensive analysis of all single-failure results. We also present a case study demonstrating that, although the failure localization algorithm was conceived for single-failure localization, it also successfully localizes a double-failure scenario, leveraging the extrapolation ability provided by ML algorithms.

The remainder of this paper is organized as follows. Section II presents the soft-failure localization framework and adaptations required for testbed operation. Section III describes the experimental setup. Section IV presents the failure localization results for all single failures and investigates the case study of double-failure localization. Lastly, Section V concludes the paper.

II. SOFT-FAILURE LOCALIZATION FRAMEWORK

The framework used for soft-failure localization is shown in Fig. 1. The original method was proposed in [21] and refined in [16]. Compared with the original method, the testbed demonstration [1] requires the new “baseline field training data” arrow shown in Fig. 1.

The telemetry collector [35] retrieves telemetry data from amplifiers and transponders. An SDN information base stores

the telemetry data consumed by the failure localization ML pipeline. The ML algorithm is composed of a shallow ANN [21]. The ANN input layer corresponds to the input telemetry data. Eventually, PCA can be applied to the telemetry inputs to speed up the training process and reduce the ANN computational complexity [16]. The output layer corresponds to all amplifiers, fiber links¹, and transponders that may fail. Its neurons use a Softmax activation function [13], [36], providing a probabilistic indication that an element has failed. The ANN hidden layer uses linear neurons, and its size is empirically adjusted according to the network dimensions. A training pipeline creates an ML training dataset through a network digital twin (NDT) [37], [38], where topology, lightpaths, spectrum allocation, and telemetry data (using analytic models of power propagation) are replicated.

In optical networks, failures are relatively infrequent, making it difficult to train supervised ML techniques using historical data. Alternatively, we generate an exhaustive training dataset using synthetic telemetry produced in the NDT. The training pipeline indicated in Fig. 1 generates artificial hard and soft failures in all amplifiers (boosters, in-line amplifiers, and preamplifiers), fiber links, and transponders. In our framework, soft and hard failures encompass amplifier gain degradation, additional fiber losses, and transponder power degradation. In practical optical networks, other components may fail (e.g., WSSs, splitters, or multiplexers and demultiplexers [39], [40]) but with lower failure rates [41], mainly the passive ones. Nevertheless, even if unmodeled failures occur in WSSs, splitters, multiplexers, and demultiplexers, we still expect a failure to be localized in the vicinity of the faulty device, similarly to the case of partial telemetry discussed in [16].

Although the synthetic telemetry dataset produced by the NDT reproduces the field telemetry data set with reasonable fidelity, statistical deviations in the field telemetry can trigger a failure localization process even in the absence of failures. To avoid false positives, in addition to the synthetic telemetry data

¹In this paper, we localize faulty fiber links, without pinpointing the exact failure coordinates. Eventually, the proposed algorithm could be combined with optical time-domain reflectometers (OTDRs) for extended capabilities.

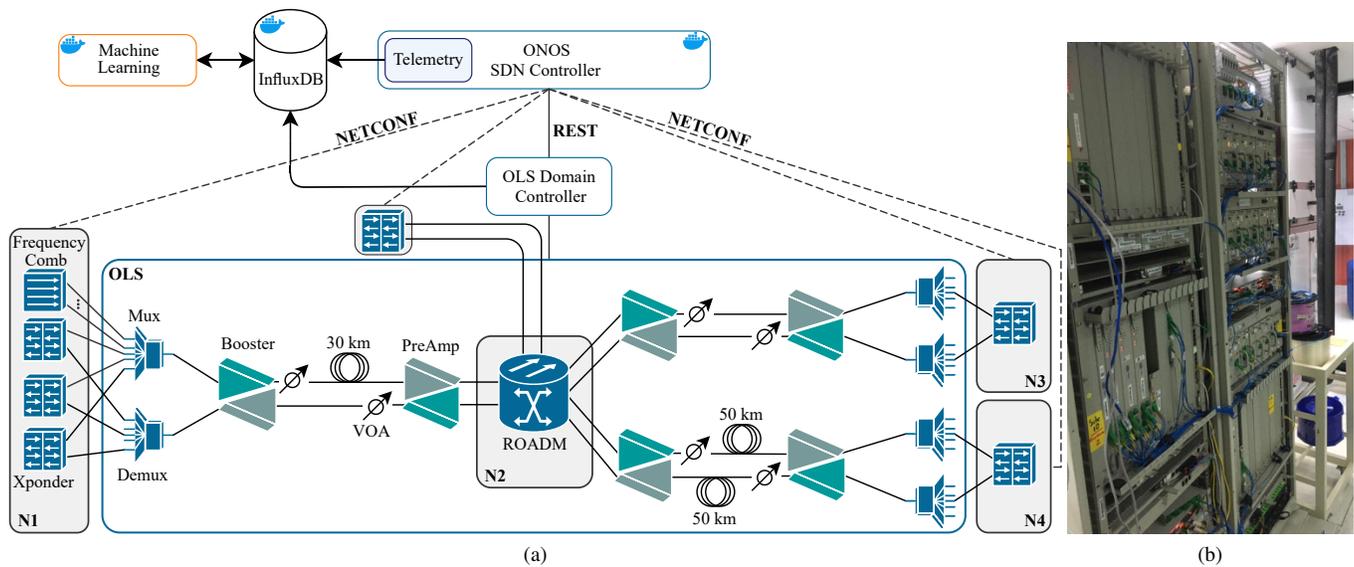


Fig. 2. Experimental setup for soft-failure generation and localization. The 4-node setup includes six bidirectional transponders and corresponding boosters and pre-amplifiers. The central node is a fully equipped reconfigurable add-drop multiplexed. A comb-generator in N1 generates unmodulated unidirectional lightpaths, with half of the lightpaths ending in N3 and the other half ending in N4. SDN telemetry data retrieved from the physical testbed is stored in an InfluxDB time series database. ML-based soft-failure localization is carried out in a public cloud. (a) Logical setup. (b) Physical testbed.

PORT ID	NAME	TYPE	ENABLED	MIN FREQ (THz)	MAX FREQ (THz)	GRID (GHz)	CURRENT OUTPUT POWER (dBm)	CURRENT INPUT POWER (dBm)	OSNR
3	Trx_1_4_1	OCH	true	190.7	195.45	50.0	0.00	-8.32	29.700000762939453
2	Trx_1_3_1	OCH	true	190.7	195.45	50.0	0.00	-7.44	31.100000381469727
1	Trx_1_2_1	OCH	true	190.7	195.45	50.0	0.00	-7.62	28.299999237060547
3	Trx_4_1_1	OCH	true	190.7	195.45	50.0	-4.98	-5.85	20.299999237060547
2	Trx_3_1_1	OCH	true	190.7	195.45	50.0	-9.98	-7.54	18.700000762939453
1	Trx_2_1_1	OCH	true	190.7	195.45	50.0	0.00	-12.44	23.799999237060547

Fig. 3. ONOS screenshot of controlled transponders implementing NETCONF.

sets, we also train the ANN with snapshots of field telemetry labeled with no failures, serving as a baseline. In this way, the ANN is forced to accept the baseline data set as a no-failure scenario, preparing the ANN for statistical variations. Adding these additional snapshots is essential for proper ANN operation in a practical scenario.

III. EXPERIMENTAL SETUP

The experimental testbed comprises four nodes (N1, N2, N3, and N4), as shown in Fig. 2a. Nodes N1, N3, and N4 are equipped with fixed optical add-drop multiplexers, while N2 is equipped with a reconfigurable add-drop multiplexer (ROADM) with broadcast and select (B&S) architecture. Nodes are interconnected by optical links with optical fibers, variable optical attenuators (VOAs), or both. Fig. 2b shows the physical setup.

Three bidirectional lightpaths are assigned in the optical network, linking nodes N1-N2, N1-N3, and N1-N4. In addition, a comb-generator in N1 generates unmodulated uni-

directional lightpaths, with half of the lightpaths ending in N3 and the other half ending in N4. The transponders are Padtec boards at 100 Gb/s and 200 Gb/s modulated with QPSK, 8-QAM, and 16-QAM formats. Transponders are connected to the ONOS SDN controller by means of NETCONF/YANG interfaces realized by an ONOS NETCONF driver developed to interact with the Padtec devices based on the `openconfig-terminal-device` YANG model with extensions. ONOS ODTN retrieves telemetry data from terminals using NETCONF get calls through a modified version of the ONOS App Roadm-GUI. Amplifier and wavelength-selective switch (WSS) streaming telemetry is carried out via an optical line system (OLS) domain controller using a proprietary protocol. All telemetry data collected by the OLS Domain Controller and the ONOS SDN controller are stored in an InfluxDB time series database. The setup can be considered as a partially disaggregated scenario, which is a natural step towards full disaggregation. Fig. 3 shows an ONOS screenshot for the six controlled transponders.



Fig. 4. Probability bar graph of the ANN outputs displayed with Grafana. In this example, we add an artificial 2-dB gain degradation to PreAmp_4_2 (preamplifier in the link interconnecting nodes N2 and N4). (a) ANN output before the failure (b) ANN output after the failure (67.1% empirical probability).

In order to keep a per-channel launch power of 0 dBm, the transponders output powers of nodes N1, N3, and N4 are adjusted in such a way to compensate for the insertion loss of add-drop multiplexers. As node N2 is a fully equipped ROADM, the target launch power of 0 dBm is obtained by activating the WSS equalization loop.

The experimental setup has 42 monitoring parameters, consisting of 24 amplifier input and output power values and 18 transponder parameters of OSNR, output power, and input power. A total of 24 devices may fail, counting six transponders, 12 amplifiers, and six unidirectional fiber links (optical links with fibers, VOAs, or both).

IV. EXPERIMENTAL EVALUATION

A. NDT construction and baseline training

The NDT is constructed using analytical models, computationally replicating the telemetry data measured with the system free of failures. The construction of the NDT first measures the fiber link attenuation by the ratio of input and output fiber powers of adjacent amplifiers. Likewise, amplifier gains are calculated from the ratio of output and input powers at a specific amplifier. The NDT is also fed with the target per-channel launch power after a power control loop in the ROADM. The per-channel transponder output power is set according to the values stored in the SDN controller.

After NDT construction, the training pipeline shown in Fig. 1 is executed. Based on [16], [21], a synthetic telemetry data set is generated considering the following soft- and hard-failure scenarios:

- Amplifier gain degradation of 3 dB, 10 dB (soft failures), and amplifier gain equal to 0 dB (hard failure).
- Transponder power degradation of 3 dB (soft failure) and output power of 0 W (hard failure).
- Additional fiber loss of 3 dB, 10 dB (soft failures), and attenuation that goes to infinity (hard failure).

To prepare the ANN for statistical variations in the no-failure condition and to avoid false positives, we also train the ANN with 22 baseline field telemetry snapshots labeled with no failures.

B. Single-failure localization

We apply the ANN approach with PCA to reduce the non-essential information from the dataset used for training. PCA reduced the number of ANN inputs from 42 to 30 parameters (reduction of 28.57%), keeping 99.9% of the input dataset

energy. Failure localization is accomplished by a shallow ANN with three layers [42], implemented in Python by the Keras library. The first layer has 30 inputs (corresponding to all collected telemetry data processed by PCA), the hidden layer has 100 linear neurons, and the output layer has 24 nonlinear neurons with the Softmax activation function [13], [36], corresponding to all network devices that may fail. The ANN output error is calculated by a categorical cross-entropy loss function [43]. Z-score normalization [44] is used to reduce the training time and contribute to numerical stability. Backpropagation is optimized by the infinite order (Adamax) [45] backpropagation algorithm. Training is performed in less than one second on Amazon Elastic Compute Cloud, at instance c5.9xlarge with 2nd generation Intel Xeon processor with a turbo frequency of 3.6 GHz, 72-GB RAM, and 36 vCPUs.

The training dataset comprises 88 input/output entries, where the failure generation block creates 66 synthetic entries, and the other 22 are collected from field baseline telemetry (without failures). After 100 training epochs, the combined ML algorithm reaches an accuracy of 100%. It is important to report that, as the ANN has a Softmax output layer, the outputs add to one even when there is no failure, distributing its failure levels uniformly among all outputs. Thus, in normal operation, the ANN outputs are approximately 1/24 (approximately 4.1% empirical failure probability). This floor level tends to zero for large optical networks with thousands of components (see [16]).

Fig. 4 shows the ANN outputs as a probability bar graph displayed at Graphana. Fig. 4a shows the ANN outputs when the network is free of failures. In this condition, the failure probability is low for all network components. Then, in this example, an artificial gain degradation of 2 dB is assigned to PreAmp_4_2 (preamplifier in the link interconnecting nodes N2 and N4). Fig. 4b shows the ANN output after the artificial failure. In this situation, the ANN output related to PreAmp_4_2 exceeds the others, reaching an empirical failure probability of 67.1%. To avoid incorrect failure localization caused by anomalous readings in the telemetry data we have employed a majority voting algorithm with a five-sample verification window (VW). If an ML output exceeds a certain failure threshold (e.g., 50% empirical failure probability), the failure index of the corresponding network component is sent to the VW. Otherwise, a no-failure index is sent to the VW. The VW operates as a FIFO (first-in, first-out) queue that stores the failure index for five consecutive time samples. The majority voting is computed for every sample entering the VW.

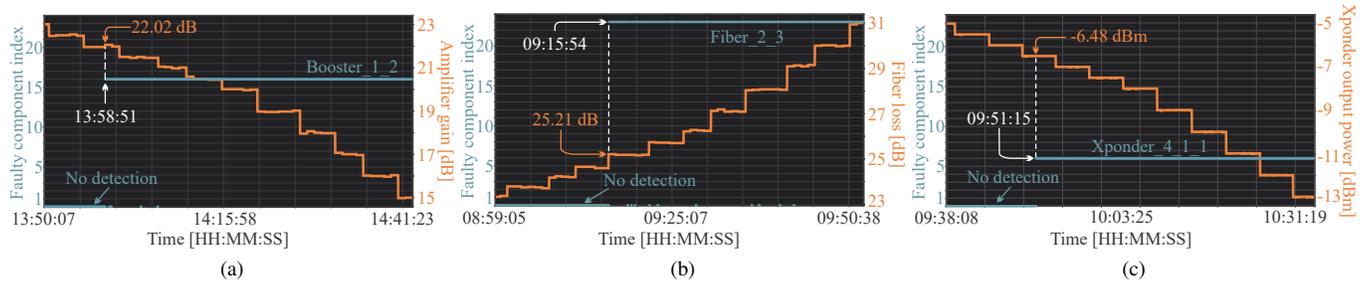


Fig. 5. Single-failure generation and localization. The orange curve is the output power level of the device being manipulated (right y-axis), while the blue curve is the faulty component index produced by the ANN (left y-axis). The faulty component index can take 25 levels, being 24 related to components and level zero corresponding to the no-failure state. (a) Booster-amplifier in the link interconnecting nodes N1 and N2 (Booster_1_2). (b) Fiber link interconnecting nodes N2 and N3 (Fiber_2_3). (c) Transponder located at node N4 that communicates with node N1 (Xponder_4_1).

TABLE I
SINGLE-FAILURE LOCALIZATION RESULTS.

Component	D_{FL} [dB]	T_{FL} [s]	Component	D_{FL} [dB]	T_{FL} [s]	Component	D_{FL} [dB]	T_{FL} [s]
Booster_1_2	0.98	185	PreAmp_2_3	1.92	80	Fiber_2_4	2.00	4
Booster_2_1	1.95	65	PreAmp_3_2	1.55	17	Fiber_4_2	1.42	5
Booster_2_3	2.13	65	PreAmp_2_4	1.46	107	Xponder_1_2	1.46	8
Booster_3_2	1.72	4	PreAmp_4_2	1.51	272	Xponder_2_1	1.48	5
Booster_2_4	2.54	4	Fiber_1_2	3.73	5	Xponder_1_3	1.97	4
Booster_4_2	1.49	15	Fiber_2_1	1.57	127	Xponder_3_1	1.51	54
PreAmp_1_2	1.49	4	Fiber_2_3	1.89	5	Xponder_1_4	1.48	6
PreAmp_2_1	2.90	5	Fiber_3_2	2.51	4	Xponder_4_1	1.50	122

D_{FL} and T_{FL} are the degradation and time to failure localization, respectively.

A component is declared as faulty whenever its index has the VW majority (i.e., a component index appears in three or more positions of the VW).

Fig. 5 presents the single-failure localization results for three randomly selected components (an amplifier, a fiber link, and a transponder). The orange curve is the output power level of the device being manipulated (right y-axis), while the blue curve is the faulty component index produced by the ANN (the left y-axis has 25 levels, being 24 for faulty components and one for no-detection indication). Fig. 5a shows the results for a gain degradation in the booster amplifier in the link interconnecting nodes N1 and N2 (Booster_1_2). After 185 s that the Booster_1_2 gain reduces from 22.48 dB to 22.02 dB (0.98 dB gain degradation), the ML algorithm localizes the faulty component. Fig. 5b presents the results of a soft failure causing extra losses in the fiber link interconnecting nodes N2 and N3 (Fiber_2_3). The correct component is correctly localized in five seconds, after the fiber loss transition from 24.59 dB to 25.51 dB, corresponding to 1.89 dB extra fiber loss. Fig. 5c depicts the results for an output power degradation in the transponder located at node N4 communicating with node N1 (Xponder_4_1). Failure localization is accomplished in 122 s after the output power degradation increases from 1 dB to 1.5 dB.

Table I summarizes the single-failure localization results for all amplifiers, fibers, and transponders in the evaluated testbed. The time to failure localization (T_{FL} [s]) depicts the time required to pinpoint a failure after reaching D_{FL} [dB] in a given component. The time T_{FL} [s] comprises the ML and majority vote processing times and the telemetry update time.

Failure localization is accomplished between 4 s and 272 s, with an average of 44 s. Also, a degradation between 0.98 dB and 3.73 dB (1.85 dB average degradation) was required to localize the faulty device. No incorrect failure localization events have been observed. Note that if a similar test were performed with longer intervals between degradation changes, the outcome would be similar. The ML approach relies on instantaneous input/output powers and OSNR values without addressing time-series properties. Different values of T_{FL} [s] are mainly dependent on heterogeneous telemetry update times for different components.

C. Double-failure localization case study

Although the proposed ML technique is trained to localize single failures, we also verified the system behavior under a double-failure case study. Different from the single-failure localization case, we analyze the ANN output directly, without a failure threshold and the majority vote algorithm.

Fig. 6 shows the ANN outputs as a probability bar graph. In this example, an output power degradation of 2 dB is assigned to Xponder_1_3 (transponder located at node N1 communicating with node N3) and Xponder_1_4 (transponder located at node N1 communicating with node N4). Fig. 6a shows the ANN outputs for the network with no failure. As in Fig. 4a, the probability of failure is low for each network component. In Fig. 6b, after assigning the output power degradation to Xponder_1_3 and Xponder_1_4, the ANN outputs related to these components exceed the others, reaching an empirical failure probability of 30.0% and 54.2%,



Fig. 6. Probability bar graph of the ANN output displayed with Grafana. In this example, we introduce a 2-dB output power degradation at Xponder_1_3 (transponder located at node N1 that communicates with node N3) and Xponder_1_4 (transponder located at node N1 that communicates with node N4). (a) ANN output before the failure. (b) ANN output after the failure, indicating an empirical probability of 30.0% (Xponder_1_3) and 54.2% (Xponder_1_4).

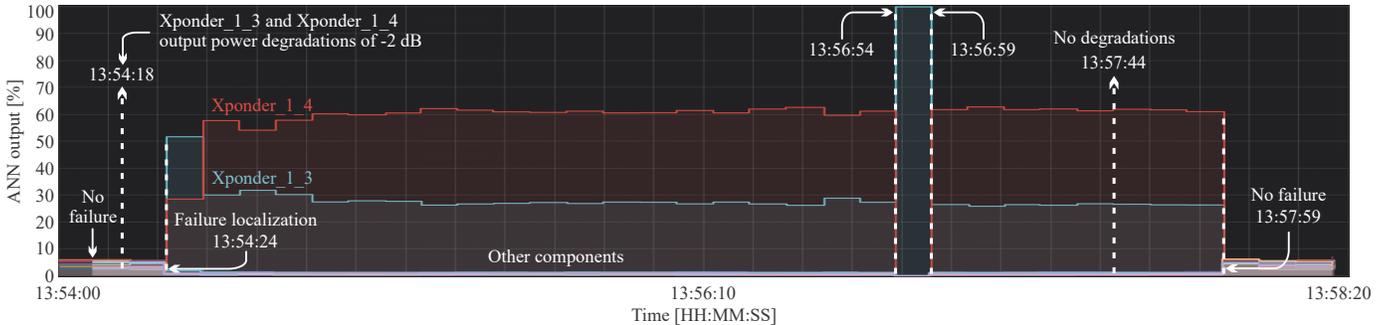


Fig. 7. Double-failure localization of Xponder_1_3 (Transponder located at node N1 that communicates with node N3) and Xponder_1_4 (Transponder located at node N1 that communicates with node N4). ML outputs displayed with Grafana in real time.

respectively. As the lightpaths originating in transponders Xponder_1_4 and Xponder_1_3 traverse different network segments with heterogeneous properties, including an extra 50-km fiber segment, the ANN produces different results. However, given the low interpretability of ANNs, it is difficult to assert the exact reason for having one value more prominent than the other.

Fig. 7 presents the ANN outputs for the double failure localization case, also displayed with Grafana. The experimental testbed starts from a no-failure state at 13:54:00. After 18 s, 2 dB degradations are assigned to transponders Xponder_1_3 and Xponder_1_4. Six seconds after creating the anomalies, the ANN outputs regarding Xponder_1_3 and Xponder_1_4 reach approximately 30% and 60%. This behavior is maintained until 13:56:54, where the Xponder_1_3 and Xponder_1_4 failure probabilities, respectively, go to almost 100% and 0%. This temporary anomalous condition is maintained for five seconds. Finally, the ANN outputs return to their previous states. At 13:57:44, the transponders power degradation is turned off, and at 13:57:59, the ANN outputs return to the initial no-failure state. Although double failures can occur in the network, usually one device fails first, triggering maintenance actions. Further work can attempt to improve double-failure localization conditioned to a pre-existing single-failure.

V. CONCLUSIONS

We present a testbed experimental demonstration of an ANN-based failure localization method using state-of-the-art SDN telemetry. While in previous theoretical and simulation

works ANN training was purely performed by synthetic data generated in a DNT, the practical demonstration reveals the need for training using a combination of synthetic and field data. Using field data for the no-failure network state accounts for statistical variations in the telemetry data, avoiding false positives in normal operation. The experimental setup also required a majority voting algorithm to filter out anomalous variations in telemetry. The results indicate a successful failure localization process for all single failures in the network. We also investigated the algorithm robustness in a double-failure case study. The results demonstrate the ANN ability to recognize double failures.

REFERENCES

- [1] R. P. Pinto *et al.*, "Demonstration of machine-intelligent soft-failure localization using SDN telemetry," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Francisco: IEEE, Jun. 2021, pp. 1–3.
- [2] D. Wang, L. Lou, M. Zhang, A. C. Boucouvalas, C. Zhang, and X. Huang, "Dealing with alarms in optical networks using an intelligent system," *IEEE Access*, vol. 7, pp. 97 760–97 770, Jul. 2019.
- [3] Q. Pham-Van *et al.*, "Demonstration of alarm correlation in partially disaggregated optical networks," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Diego: IEEE, Mar. 2020, pp. 1–3.
- [4] D. Kreutz, F. M. V. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proc. IEEE*, vol. 103, no. 1, pp. 14–76, Jan. 2015.
- [5] L. Velasco *et al.*, "End-to-end intent-based networking," *IEEE Commun. Mag.*, vol. 59, no. 10, pp. 106–112, Oct. 2021.
- [6] L. Velasco, S. Barzegar, F. Tabatabaeimehr, and M. Ruiz, "Intent-based networking and its application to optical networks [Invited Tutorial]," *J. Opt. Commun. Netw.*, vol. 14, no. 1, pp. A11–A22, Jan. 2022.
- [7] A. S. Thyagaturu, A. Mercian, M. P. McGarry, M. Reisslein, and W. Kellerer, "Software defined optical networks (SDONs): A comprehensive survey," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 4, pp. 2738–2786, Jul. 2016.

- [8] M. N. Hall, K.-T. Foerster, S. Schmid, and R. Durairajan, "A survey of reconfigurable optical networks," *Opt. Switch. Netw.*, vol. 41, pp. 1–19, Sep. 2021.
- [9] T. Szyrkowiec *et al.*, "Automatic intent-based secure service creation through a multilayer SDN network orchestration," *J. Opt. Commun. Netw.*, vol. 10, no. 4, pp. 289–297, Apr. 2018.
- [10] S. Shahkarami, F. Musumeci, F. Cugini, and M. Tornatore, "Machine-learning-based soft-failure detection and identification in optical networks," in *Proc. Opt. Fiber Commun. Conf. Expo.* San Diego: IEEE, Mar. 2018, pp. 1–3.
- [11] D. Rafique, T. Szyrkowiec, H. Griefser, A. Autenrieth, and J.-P. Elbers, "Cognitive assurance architecture for optical network fault management," *J. Lightw. Technol.*, vol. 36, no. 7, pp. 1443–1450, Apr. 2018.
- [12] A. P. Vela *et al.*, "Soft failure localization during commissioning testing and lighthouse operation," *J. Opt. Commun. Netw.*, vol. 10, no. 1, pp. A27–A36, Jan. 2018.
- [13] H. Lun *et al.*, "Soft failure identification for long-haul optical communication systems based on one-dimensional convolutional neural network," *J. Lightw. Technol.*, vol. 38, no. 11, pp. 2992–2999, Jun. 2020.
- [14] S. Barzegar *et al.*, "Soft-failure localization and device working parameters estimation in disaggregated scenarios," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Diego: IEEE, Mar. 2020, pp. 1–3.
- [15] S. Varughese, D. Lippiatt, T. Richter, S. Tibuleac, and S. E. Ralph, "Low complexity soft failure detection and identification in optical links using adaptive filter coefficients," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Diego: IEEE, Mar. 2020, pp. 1–3.
- [16] K. S. Mayer, J. A. Soares, R. P. Pinto, C. E. Rothenberg, D. S. Arantes, and D. A. A. Mello, "Machine-learning-based soft-failure localization with partial software-defined networking telemetry," *J. Opt. Commun. Netw.*, vol. 13, no. 10, pp. E122–E131, Oct. 2021.
- [17] S. Barzegar, M. Ruiz, A. Sgambelluri, F. Cugini, A. Napoli, and L. Velasco, "Soft failure detection, localization, identification, and severity prediction by estimating QoT model input parameters," *IEEE Trans. Netw. Service Manag.*, vol. 18, no. 3, pp. 2627–2640, Sep. 2021.
- [18] F. Musumeci, "Machine learning for failure management in optical networks," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Francisco: IEEE, Jun. 2021, pp. 1–28.
- [19] Z. Wang *et al.*, "Failure prediction using machine learning and time series in optical network," *Opt. Express*, vol. 25, no. 16, pp. 18553–18565, Aug. 2017.
- [20] F. Musumeci, C. Rottondi, G. Corani, S. Shahkarami, F. Cugini, and M. Tornatore, "A tutorial on machine learning for failure management in optical networks," *J. Lightw. Technol.*, vol. 37, no. 16, pp. 4125–4139, Aug. 2019.
- [21] K. S. Mayer, J. A. Soares, R. P. Pinto, C. E. Rothenberg, D. S. Arantes, and D. A. A. Mello, "Soft failure localization using ML with SDN-based network-wide telemetry," in *Proc. Eur. Conf. Opt. Commun.* Brussels: IEEE, Dec. 2020, pp. 1–4.
- [22] F. N. Khan, Q. Fan, C. Lu, and A. P. T. Lau, "An optical communication's perspective on machine learning and its applications," *J. Lightw. Technol.*, vol. 37, no. 2, pp. 493–516, Feb. 2019.
- [23] N. Guo, L. Li, L. Xiang, S. K. Bose, and G. Shen, "What if AI fails: protection against failure of AI-based QoT prediction," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Diego: IEEE, Mar. 2020, pp. 1–3.
- [24] H. Lun *et al.*, "Machine-learning-based telemetry for monitoring long-haul optical transmission impairments: methodologies and challenges [Invited]," *J. Opt. Commun. Netw.*, vol. 13, no. 10, pp. E94–E108, Oct. 2021.
- [25] C. Zhang *et al.*, "Cause-aware failure detection using an interpretable XGBoost for optical networks," *Opt. Express*, vol. 29, no. 20, pp. 31974–31992, Sep. 2021.
- [26] T. Tanaka, T. Inui, S. Kawai, S. Kuwabara, and H. Nishizawa, "Monitoring and diagnostic technologies using deep neural networks for predictive optical network maintenance [Invited]," *J. Opt. Commun. Netw.*, vol. 13, no. 10, pp. E13–E22, Oct. 2021.
- [27] S. Liu, D. Wang, C. Zhang, L. Wang, and M. Zhang, "Semi-supervised anomaly detection with imbalanced data for failure detection in optical networks," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Francisco: IEEE, Jun. 2021, pp. 1–3.
- [28] H. Lun *et al.*, "GAN based soft failure detection and identification for long-haul coherent transmission systems," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Francisco: IEEE, Jun. 2021, pp. 1–3.
- [29] S. Varughese, D. Lippiatt, T. Richter, S. Tibuleac, and S. E. Ralph, "Identification of soft failures in optical links using low complexity anomaly detection," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, San Diego, Mar. 2019, pp. 1–3.
- [30] K. Sun, Z. Yu, L. Shu, Z. Wan, and K. Xu, "Generalized soft failure identification enabled by digital residual spectrum and autoencoder," in *Proc. Opt. Fiber Commun. Conf. Exhib.* San Francisco: IEEE, Jun. 2021, pp. 1–3.
- [31] L. Shu, Z. Yu, Z. Wan, J. Zhang, S. Hu, and K. Xu, "Dual-stage soft failure detection and identification for low-margin elastic optical network by exploiting digital spectrum information," *J. Lightw. Technol.*, vol. 38, no. 9, pp. 2669–2679, Oct. 2020.
- [32] F. Musumeci *et al.*, "Domain adaptation and transfer learning for failure detection and failure-cause identification in optical networks across different lightpaths [Invited]," *J. Opt. Commun. Netw.*, vol. 14, no. 2, pp. A91–A100, Feb. 2022.
- [33] T. Pan, E. Song, C. Jia, W. Cao, T. Huang, and B. Liu, "Lightweight network-wide telemetry without explicitly using probe packets," in *Proc. IEEE Conf. Comp. Commun. Works.* Toronto: IEEE, Jul. 2020, pp. 1354–1355.
- [34] H. Date, T. Kubo, T. Kawasaki, and H. Maeda, "Silent failure localization on optical transport system," *IEEE Photon. Technol. Lett.*, vol. 33, no. 13, pp. 649–651, Jul. 2021.
- [35] J. Kundrát *et al.*, "Opening up ROADMs: streaming telemetry [Invited]," *J. Opt. Commun. Netw.*, vol. 13, no. 10, pp. E81–E93, Oct. 2021.
- [36] Y.-L. He, X.-L. Zhang, W. Ao, and J. Z. Huang, "Determining the optimal temperature parameter for softmax function in reinforcement learning," *Appl. Soft Comput.*, vol. 70, pp. 80–85, Sep. 2018.
- [37] D. Wang *et al.*, "The role of digital twin in optical communication: Fault management, hardware configuration, and transmission simulation," *IEEE Commun. Mag.*, vol. 59, no. 1, pp. 133–139, Jan. 2021.
- [38] A. Mahajan, K. K. Christodouloupoloulos, R. Martínez, R. Muñoz, and S. Spadaro, "Quality of transmission estimator retraining for dynamic optimization in optical networks," *J. Opt. Commun. Netw.*, vol. 13, no. 4, pp. B45–B59, Apr. 2021.
- [39] S. Barzegar, M. Ruiz, and L. Velasco, "Soft-failure localization and time-dependent degradation detection for network diagnosis," in *Proc. Inter. Conf. Transp. Opt. Netw.* Bari: IEEE, Jul. 2020, pp. 1–4.
- [40] H. Lun *et al.*, "ROADM-induced anomaly localization and evaluation for optical links based on receiver DSP and ML," *J. Light. Technol.*, vol. 39, no. 9, pp. 2696–2703, May 2021.
- [41] D. A. A. Mello, D. A. Schupke, M. Scheffel, and H. Waldman, "Availability maps for connections in WDM optical networks," in *Inter. Works. Design Reli. Commun. Netw.* Naples: IEEE, Oct. 2005, pp. 1–8.
- [42] M. Soltanolkotabi, A. Javanmard, and J. D. Lee, "Theoretical insights into the optimization landscape of over-parameterized shallow neural networks," *IEEE Trans. Inf. Theory*, vol. 65, no. 2, pp. 742–769, Feb. 2019.
- [43] X. Li, L. Yu, D. Chang, Z. Ma, and J. Cao, "Dual cross-entropy loss for small-sample fine-grained vehicle classification," *IEEE Trans. Veh. Technol.*, vol. 68, no. 5, pp. 4204–4212, May 2019.
- [44] A. Jain, K. Nandakumar, and A. Ross, "Score normalization in multimodal biometric systems," *Pattern Recognit.*, vol. 38, no. 12, pp. 2270–2285, Dec. 2005.
- [45] D. P. Kingma and L. J. Ba, "Adam: a method for stochastic optimization," in *Inter. Conf. Learn. Rep.*, San Diego, May 2015, pp. 1–15.