

ML-based Soft-failure Localization with Partial SDN Telemetry

KAYOL S. MAYER^{*}, JONATHAN A. SOARES, ROSSANO P. PINTO, CHRISTIAN E. ROTHENBERG,
DALTON S. ARANTES, AND DARLI A. A. MELLO

School of Electrical and Computer Engineering (FEEC), University of Campinas (UNICAMP), Av. Albert Einstein 400, Campinas, 13083-852, Brazil

*Corresponding author: kayol@decom.fee.unicamp.br

Compiled August 23, 2021

Soft-failure localization frameworks typically use if-else rules to localize failures based on the received telemetry data. However, in certain cases, particularly in disaggregated networks, some devices may not implement telemetry, or their telemetry may not be readily available. Alternatively, machine-learning-based (ML-based) frameworks can automatically learn complex relationships between telemetry and the fault location, incorporating information from the telemetry data collected network-wide (R3.-2#). This paper evaluates an ML-based soft-failure localization framework in scenarios of partial telemetry. The framework is based on an artificial neural network (ANN) trained by optical signal and noise power models that simulate the network telemetry upon all possible failure scenarios. The ANN can be trained in less than two minutes, allowing it to be retrained according to the available partial telemetry data. The ML-based framework exhibits excellent performance in scenarios of partial telemetry, practically interpolating the missing data. We show that in the rare cases of incorrect failure localization, the actual failure is in the localized device's vicinity. We also show that ANN training is accelerated by principal component analysis and can be carried out using cloud-based services. Finally, the evaluated ML-based framework is emulated in an SDN-based setup using the gNMI protocol for streaming telemetry.

© 2021 Optical Society of America

<http://dx.doi.org/10.1364/ao.XX.XXXXXX>

1. INTRODUCTION

In optical networks, efficient failure localization can mitigate service interruptions and avoid economic and social losses. Until recently, failure localization in optical networks has been carried out based on alarm correlation. However, recent advances in management plane technologies (model-based configuration and monitoring), software-defined networking (SDN) [1], and companion telemetry streaming services [2], offer increasingly advanced automation and intelligence levels. One of these advances is the management of soft-failures in optical networks. Soft-failures cause perceptible variations on network parameters but are not severe enough to generate alarms and disrupt the service [3]. Eventually, soft-failures evolve into hard-failures. Therefore, the early response to soft-failures avoid service disruptions and speed up maintenance actions. The management of soft-failures in optical networks has been classified into the processes of *detection*, *localization*, and *identification* [4]. Detection triggers alerts after telemetry data exceeds thresholds, such

as bit error rate (BER) or power levels. Localization pinpoints the faulty device, e.g., an amplifier, transponder, or fiber link. Identification finds the cause of the failure, such as a pump laser malfunctioning or excessive fiber bending. This paper addresses the problem of failure localization.

Recently, several approaches have been proposed for soft-failure management in optical networks [4–6]. Barzegar et al. [7] localize soft failures by monitoring the end-to-end performance of active lightpaths and looking for correlations [8]. Vela et al. [9] use spectrum analyzers and optical test channels to identify soft failures during commissioning testing and operation. Shahkarami et al. [10] monitor the BER to identify soft-failures in an experimental setup. Also dedicated to soft-failure identification, Lun et al. [11] apply a convolutional neural network to receiver DSP parameters. Similarly, Varughese et al. [12] apply support vector machines (SVM) to DSP adaptive coefficients for soft-failure identification. In Shu et al. [13], soft-failure identification is carried out by analyzing the digital spectrum of certain lightpaths. Wang et al. [8] and Rafique et al. [3] tackle

soft-failure prediction using time series. Although soft-failure detection and identification has received considerable attention in recent works, soft-failure localization is still an underexplored topic. Panayiotou et al. [14] apply ML to localize hard failures in fiber links based on historical data, but soft failures are not considered. Likewise, Srinivasan et al. [15] investigate failure localization and identification in complex networks without addressing optical layer parameters. Li et al. [16] demonstrate SDN failure localization based on alarm knowledge graphs but do not tackle soft-failures. Barzegar et al. [17] investigate soft-failure localization by correlating the signal-to-noise ratio (SNR) of lightpaths at reception. This approach localizes faulty links but does not pinpoint particular devices, such as a specific inline amplifier in a cascade of fiber spans. A soft-failure (e.g., an amplifier with degrading gain) eventually triggers anomalies in several network parameters, and localizing the original failure is a network-wide [18] process.

Soft-failure detection can be implemented by monitoring equipment telemetry and using thresholds [19] (R3-6#). For example, the attenuation on a link increasing by few dBs, or the input power on a transponder decreasing below a certain threshold, are clear failure indicators. Still, failures in critical devices can affect parameters in several portions of the network. This is the situation of amplifiers and fiber links, which are traversed by several lightpaths routed all over the network. One solution to isolate the failure is to build a dependency tree that prioritizes all failures, assigning a higher priority to elements that are higher in the dependency tree. The root cause is identified as the device with the highest priority.

Although the direct programming of if-else rules can be used for soft-failure localization, new ML-based technologies can automatically learn relationships based on training, eventually grasping complex associations that are not evident to a human operator. In addition, ML-based soft-failure localization manipulates the telemetry data in a more fluid fashion, contrasting with the hard thresholds involved in direct detection. We proposed in [20] a framework to localize soft-failures using an artificial neural network (ANN) applied to network-wide parameters following a gNMI-based SDN telemetry service architecture [2]. We experimentally validated the proposed framework in an emulated network scenario.

In this paper, we extend the results of [20] by evaluating the algorithm in scenarios of partial network telemetry, i.e., when part of the telemetry data is missing. This can be the case, e.g., with incompatible legacy non-SDN equipment, disaggregated environments with software and hardware non-uniformity, non-responsive devices, or issues in the network carrying the telemetry data. Certain failure localization strategies can also be modeled as partial telemetry. For example, using only the SNR for failure localization, such as in [17], and failure localization in networks with black-box amplifiers without monitoring, such as in [21], are also particular cases of partial telemetry. ML-based techniques allow retraining the failure localization algorithm to account for partial information, localizing failures using information distributed throughout the network. To our knowledge, this is the first work addressing soft-failure localization in scenarios of partial telemetry. In addition, this paper evaluates the use of principal component analysis (PCA) to speed up the ANN training process.

The paper is organized as follows. Section 2 presents the direct and ML-based network-wide soft-failure localization frameworks investigated in this paper. Section C describes the evaluation setup. Section D presents the soft-failure localization with

partial telemetry results. Section 5 presents the results of an emulated soft-failure localization setup scenario using state-of-the-art streaming-based telemetry. Lastly, Section 6 concludes the paper.

2. NETWORK-WIDE SOFT-FAILURE LOCALIZATION

A. Full and partial SDN telemetry

We investigate network-wide soft-failure localization based on full and partial telemetry. In this paper, we denote as *telemetry* the periodic and frequent retrieval of performance monitoring data by a centralized telemetry collector. Modern SDN streaming telemetry approaches [2] achieve high transfer rates with low overheads [21]. Although this paper focuses on gNMI-based streaming telemetry services, periodic and frequent telemetry can be eventually implemented in legacy infrastructure by polling monitored equipment.

In optical networks, failures in certain network devices can generate anomalies in several devices distributed over the entire network. For soft-failure localization, the most relevant monitored parameters are those directly related to the optical layer, such as input and output power, optical SNR, and BER. Recently, other parameters have been used for failure management, such as adaptive filter coefficients [11] and even temperature [8]. In this paper, we limit ourselves to certain parameters that are well-represented in a network digital mirror, or network digital twin [22]. We evaluate the input and output powers of amplifiers and transponders and the OSNR estimated by coherent transponders.

In the case of *full telemetry* (100%), we assume that the streaming telemetry collector has the current telemetry data of all monitored devices in the network. In the case of *partial telemetry*, we assume that a randomly selected fraction (e.g., 90%) of the telemetry data (input power, output power, and OSNR) is not available. We evaluate the ability of the ML-based algorithm to interpolate the missing data based on processing of the remaining data available.

B. Direct soft-failure localization

A baseline direct soft-failure localization framework is depicted in Fig. 1(a). An SDN streaming telemetry collector (STC) receives telemetry data and stores it in a time series database. The telemetry data is then analyzed by if-else rules to localize the correct failure.

A fiber link attenuation is calculated by the ratio of its input and output powers. The amplifier gain is calculated by the ratio of output and input powers. The transponder OSNR, input power, and output power are measured directly. After detecting potential failures by a simple threshold, the direct failure localization framework is carried out. If-else rules are constructed from a dependence tree based on a priority list for faulty devices. Upstream devices receive higher priority, while downstream devices receive lower priority. The faulty device is localized as the one with the highest priority. For example, suppose that an optical fiber link or amplifier fails. In that case, anomalies are detected in the optical signal to noise ratio (OSNR) at the input of all transponders whose lightpaths traverse the faulty device. As these transponders are downstream the faulty device, they receive a lower priority, while the faulty device receives a higher priority.

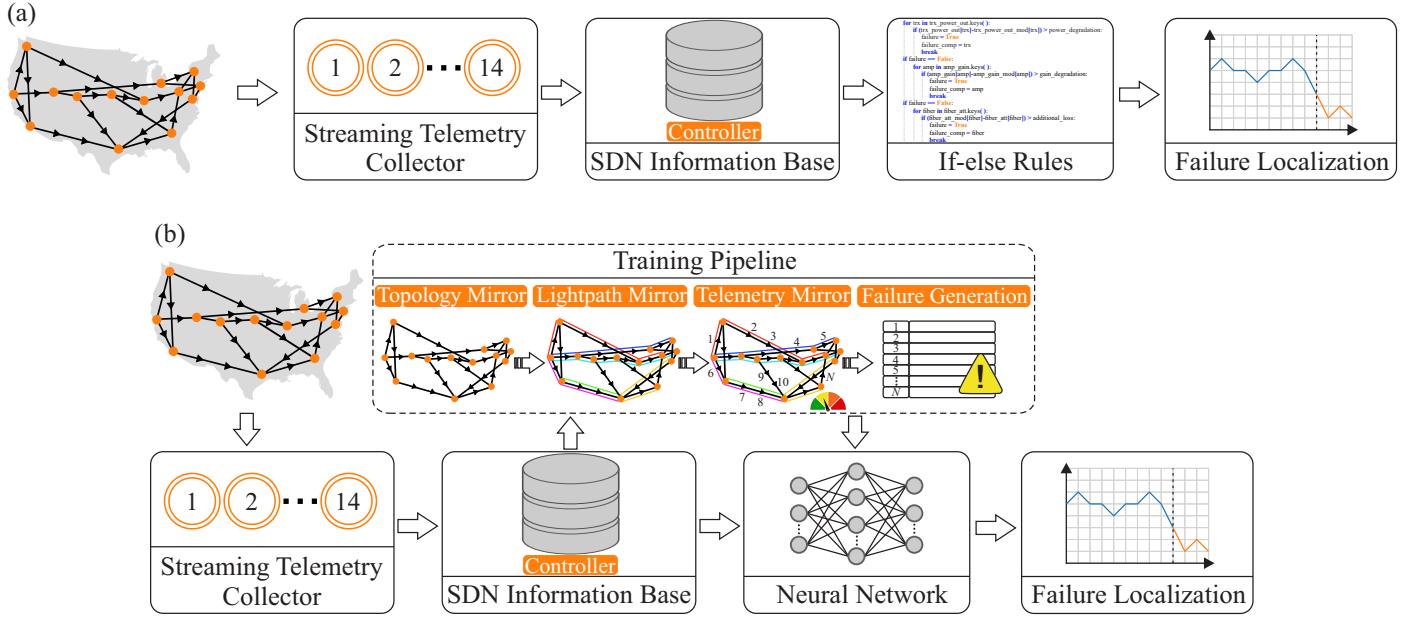


Fig. 1. Network-wide soft-failure localization frameworks. (a) Direct soft-failure localization. An STC receives the telemetry data and stores it in a time series database. The soft-failure is localized using if-else rules based on a dependence tree. (b) ML-based soft-failure localization. An STC receives the telemetry data and stores it in a time series database. In parallel, a mirror of the telemetry data is constructed using synthetic signal and noise power models. In the mirror, all possible failures are generated and used to train a neural network. The soft-failure is localized by applying the received telemetry data to the neural network.

C. ML-based soft-failure localization

Figure 1(b) presents the ML-based framework evaluated in this paper. The failure localization workflow starts by creating a snapshot mirror of the optical network using the current SDN information base, including physical topology (topology mirror), routed lightpaths (lightpath mirror), and telemetry data (telemetry mirror). This information is sent to a failure generation simulator, which produces synthetic telemetry data for all scenarios considering any possible failure in the network as inputs to train the ANN used for failure localization. The training phase must be carried out wherever there is a change in the network configuration, such that the network mirror ceases to represent a desirable network state. This can occur, e.g., upon lightpath activation or deactivation, or protection switching. The monitored devices send telemetry data to the SDN STC, which stores the collected data into a time series database. After training, the ANN module is continuously fed with the telemetry data from the time-series database to localize potential network failures. In practical scenarios, the time-series database also communicates with the network mirror to feed telemetry data for improving signal propagation models and achieving more reliable mirroring.

In partial-telemetry scenarios, not all telemetry data is available for failure localization. Certain devices may not implement SDN telemetry, and others may become non-responsive. If non-responsive devices stop sending telemetry data to the SDN STC for a certain period of time, the SDN controller can request ML retraining. As the neural network size depends on the number of inputs, retraining optimizes performance. In partial telemetry retraining, the topology and lightpath mirrors remain the same while missing data is removed from the telemetry mirror. As we will see later, retraining can be carried out in less than two minutes, even for large-scale networks containing thousands of

components. (R3.-6#)

Figure 2 shows the signal and noise power models used to generate synthetic telemetry data in simulations. (R3.-6#) We neglect nonlinear effects in the fiber power and noise model and consider a relatively low per-channel launch power. The per-channel launch power is kept constant by a feedback control loop involving a wavelength-selective (WSS) switch and an optical channel monitor (OCM). The power model keeps track of the signal and noise powers and calculates the OSNR at the end of the link. Reconfigurable optical add-drop multiplexers implement the broadcast and select architecture. More elaborate propagation models could be employed, e.g., using ML-based QoT estimation tools. However, the detailed evaluation of these models is out of the scope of this paper.

3. SIMULATION SETUP

A. Topology and traffic workload

To evaluate our soft-failure localization proposal in differently sized and connected topologies, we select two well-known and representative topologies: the 14-node 22-link National Science Foundation Network (NSFNet) [23], presented in Fig. 3(a), and the 17-node 26-link German network (GNet) [24], presented in Fig. 3(a). The NSFNet and GNet networks are used as references to demonstrate the concept of soft-failure localization using partial SDN telemetry. Although both networks exhibit similar connectivity, the NSFNet has a considerably higher number of fiber spans and in-line amplifiers, complicating the failure localization process.

The simulation-based evaluation is based on the following assumptions: (i) C-band (R3.-4#) 4.8-THz optical link spectrum corresponding to 96 50-GHz frequency slots (FS), (ii) 1,000 sequential uniformly distributed demands [25], and bandwidth

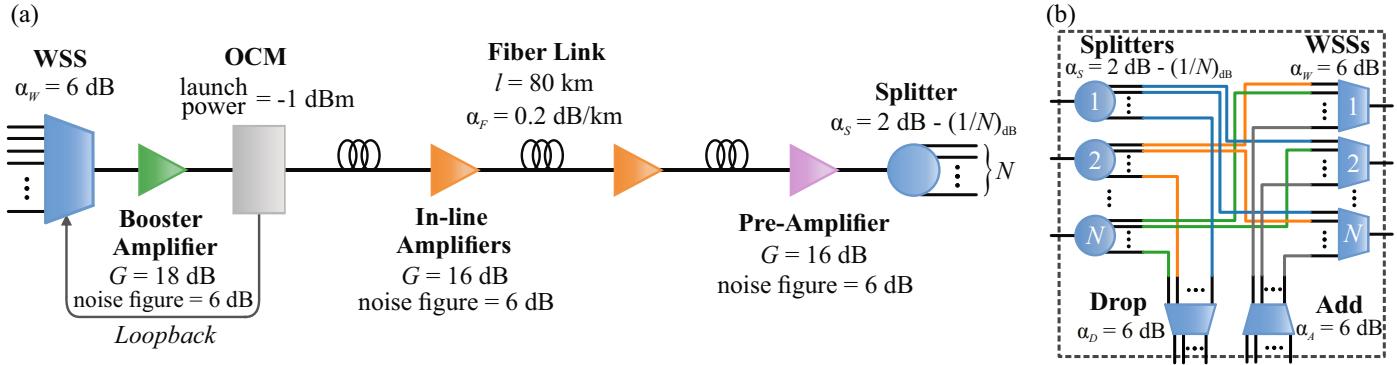


Fig. 2. Synthetic telemetry generation models. (a) Optical signal and noise power model. The optical per-channel launch power is maintained constant by a feedback loop involving WSS and OCM. (b) ROADM, assuming a broadcast and select architecture.

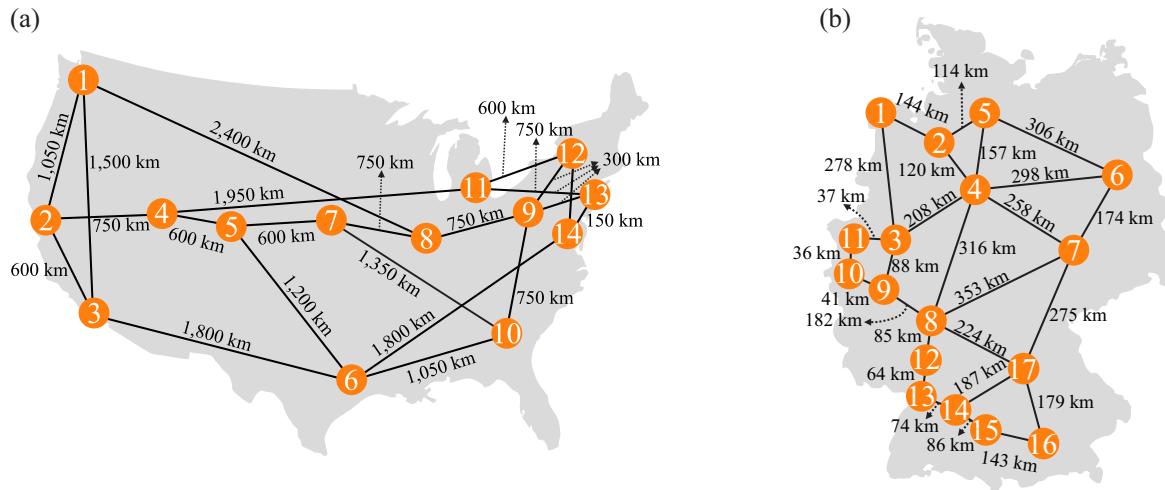


Fig. 3. Network topologies: (a) NSFNet. (b) GNet.

uniformly distributed between 1 to 4 FSs, (iii) shortest-path routing based on Dijkstra's algorithm [26], (iv) first-fit algorithm for wavelength assignment [27], (v) broadcast and select (B&S) reconfigurable add-drop multiplexers (ROADMs) with a per-channel power control loop based on optical channel monitors (OCMs) and wavelength selective switches (WSSs), and (vi) a control loop ensuring -1 dBm launch power per channel.

Figure 2 summarizes the simulation parameters, where N is the node degree, and l is the fiber length per span. We assume 80-km spans with $\alpha_F = 0.2 \text{ dB/km}$ attenuation, except for the last one, which ranges between 50 km and 120 km to result in the desired total span length. The NSFNet (GNet) topology features 524 (110) unidirectional fiber spans, with 480 (58) in-line amplifiers, 44 (52) pre-amplifiers, and 44 (52) booster amplifiers. We load the physical topology with bidirectional connections generated considering uniform traffic. Out of the 1,000 demands, 393 (422) are accepted and delivered by 786 (844) transponders.

B. Telemetry setup

The evaluation scenario computes 3,494 (3,018) monitoring parameters, consisting of 1,136 (486) amplifier input and output power values and 2,358 (2,532) transponder parameters of OSNR, output power, and input power. We assume coherent receivers that are able to estimate OSNRs from the received constellation. A total of 1,878 (1,116) devices may fail, count-

ing 786 (844) transponders, 568 (58) amplifiers, and 524 (110) unidirectional fiber spans.

Transponder and amplifier telemetry data (amplifier input and output powers and transponder OSNR, input and output powers) are streamed by an individual telemetry server per node. Each node streams telemetry data to the STC, which updates the SDN information base. The ANN module retrieves telemetry data from the SDN information base and localizes a faulty device.

C. ANN design and training

This section details the ANN design and training process for the NSFNet case study. The results for the GNet are analogous and are omitted here for the sake of clarity. Failure localization is accomplished by a shallow ANN with three layers [28], implemented in Python by the Keras library. The first layer has 3,494 inputs (corresponding to all collected telemetry data), the hidden-layer has 1,000 linear neurons, and the output-layer has 1,878 nonlinear neurons with the Softmax activation function [11, 29, 30], corresponding to all network devices that may fail. The ANN output error is calculated by a categorical cross-entropy loss function [31]. The use of Z-score normalization [32] abbreviates the training time and contributes to numerical stability. Backpropagation is optimized by the infinite order (Adamax) [33] backpropagation algorithm. Also, a batch size of

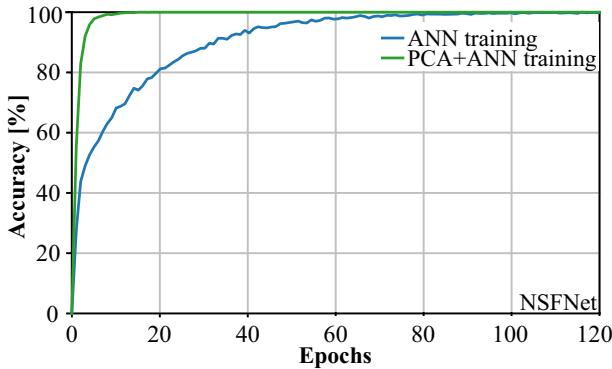


Fig. 4. Accuracy during training phase of ML algorithms for the NSFNet. Results for the GNet are equivalent.

100 samples is used to accelerate training. A device is identified as faulty if its output exceeds 0.5. The training is performed 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. Training, shown in the blue curve in Fig. 4, is performed with the following soft-failure scenarios [20]:

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

This work is focused on localizing failures in amplifiers, transponders, and unidirectional fiber spans, which include most components in an optical network. Other components may fail (e.g. WSSs, splitters, or multiplexers and demultiplexers [34, 35]). However, in general, these components, particularly the passive ones, have lower failure rates [36]. In any case, even if unmodelled failures occur, we still expect a failure to be localized in the vicinity of the faulty device.

D. Improvements through principal component analysis

To reduce the non-essential information from the dataset used for hard- and soft-failure training, the principal component anal-

ysis (PCA) [37, 38] technique is applied to reduce the number of input parameters, keeping 99.99% of the input dataset energy. PCA reduces the number of neural network inputs to 2,245 parameters, a reduction of 35.75%. After 120 training epochs (44-min dataset generation, 0.09-min PCA dimensionality reduction, and 1.12-min ANN training time), the combined ML algorithm reaches a categorical accuracy of 100%, as shown in the green curve in Fig. 4. The PCA+ANN technique achieves a stable accuracy after only 10 training epochs. In contrast, the ANN technique without dimensionality reduction requires more than 70 training epochs to reach a stable accuracy.

4. SIMULATION RESULTS

We first evaluate the effectiveness of the ML-based soft-failure localization algorithm using PCA to accelerate the training phase. The results are presented in Figs. 5(a)-(c). We assess the failure localization performance by evaluating the relative frequency that the correct device is identified as faulty (blue curves), the incorrect device is identified as faulty (orange curves), and no failure is identified (green curves). The ANN performance is shown by the solid curves, while the PCA+ANN result is shown by the dashed lines. In the x-axis, a degradation equal to zero indicates that the system is operating properly and that all values collected by telemetry correspond exactly to those stored in the baseline network mirror. A degradation in amplifier gain, fiber loss, or transponder power causes anomalies in the network telemetry that will differ from the baseline mirror. These anomalies affect the output of the ANN, which attempts to pinpoint the exact root cause. For a degradation close to 0, there is no detection, as it is too subtle to be detected by the ANN. For a higher degradation, the correct localization frequency starts to increase as anomalies start to generate significant variations at the ANN output. Fig. 5(a) shows that the ANN correctly localizes all soft-failures in amplifiers for gain degradations greater than 2 dB. Fig. 5(b) shows that it also localizes all soft-failures in fiber links causing losses greater than 2 dB. Likewise, Fig. 5(c) shows that all transponder failures are correctly localized if their power reduction is greater than 2 dB. The results indicate that PCA accelerates the training process without affecting the ANN ability to localize soft-failures. Therefore, all remaining results are generated using PCA.

Compared with direct failure localization frameworks, ML-based failure localization has the benefit of localizing failures even in the absence of the full telemetry dataset ([R3-2#](#)). There-

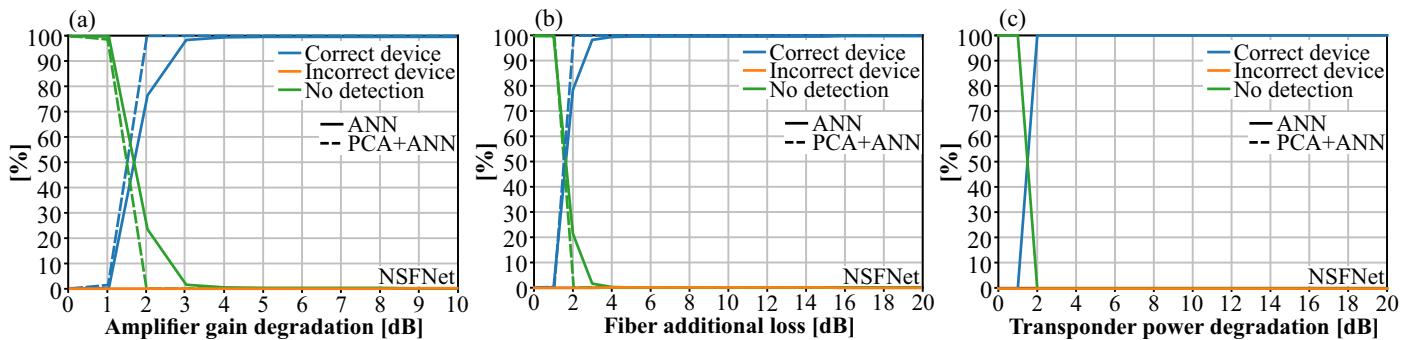


Fig. 5. Soft-failure localization results with ANN and PCA+ANN. (a) Amplifier failure localization. (b) Fiber failure localization. (c) Transponder failure localization. The figures show the relative frequency that the correct device is identified as faulty (blue curves), the incorrect device is identified as faulty (orange curves), and no failure is identified (green curves). PCA accelerates the training phase while retaining the failure localization accuracy.

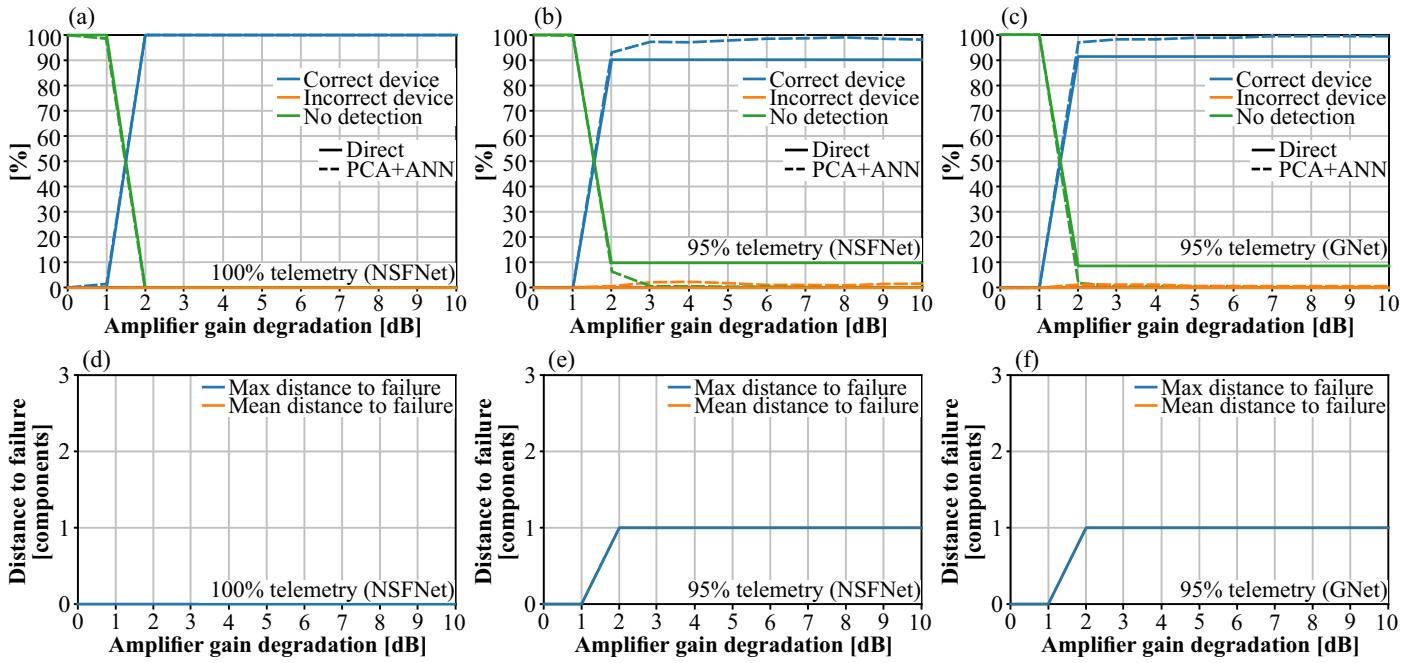


Fig. 6. Amplifier soft-failure localization results with complete and partial information of 95% telemetry, using direct and ML-based failure localization. (a), (d) NSFNet with 100% telemetry (results for the GNet are equivalent); (b), (e) NSFNet with 95% telemetry; (c), (f) GNet with 95% telemetry. Figures (a), (b), and (c) show the relative frequency that the correct device is identified as faulty (blue curves), the incorrect device is identified as faulty (orange curves), and no failure is identified (green curves). The distance to failure metric in Figs. (d), (e), and (f) applies to cases where the ML-based algorithm detects a failure in an incorrect device. The figures show the maximum distance to failure (blue curve) and the mean distance to failure (orange curve).

fore, we evaluate the ability of the ML-based framework to interpolate missing data by randomly removing part of the available telemetry. The ANN used for soft-failure localization is retrained to take into account the missing information.

Figs. 6(a)-(f) show the results of the investigated failure localization frameworks to detect an amplifier gain degradation. Figs. 6(a), (d) show that, with 100% telemetry, the algorithm detects all amplifiers whose gain degrades by more than 2 dB. Although we only show the 100% telemetry results for the NSFNet, the results for the GNet are equivalent. In Figs. 6(b), (e) (NSFNet) and 6(c), (f) (GNet) we randomly remove 5% of the telemetry data, leaving 95% for failure localization. Figs. 6(b), (c) shows that the direct failure localization framework is unable to operate in the absence of data, exhibiting a constant false-negative probability. The ML-based framework, in turn, recognizes the amplifier failure correctly with a probability higher than 99% for amplifier gains degraded by more than 3 dB. However, there is a residual probability of 1% that an incorrect device is localized. Fortunately, Figs. 6(e), (f) show that even if a device failure is incorrectly localized, it is just one device apart from the actual failure.

Figs. 7(a)-(f) show the failure localization results in case of an additional fiber loss. Figs. 7(a), (d) show that, with 100% telemetry, the algorithm detects all fibers whose attenuation degrades by more than 2 dB. Although we only show the 100% telemetry results for the NSFNet, the results for the GNet are equivalent. In Figs. 7(b), (e) (NSFNet) and 7(c), (f) (GNet) we randomly remove 5% of the telemetry data, leaving 95% for failure localization. Figs. 7(b), (c) show again that the direct failure localization framework is unable to operate in the absence of data, exhibiting a constant false-negative probability. For

both the NSFNet and the GNet, the ML-based algorithm fully interpolates missing data for high additional losses, correctly localizing the failure. For low and intermediate additional losses, there is a probability of incorrect localization. Again, Figs. 7(e), (f) show that the localized failure is always one or two devices away from the actual one. Comparing the results of the NSFNet with those for the GNet confirms that the algorithm exhibits slightly better performance for networks with short distances and consequently a lower number of devices.

Figs. 8(a)-(f) show the failure localization results in case of a transponder with reduced output power. Figs. 8(a), (d) show that, with 100% telemetry, the algorithm detects all transponders whose attenuation degrades by more than 2 dB. Although we only show the 100% telemetry results for the NSFNet, the results for the GNet are equivalent. In Figs. 8(b), (e) (NSFNet) and 8(c), (f) (GNet) we randomly remove 5% of the telemetry data, leaving 95% for failure localization. Figs. 8(b), (c) shows that both the direct and ML-based failure localization frameworks exhibit a loss in performance with partial telemetry, expressed in terms of false-negatives. As shown in Figs. 8(e), (f), incorrectly localized failures are not observed. **The ML-based algorithm outperforms the direct framework for a transponder power degradation higher than 10 dB, effectively identifying the faulty device and reducing false negatives. From this point, the WSS/OCM control loops lose the ability to equalize the power spectrum, and the transponder power degradation is sensed in other network devices. (R3.-5#)**

Table 1 extends the results to 90%, 85%, and 80% partial telemetry, considering a soft-failure scenario of 3 dB degradation (fiber loss, amplifier gain, and transponder power). We again assume that, upon detecting missing data, the direct scheme

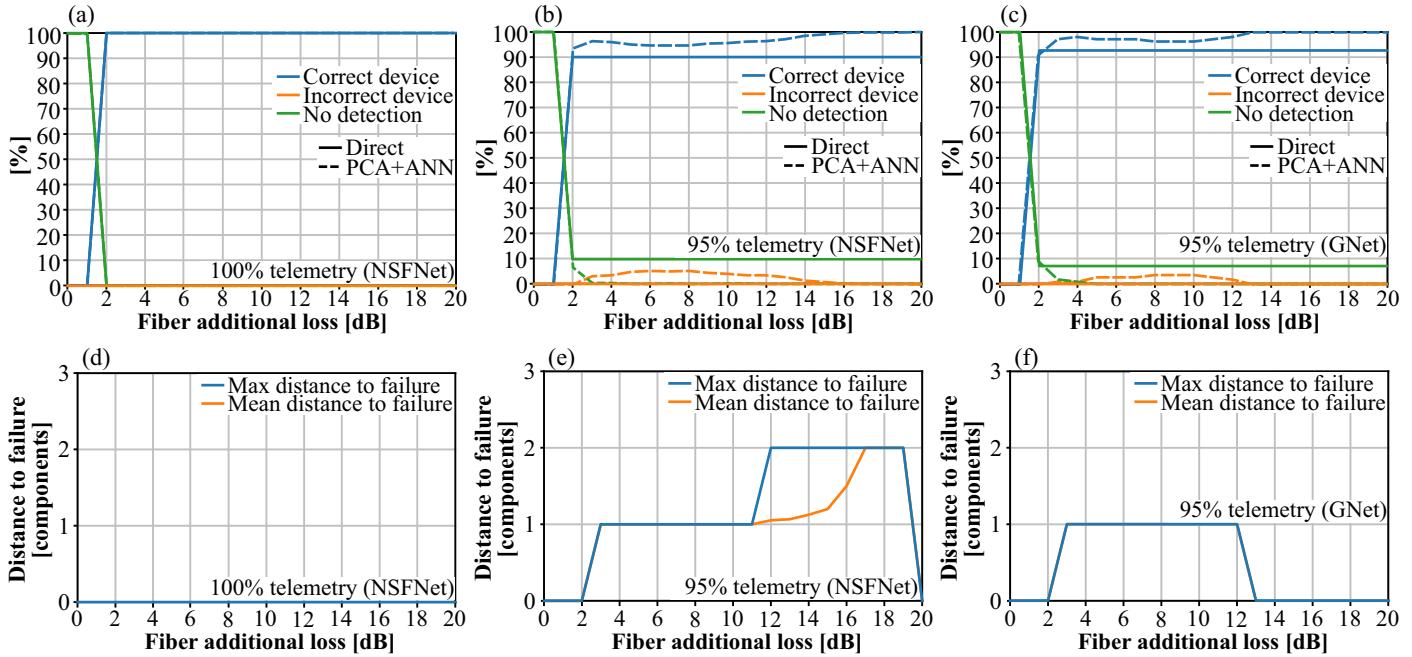


Fig. 7. Fiber link soft-failure localization results with complete and partial information of 95% telemetry, using direct and ML-based failure localization. (a), (d) NSFNet with 100% telemetry (results for the GNet are equivalent); (b), (e) NSFNet with 95% telemetry; (c), (f) GNet with 95% telemetry. Figures (a), (b), and (c) show the relative frequency that the correct device is identified as faulty (blue curves), the incorrect device is identified as faulty (orange curves), and no failure is identified (green curves). The distance to failure metric in Figs. (d), (e), and (f) applies to cases where the ML-based algorithm detects a failure in an incorrect device. The figures show the maximum distance to failure (blue curve) and the mean distance to failure (orange curve).

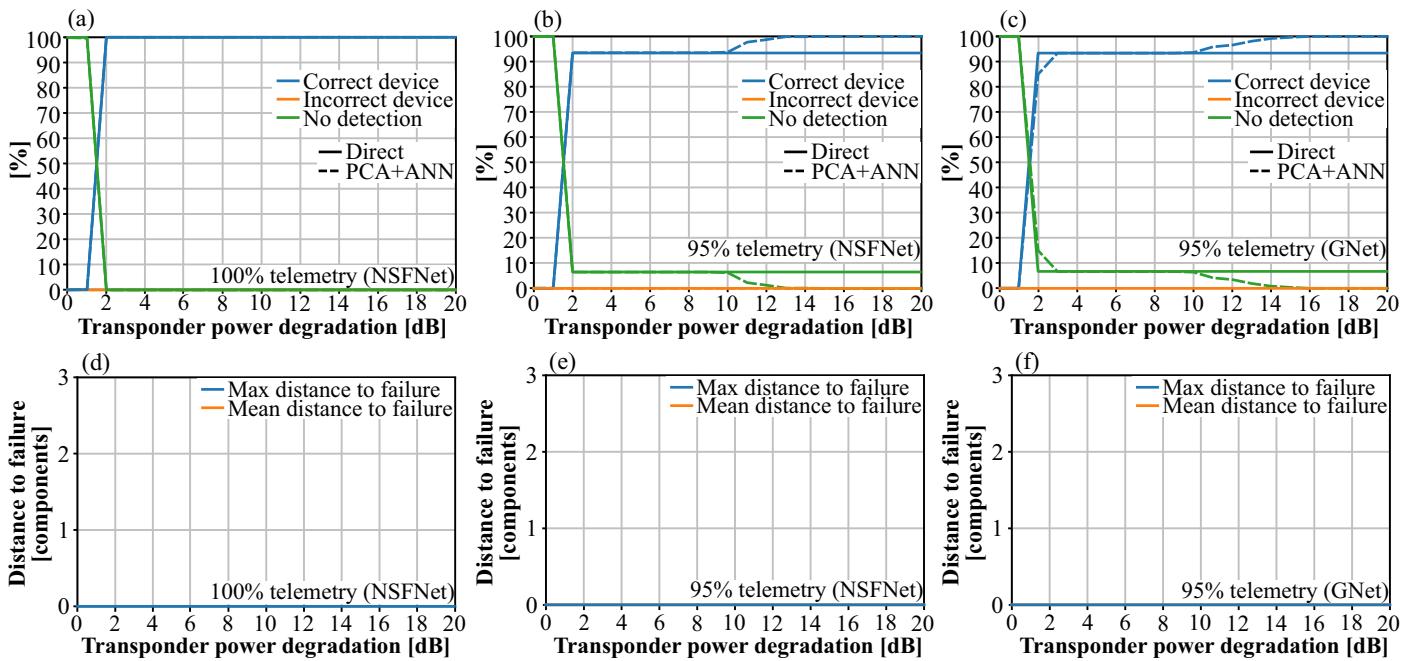


Fig. 8. Transponder soft-failure localization results with complete and partial information of 95% telemetry, using direct and ML-based failure localization. (a), (d) NSFNet with 100% telemetry (results for the GNet are equivalent); (b), (e) NSFNet with 95% telemetry; (c), (f) GNet with 95% telemetry. Figures (a), (b), and (c) show the relative frequency that the correct device is identified as faulty (blue curves), the incorrect device is identified as faulty (orange curves), and no failure is identified (green curves). The distance to failure metric in Figs. (d), (e), and (f) applies to cases where the ML-based algorithm detects a failure in an incorrect device. The figures show the maximum distance to failure (blue curve) and the mean distance to failure (orange curve).

Table 1. Soft-failure localization results with complete and 95%, 90%, 85%, and 80% partial telemetry, using direct and ML-based failure localization. The results show the relative frequency that the correct device is identified as faulty, the incorrect device is identified as faulty, and no failure is identified. The generated soft failure corresponds to a 3 dB-degradation in amplifier gain, fiber loss, and transponder power. For all incorrect localization cases using the ML-based technique, the actual failure is just one device apart.

| Algorithm | Telemetry | Localization | NSFNet | | | GNet | | |
|--------------|-----------|--------------|-----------|------------|-------------|-----------|------------|-------------|
| | | | Amplifier | Fiber link | Transponder | Amplifier | Fiber link | Transponder |
| Direct-based | 100.00% | Correct | 100.00% | 100.00% | 100.00% | 100.00% | 100.00% | 100.00% |
| | | Incorrect | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% |
| | | No detection | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% |
| | 95% | Correct | 90.14% | 90.27% | 93.51% | 91.36% | 92.73% | 93.36% |
| | | Incorrect | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% |
| | | No detection | 9.86% | 9.73% | 6.49% | 8.64% | 7.27% | 6.64% |
| | 90% | Correct | 82.75% | 81.68% | 88.30% | 82.72% | 82.73% | 88.74% |
| | | Incorrect | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% |
| | | No detection | 17.25% | 18.32% | 11.70% | 17.28% | 17.27% | 11.26% |
| ML-based | 85% | Correct | 74.30% | 72.90% | 84.99% | 74.07% | 72.73% | 84.83% |
| | | Incorrect | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% |
| | | No detection | 25.70% | 27.10% | 15.01% | 25.93% | 27.27% | 15.17% |
| | 80% | Correct | 64.96% | 62.79% | 80.15% | 62.96% | 61.82% | 79.98% |
| | | Incorrect | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% |
| | | No detection | 35.04% | 37.21% | 19.85% | 37.04% | 38.18% | 20.02% |
| | 100.00% | Correct | 100.00% | 100.00% | 100.00% | 100.00% | 100.00% | 100.00% |
| | | Incorrect | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% |
| | | No detection | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% |
| ML-based | 95% | Correct | 97.54% | 95.99% | 93.51% | 96.91% | 98.18% | 93.36% |
| | | Incorrect | 1.41% | 3.05% | 0.00% | 1.85% | 0.91% | 0.00% |
| | | No detection | 1.05% | 0.96% | 6.49% | 1.24% | 0.91% | 6.64% |
| | 90% | Correct | 94.01% | 91.41% | 88.30% | 95.06% | 93.64% | 88.74% |
| | | Incorrect | 3.17% | 5.92% | 0.00% | 2.47% | 5.45% | 0.00% |
| | | No detection | 2.82% | 2.67% | 11.70% | 2.47% | 0.91% | 11.26% |
| | 85% | Correct | 89.44% | 88.93% | 84.99% | 91.88% | 86.36% | 84.83% |
| | | Incorrect | 4.75% | 5.15% | 0.00% | 2.46% | 10.00% | 0.00% |
| | | No detection | 5.81% | 5.92% | 15.01% | 5.66% | 3.64% | 15.17% |
| | 80% | Correct | 85.56% | 81.49% | 80.15% | 81.48% | 79.09% | 79.98% |
| | | Incorrect | 5.28% | 7.63% | 0.00% | 3.09% | 8.18% | 0.00% |
| | | No detection | 9.16% | 10.88% | 19.85% | 15.43% | 12.73% | 20.02% |

does not attempt to infer missing data and, therefore, its incorrect failure localization frequency remains at zero.

The results show that the trends observed in Figs. 6, 7, and 8 are also preserved for higher percentages of missing telemetry. At 80% telemetry, some differences between transponder failure localization, and amplifier and fiber link failure localization, become more evident. In transponder failure localization, the di-

rect and ML-based schemes exhibit similar performance. This is because, in transponder failure localization, there is less information from other devices for the ML-based scheme to localize the failure. As shown in Fig. 8, the ML-based strategy only improves its performance for a higher degradation (e.g., higher than 10 dB), when the WSS/OCM control loops lose the ability to equalize the signal power and the transponder power degradation can

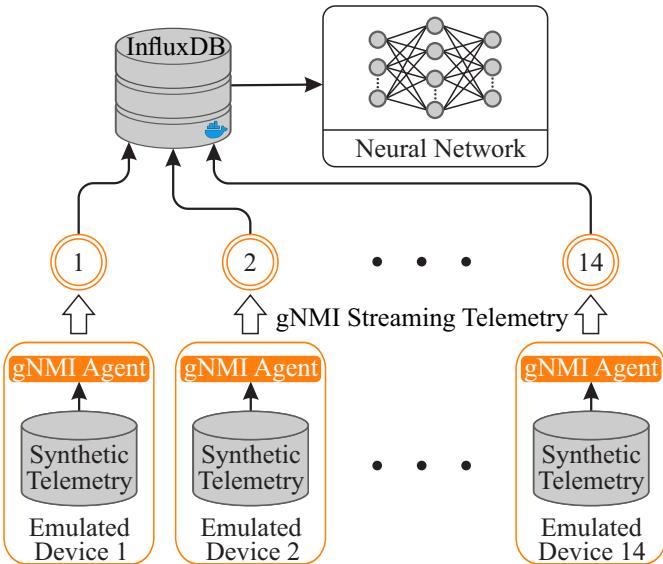


Fig. 9. Emulation setup composed of emulated devices. Each emulated device offers a gNMI agent that is able to provide streaming telemetry service to gNMI collectors. Each collector stores the received telemetry in an InfluxDB time series database. An ANN reads from InfluxDB and is able to localize the soft-failure.

be sensed in other network devices (R3.-5#). For amplifier and fiber link localization, the ML-based scheme exhibits a considerably higher correct localization frequency ($\approx 80\%$ for ML-based *versus* $\approx 60\%$ for the direct scheme). Additionally, we observed in all cases of incorrect localization (3% to 8%) that the actual failure is just one device apart. Therefore, the ML-based scheme can also be examined as an algorithm that localizes the vicinity of a true failure.

5. EMULATION EXPERIMENTS

We validate the consistency between telemetry and failure localization using synthetic data generated by emulated NSFNet devices. Fig. 9 illustrates the emulation setup, corresponding to the 14 NSFNet nodes. Each emulated device comprises a synthetic telemetry database and one gNMI Agent (server). The synthetic data stored in the telemetry database is generated using the signal and noise power propagation models depicted in Fig. 2 (R3.-5#) (R3.-6#). For each emulated device, a corresponding collector is used to subscribe to telemetry updates using gNMI streaming telemetry. When a collector receives new data, it forwards it to an InfluxDB time-series database hosted in a docker container. Another process runs an ML algorithm (PCA+ANN solution presented in Section 2) that continuously reads the telemetry from InfluxDB. We emulate a gain degradation in the 5th amplifier in the link interconnecting NSFNet nodes 1 and 2 (Amp_1_2_5). The nominal gain value is 16 dB. To generate the synthetic telemetry, a simulator is configured to generate datasets corresponding to Amp_1_2_5 gain degradations ranging from 0 dB (normal operation) to 4 dB in steps of 0.5 dB.

Fig. 10 presents the emulation results. Telemetry latency issues are neglected. The blue curve with crosses shows the relative time in which the telemetry data in Amp_1_2_5 gain degradation is generated at all NSFNet nodes. The green line

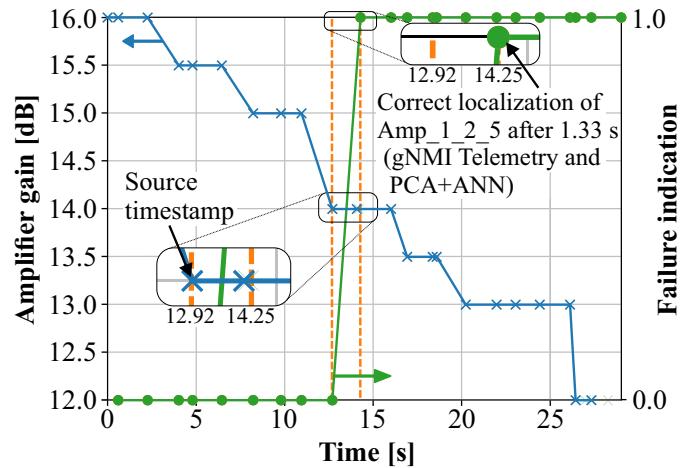


Fig. 10. Emulation result of a failure in Amp_1_2_5. Telemetry data is computer-emulated, transmitted by the gNMI protocol, and collected by an InfluxDB database. The database feeds the PCA+ANN to localize the failure. The pipeline between failure emulation and PCA+ANN detection takes 1.33 s.

with circles shows the output of the PCA+ANN after streaming telemetry and database processing. The PCA+ANN failure localization is triggered when the amplifier gain reaches 14 dB. The entire streaming telemetry, database processing, and ML algorithm processing are executed between instants 12.92 and 14.25, taking in total 1.33 second.

6. CONCLUSIONS

We evaluate an ML-based framework to localize soft-failures in optical networks in scenarios of partial telemetry. The ML-based framework uses an ANN trained with synthetic telemetry data simulating the occurrence of hard- and soft-failures. PCA is used to accelerate the training process. We show that the ML-based framework interpolates missing data, correctly localizing the failure in partial telemetry scenarios. In the rare events where the ML-based algorithm localizes an incorrect failure, the localized failure is in the vicinity of the actual failure. The proposed ML-based algorithm exhibits a suitable real-time performance in an emulated large-scale network scenario using state-of-the-art streaming telemetry.

ACKNOWLEDGMENTS

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.

Jonathan A. Soares is supported in part by the National Council for Scientific and Technological Development – Brasil (CNPq) – Grant Number 132545/2019-5.

This is an extended version of [20].

REFERENCES

- D. Kreutz, F. M. V. Ramos, P. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," Proc. IEEE **103**, 14–76 (2015).

2. F. Paolucci and A. Sgambelluri, "Telemetry in disaggregated optical networks," in *2020 International Conference on Optical Network Design and Modeling (ONDM)*, (2020), pp. 1–3.
3. D. Rafique, T. Szyrkowiec, H. Grießer, A. Autenrieth, and J. Elbers, "Cognitive assurance architecture for optical network fault management," *J. Light. Technol.* **36**, 1443–1450 (2018).
4. 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. Light. Technol.* **37**, 4125–4139 (2019).
5. F. N. Khan, Q. Fan, C. Lu, and A. P. T. Lau, "An optical communication's perspective on machine learning and its applications," *J. Light. Technol.* **37**, 493–516 (2019).
6. 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 *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, (2020), pp. 1–3.
7. S. Barzegar *et al.*, "Soft-failure localization and device working parameters estimation in disaggregated scenarios," in *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, (2020), pp. 1–3.
8. Z. Wang *et al.*, "Failure prediction using machine learning and time series in optical network," *Opt. Express* **25**, 18553–18565 (2017).
9. A. P. Vela *et al.*, "Soft failure localization during commissioning testing and lightpath operation," *J. Opt. Commun. Netw.* **10**, A27–A36 (2018).
10. S. Shahkarami, F. Musumeci, F. Cugini, and M. Tornatore, "Machine-learning-based soft-failure detection and identification in optical networks," in *2018 Optical Fiber Communications Conference and Exposition (OFC)*, (2018), pp. 1–3.
11. H. Lun *et al.*, "Soft failure identification for long-haul optical communication systems based on one-dimensional convolutional neural network," *J. Light. Technol.* **38**, 2992–2999 (2020).
12. S. Varughese, D. Lippatti, T. Richter, S. Tibuleac, and S. E. Ralph, "Identification of soft failures in optical links using low complexity anomaly detection," in *2019 Optical Fiber Communications Conference and Exhibition (OFC)*, (2019), pp. 1–3.
13. 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. Light. Technol.* **38**, 2669–2679 (2020).
14. T. Panayiotou, S. P. Chatzis, and G. Ellinas, "Leveraging statistical machine learning to address failure localization in optical networks," *J. Opt. Commun. Netw.* **10**, 162–173 (2018).
15. S. M. Srinivasan, T. Truong-Huu, and M. Gurusamy, "Machine learning-based link fault identification and localization in complex networks," *IEEE Internet Things J.* **6**, 6556–6566 (2019).
16. Z. Li *et al.*, "Demonstration of alarm knowledge graph construction for fault localization on ONOS-based SDON platform," in *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, (2020), pp. 1–3.
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. Serv. Manag.* **Early Access** (2021).
18. T. Pan, E. Song, C. Jia, W. Cao, T. Huang, and B. Liu, "Lightweight network-wide telemetry without explicitly using probe packets," in *IEEE INFOCOM 2020*, (2020), pp. 1354–1355.
19. A. Sadasivarao *et al.*, "Demonstration of extensible threshold-based streaming telemetry for open DWDM analytics and verification," in *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, (2020), pp. 1–3.
20. K. S. Mayer, J. A. Soares, R. P. Pinto, C. E. Rothenberg, D. S. Arantes, and D. A. A. Mello, "Soft failure localization using machine learning with SDN-based network-wide telemetry," in *46th European Conference on Optical Communication (ECOC 2020)*, (2020), pp. 1–4.
21. F. Paolucci, A. Sgambelluri, F. Cugini, and P. Castoldi, "Network telemetry streaming services in SDN-based disaggregated optical networks," *J. Light. Technol.* **36**, 3142–3149 (2018).
22. D. Wang *et al.*, "The role of digital twin in optical communication: Fault management, hardware configuration, and transmission simulation," *IEEE Commun. Mag.* **59**, 133–139 (2021).
23. Y. Zhao *et al.*, "Static resource allocation of advanced reservation requests in elastic optical networks," *Appl. Opt.* **59**, 1420–1429 (2020).
24. A. Betker *et al.*, "Reference transport network scenarios," *MultiTeraNet Rep.* pp. 1–15 (2003).
25. A. C. Jatoba-Neto, D. A. A. Mello, C. E. Rothenberg, S. Ö. Arik, and J. M. Kahn, "Scaling SDM optical networks using full-spectrum spatial switching," *J. Opt. Commun. Netw.* **10**, 991–1004 (2018).
26. I. Katib and D. Medhi, "IP/MPLS-over-OTN-over-DWDM multilayer networks: an integrated three-layer capacity optimization model, a heuristic, and a study," *IEEE Trans. Netw. Serv. Manag.* **9**, 240–253 (2012).
27. Y. Xiao, J. Zhang, Z. Gao, and Y. Ji, "Service-oriented DU-CU placement using reinforcement learning in 5G/B5G converged wireless-optical networks," in *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, (2020), pp. 1–3.
28. M. Soltanolkotabi, A. Javanmard, and J. D. Lee, "Theoretical insights into the optimization landscape of over-parameterized shallow neural networks," *IEEE Trans. Inf. Theory* **65**, 742–769 (2019).
29. Y. He, X. Zhang, W. Ao, and J. Z. Huang, "Determining the optimal temperature parameter for Softmax function in reinforcement learning," *Appl. Soft Comput.* **70**, 80–85 (2018).
30. H. Lun *et al.*, "Soft failure identification in optical networks based on convolutional neural network," in *45th European Conference on Optical Communication (ECOC 2019)*, (2019), pp. 1–3.
31. 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.* **68**, 4204–4212 (2019).
32. A. Jain, K. Nandakumar, and A. Ross, "Score normalization in multimodal biometric systems," *Pattern Recognit.* **38**, 2270–2285 (2005).
33. D. P. Kingma and L. J. Ba, "Adam: a method for stochastic optimization," in *International Conference on Learning Representations (ICLR) 2015*, (2015), pp. 1–15.
34. S. Barzegar, M. Ruiz, and L. Velasco, "Soft-failure localization and time-dependent degradation detection for network diagnosis," in *22nd International Conference on Transparent Optical Networks (ICTON)*, (2020), pp. 1–4.
35. H. Lun *et al.*, "ROADM-induced anomaly localization and evaluation for optical links based on receiver DSP and ML," *J. Light. Technol.* **39**, 2696–2703 (2021).
36. D. A. A. Mello, D. A. Schupke, M. Scheffel, and H. Waldman, "Availability maps for connections in WDM optical networks," in *International Workshop on Design of Reliable Communication Networks, 2005 (DRCN 2005)*, (2005), pp. 1–8.
37. J. Lever, M. Krzywinski, and N. Altman, "Principal component analysis," *Nat. Methods* **14**, 641–642 (2017).
38. H. Zhao, J. Zheng, J. Xu, and W. Deng, "Fault diagnosis method based on principal component analysis and broad learning system," *IEEE Access* **7**, 99263–99272 (2019).