Confidential Detection of Multiple Failures in Optical Networks: an Experimental Evaluation

M. F. Silva[1], A. Sgambelluri[2], A. Pacini[2,*], F. Paolucci[3], A. Green[1], D. Mascarenas[1], and L. Valcarenghi[2]

[1] *Los Alamos National Laboratory, New Mexico, USA*
[2] *Scuola Superiore Sant'Anna, Pisa, Italy*
[3] *CNIT, Pisa, Italy*

[*] *alessandro.pacini@santannapisa.it*

Abstract: This paper presents a Machine Learning technique based on Principal Component Analysis (PCA) combined with telemetry data scrambling to detect multiple types of failure in optical networks while preserving data confidentiality. Experiments in an optical testbed show the effectiveness of the proposed solution. © 2022 The Author(s)

1. Introduction

Preserving data confidentiality is becoming of paramount importance in many fields (e.g., confidential computing [1] and in optical networks). With emerging trend of network automation or zero touch networking some functions, such as failure detection, can be offloaded to third party applications based on Artificial Intelligence/Machine Learning. In this case, network providers might not want to disclose confidential data related to their network state but they would like that third party applications allow them to effectively detect network failures.

In [2] the authors proposed a method for maintaining confidentiality of telemetry data to be elaborated by a third party app. The solution stems from methods already proposed in image processing [3, 4] where, although pixels are scrambled, Machine Learning (ML) techniques are capable of classifying the original image. However, the evaluation reported in [2] showed the effectiveness of the proposed technique considering a single specific failure only and a specific set of features, such as the Bit Error Rate (BER) and the Optical Signal to Noise Ratio (OSNR) at the receiver. This paper experimentally that the aforementioned method is effective in preserving the data confidentiality and detecting multiple failures along a WDM network testbed shared by multiple lightpaths.

2. Feature scrambling-based confidentiality and PCA-based failure detection

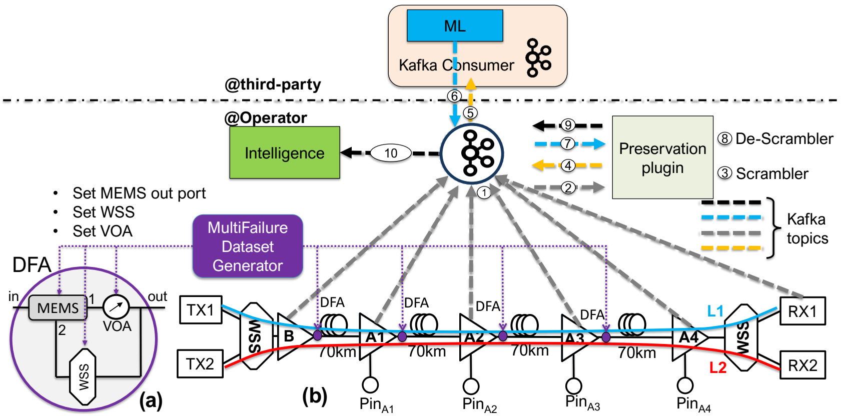

Fig. 1. a) Distributed Failure Actuator (DFA) internal architecture and configurability; b) Experimental optical testbed and telemetry system.

As originally presented in [2] the proposed method for confidential failure detection combines feature scrambling, for preserving confidentiality, and Principal Component Analysis (PCA), for failure detection. Considering Fig.1, data are collected from the optical network and sent to the preservation plugin (step 1-2). Data confidentiality is obtained by scrambling, at the operator premises, the elements of the matrix $\mathbf{X} \in \mathbb{R}^{n \times m}$ consisting of m telemetry parameters collected n times from different network devices (step 3).

The scrambled feature matrix $\tilde{\mathbf{X}}$ is then sent to the third-party app (step 4-5) that finds its principal components by means of Principal Components analysis (PCA) (see [5] for details). Because PCA is rotational invariant, the ordering of the measurement instances has no effect on the resulting calculation of the principle components and principle directions. Thus, applying PCA to $\mathbf{X}$ or $\tilde{\mathbf{X}}$ results in the same manifold space but with different rotations.

The PCA reconstructed scrambled matrix $\tilde{\mathbf{X}}_R$ is then sent back to the operator (step 6-7), where the preservation plugin de-scrambles the data (step 8) and sends them to the Intelligence block (step 9-10). The intelligence computes the error $\mathbf{E} = \mathbf{X} - \tilde{\mathbf{X}}_R$, in the form of the Euclidean distance. If the error is higher than a given threshold, during the training, a failure is detected. The novelty of this paper is that the method is successfully applied for the first time to a scenario where failures can occur in different links and the path is shared by multiple lightpaths.

3. Experimental Evaluation Scenario

In the bottom of Fig.1 the testbed considered for the experimental assessment of the proposed solution is shown. It consists of a multispan link (i.e., 4 spans) with a single-mode-fiber spool of 70km each. After each span one amplifier (A_i) compensates the loss experienced by the optical signal (around 15dB). The amplifiers work in gain mode, applying the configured gain to the signal. The link is used to transmit two lightpaths, $L1$ and $L2$, generated by two commercial 100G muxponders. The channels are aggregated with a Wavelength Selective Switch (WSS) before being amplified by the booster amplifier (i.e., B). The booster amplifier is configured in power mode in a way that each channel at the ingress of span1 has a launch power of 0dBm, avoiding non-linear effects during the transmission. After the 4 spans link, the receiver side is present (right side of the figure). The two channels are filtered with a second WSS and are sent to the respective coherent receiver (i.e., RX1 and RX2), where the Digital Signal Processing (DSP) is performed and the channel performance is analyzed. At each span, a Distributed Failure Actuator (DFA) has been connected to emulate failures. The DFA has been implemented with a Micro Electro-Mechanical Systems (MEMS), able to switch the optical transmission from an input port to two possible output ports, and a WSS to allow different failure emulation (i.e., incremental loss, narrow filtering). As shown in the insight 1.a the two available physical paths present the same attenuation, keeping in consideration the insertion loss of MEMS (e.g., around 3dB) and WSS (e.g., around 8dB). By relying on the DFAs and the MultiFailure Dataset Generator, a failure can be automatically generated in one of the different spans, allowing the implementation of different scenarios in a single testbed (i.e., from a 280km multi-span link to a network with 5 ROADMs connected by 4 single-span links). The Generator is able to configure, for each DFA, failure/non failure condition, along with different types of failure. In the experimental validation three different types of failure have been considered: (i) the generalized degradation affecting the two optical channels (i.e., 10dB), (ii) the degradation of only $L1$, (iii) the degradation of only $L2$. The described testbed is connected to a new-generation telemetry-enabled monitoring platform [6], able to collect the Key Performance Indicators (KPIs) from the traversed devices. More specifically, the system samples the input power level at each amplifier (e.g., Pin_{Ai} in the figure, with i=1,2,3,4) and the coherent data at receiver side (i.e., OSNR and instantaneous pre-FEC Bit Error rate) for each installed lightpath, with a period of 5 seconds. The platform consists of a distributed event streaming system based on $Kafka$, with high-performance data pipeline and data processing. In particular, the platform allows the streaming of data from $KafkaProducers$ to $KafkaConsumers$ using $Topics$ (see Fig. 1). This type of architecture is able to perform on-the-fly data processing in the form of plugin attached to the system by using standalone $Consumer-Processor-Producer$ nodes. In this specific scenario, the plugin concept has been applied to implement the scrambling (step 3 in the figure) and the de-scrambling (step 8) operations to preserve the data confidentiality between operator and third-party app. To test the PCA-based algorithm 4 different datasets have been collected, considering a failure in each of the 4 spans. Each failure has duration of 1 minute, with random inter-failure time between 1 and 4 minutes. Each experiment considers four main phases: (1) the normal condition, (2) five failures on $L1$ only, (3) five failures on $L2$ only, (4) failures on both $L1$ and $L2$.

4. Results

All the experiments are performed on a Ubuntu PC equipped with a CPU Intel$^{®}$ Core^{TM} i7-8750H, and 32 GB of RAM. Fig. 2 introduces the performance of the proposed approach in terms of data reconstruction (Fig. 2a) and failure detection (Fig. 2b) along with the confusion matrix for the original dataset (Fig. 2c) and the scrambled one (Fig. 2d). In the training ($Train$) and validation (Val) phase (depicted in black and gray colors) a specific dataset is considered that contains the key parameters in normal conditions, with no failures; 80% is used for training while 20% is used for validation. It can be verified that the model adequately reproduces the normal pattern of the system, which is corroborated by the small values of the mean squared error for the training data in Fig. 2a.

For the test data, Fig. 2b reports the outlier detection results, when the aforementioned failure sequence is generated respectively in span2 (S2 in blue) and in span3 (S3 in green). In the former case, outliers are found for the input power at ampli2, ampli3, ampli4 and the OSNR collected at two cards. Indeed, according to the failures sequence, five failures affect $L1$ (i.e., card1 OSNR), then five failures affect $L2$ (i.e., card2 OSNR) and finally five failures are generated to degrade both $L1$ and $L2$. Because the failures are generated after span1, span1 is not affected by the failure (i.e., very few outliers are detected). In the S3 test, outliers are found for the input power at

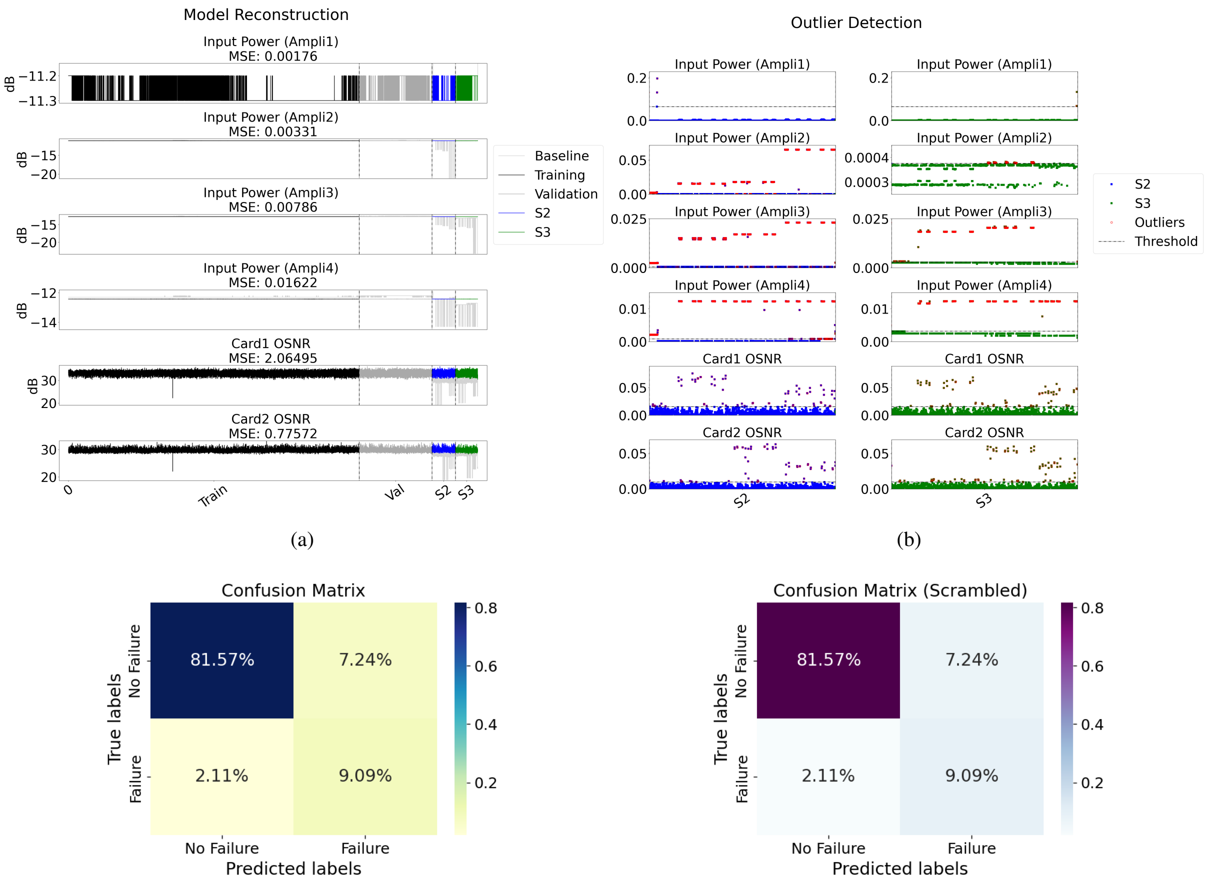

(a)

(b)

(c)

(d)

Fig. 2. Model reconstruction (a), detection (b), original (c) and scrambled confusion matrix (d).

ampli3, ampli4 and the OSNR collected at two cards. Since the same failure sequence is generated, similar outlier detection pattern can be observed, with the difference that since the failures are generated after span2, span1 and span2 do not present outliers (i.e., very few outliers are detected).

To check the correspondence between a detected soft failure and the actual condition of the system the confusion matrix is used. For simplicity, Fig. 2c and Fig. 2d show the confusion matrix when the data is not and is scrambled respectively. After verifying the values at the main diagonal, for both cases, the model performs accurate predictions of failure and no failure conditions with 90.66% of accuracy in the predictions.

5. Conclusions

A technique based on PCA and telemetry data scrambling is adopted to confidentially detect multiple failures in Optical Networks. The experimental results collected on an optical testbed show the effectiveness of the solution.

Acknowledgment. Funding received from the ECSEL JU project BRAINE (grant agreement No 876967). The JU receives support from the EU Horizon 2020 research and innovation programme and the Italian Ministry of University, and Research (MUR).

References

1. F. Y. Rashid, “The rise of confidential computing: Big tech companies are adopting a new security model to protect data while it’s in use - [news],” IEEE Spectr. **57**, 8–9 (2020).
2. M. Felipe Silva, A. Pacini, A. Sgambelluri, F. Paolucci, and V. Valcarenghi, “Confidentiality-preserving machine learning scheme to detect soft-failures in optical communication networks,” in *ECOC 2022,* (2022), pp. 1–4.
3. M. Tanaka, “Learnable image encryption,” in *2018 IEEE ICCE-TW Conference,* (2018), pp. 1–2.
4. B. Martinez, A. Green, M. F. Silva, Y. Yang, and D. Mascareñas, “Sparse and random sampling techniques for high-resolution, full-field, bss-based structural dynamics identification from video,” Sensors **20** (2020).
5. F. Paolucci, A. Sgambelluri, M. Felipe Silva, A. Pacini, P. Castoldi, L. Valcarenghi, and F. Cugini, “Peer-to-peer disaggregated telemetry for autonomic machine-learning-driven transceiver operation,” J. Opt. Commun. Netw. **14** (2022).
6. A. Sgambelluri, A. Pacini, F. Paolucci, P. Castoldi, and L. Valcarenghi, “Reliable and scalable kafka-based framework for optical network telemetry,” J. Opt. Commun. Netw. **13**, E42–E52 (2021).