Service-triggered Failure Identification/Localization Through Monitoring of Multiple Parameters

M. Ruiz[1*], F. Fresi[2], A. P. Vela[1], G. Meloni[2], N. Sambo[3], F. Cugini[2], L. Poti[2], L. Velasco[1], P. Castoldi[3]

(1) Universitat Politècnica de Catalunya (UPC), Barcelona (Spain), mruiz@ac.upc.edu
(2) CNIT, Pisa, (Italy)
(3) Scuola Superiore Sant'Anna, Pisa, (Italy)

Abstract Failures in the optical layer might impact the quality of supported services. We experimentally characterize several failure causes and propose an effective machine learning-based algorithm to localize and identify the most probable cause of failure impacting a given service.

Introduction

Service layer connections are usually set up on virtual network topologies, where virtual links (*vlinks*) are supported by lightpaths in the optical layer. Thus, errors in lightpaths translate in errors in those connections that might cause packet losses and retransmissions and lead to unacceptable Quality of Service (QoS). Physical layer monitoring is key to verify the fulfilment of service level agreements (SLA)[1] and, in case of faults or degradations, to localize the failed elements[2,3] and to take actions for preserving the services. Information retrieved by commonly used power monitors can be combined with monitoring information accessible through emerging transponders based on coherent detection[4]. In particular, such transponders offer the possibility to monitor several parameters associated to connections or to the traversed links: e.g. pre-forward error correction bit error rate (pre-FEC BER) or linear dispersion.

In this paper, we study the effects on Quality of Transmission (QoT) monitoring parameters of several failures on the optical layer, specifically those of *tight filtering* and *inter-channel interference*; collected QoT monitoring parameters include received power (P_{Rx}) and pre-FEC BER. We propose an algorithm that analyses monitoring time series and, based on the expected patterns for the considered failure causes obtained in our experiments, localizes and identifies the most probable cause of failure at the optical layer affecting a given service.

Failure Identification/Localization

For illustrative purposes, Fig. 1 presents monitoring data series for the possible causes of failure affecting a given optical connection. Tight filtering happens when a too much narrow filter configuration distorts the signal. Such effect may become even more relevant when the signal drifts (e.g. due to a laser drift) toward the rising edge of the filter[5].

As it will be shown in the experimental demonstration, in case of laser drift, P_{Rx} decreases because of the filtered power and BER is degraded; in fact periods with degraded BER are followed by others with normal BER, which makes difficult to localize the failure cause. In fact, BER degradation is not always caused by P_{Rx} decrease, as shown in the inter-channel interference example, where the allocation of a neighboring lightpath results in a sudden increment observed in the target lightpath. Hence, failure localization entails deep analysis of monitoring data from several lightpaths.

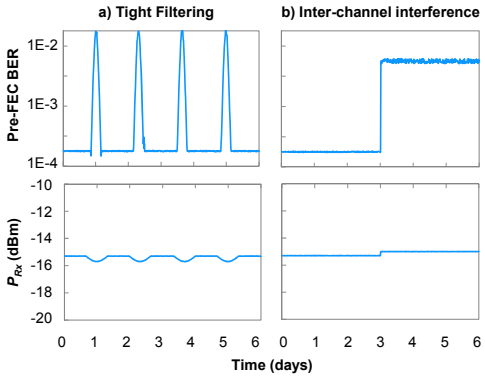

Fig. 1. BER and P_{Rx} monitoring time series for the considered failures.

Proposed Algorithm

With the above in mind, we propose a probabilistic failure localization algorithm based on Bayesian Networks (BN)[6]. A BN is a directed acyclic graph where nodes represent features and edges the conditional dependency between a pair of features. Each node is associated to a probability function that takes values from the parent nodes and returns the probability of the feature represented by the node.

The proposed BN is trained to locate different causes of failures and returns its probability. Before training, several experimental tests for each of the possible causes of failure, as well as for the no failure case, need to be carried out to obtain monitoring data series similar to the ones in Fig. 1. Then, those data series are transformed into relevant *descriptive features* collecting the main characteristics of data series, such as minimum, maximum, average, trend,

Table 1. Failure Localization Algorithm

INPUT: *s, BN*
OUT: *A*

1: *P*←getLightpaths(*s*)
2: *A*←Ø
3: **for each** *p* in *P* **do**
4: *H*←getMonitorDataSeries(*p*)
5: *F*←*BN*.computeFeatures(*H*)
6: *F'*←*BN*.discretize(*F*)
7: *D*←*BN*.predict(*F'*)
8: *A.p*←sortProblemList(*D*)
9: **return** *A*

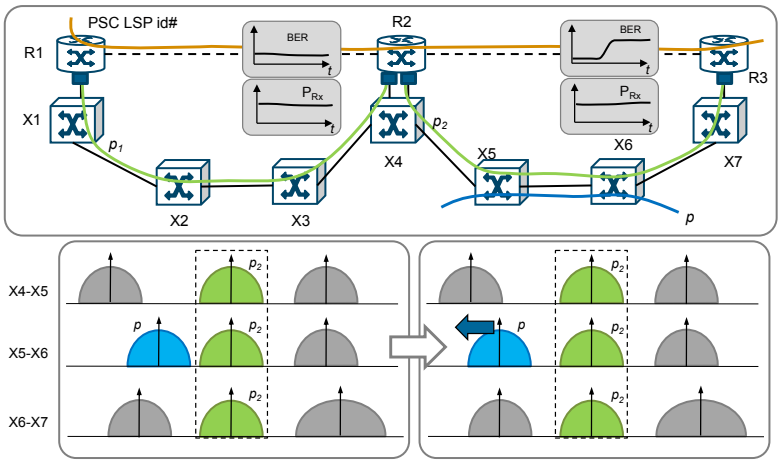

Fig. 2. Example of failure localization caused by inter-channel interference.

stepped change presence and size, etc. Since BNs require categorical features (i.e. with a finite range of levels), continuous features can be easily discretized by applying a clustering algorithm to find the number and ranges of each of the levels. The type of failure is also added as *response feature*. The proposed algorithm that integrates the BN (Table 1) receives as input the affected service connection *s* and the previously trained BN. After retrieving the set *P* of lightpaths supporting *s* from the operational database (line 1 in Table 1), every single lightpath is sequentially processed as follows: first, the available BER and P_{Rx} monitoring data series are retrieved from the monitoring repository in the form of variable-length time series of continuous data (line 4). Then, continuous features are computed and transformed into categorical values (lines 5-6). The prediction *D* returns, for each of the failures, the probability that it actually occurs (line 7). Such probabilities allow sorting the list of failures for every single lightpath, which is returned (lines 8-9).

Fig. 2 shows an example of the proposed failure localization algorithm. Let us imagine that a service using a connection between R1 and R3 has detected and notified service degradation to the service provider. The monitoring data of the two lightpaths supporting the service connection are analyzed. As observed, *p1* monitoring data series show an almost constant trend for both power and BER, whereas *p2* BER suffered a steeped increase at some point in the past. From this available data, the failure localization algorithm returns no failure with probability 95% for *p1* and identifies interference with 70% and tight filtering with 25% for *p2*.

According to the probabilities above, the scope of network reconfiguration is firstly focused on *p2* and those lightpaths sharing an optical link in the route of *p2*. A deeper analysis identifies that *p2* BER increment is correlated with lightpath *p* set-up. Therefore, by slightly shifting *p* in the spectrum away from *p2*, its BER should be improved and the detected service degradation eventually reduced. However, a monitoring period after reconfiguration is needed to verify that BER degradation has been completely solved. In the case that *p2* BER has not reduced to normal values after *p2* shifting, the second action in the list is taken, which might consists in making filters wider to overcome the probable tight filtering failure. With this second reconfiguration step, the service degradation should be finally solved.

Experimental demonstration

We demonstrated the effectiveness of the proposed algorithm by exploiting monitoring information retrieved by the experimental testbed shown in Fig. 3. A Nyquist wavelength division multiplexing (NWDM) signal is assumed. In particular a digital-to-analog converter (DAC) is used to periodically output pulse shaped electrical signals which drove the Mach Zehnder based IQ-modulators. A root raised cosine (RRC) with a roll-off of 0.2 and a bandwidth of 15 GHz is used to confine signal bandwidth. Two single polarization IQ-modulators are used to modulate two external cavity lasers (ECL) and generate two quadrature phase-shift keying (QPSK) at a gross baud rate of 30 Gbaud (i.e. 60 Gb/s gross bit rate). Next, the bit rate is doubled by a polarization multiplex emulation stage, thus obtaining two 120 Gb/s polarization multiplexed (PM)-QPSK signals. The two modulated lasers are then multiplexed by means of a spectrum selective switch (SSS) configured to reserve a 37.5 GHz frequency slot for each channel. The signals are then transmitted through 60 km of standard optical fiber and the desired channel is optically filtered, pre-amplified and received exploiting coherent detection. The bandwidth of the optical filter at the receiver is set to 25 GHz (which correspond to around ten switch&select nodes, according to

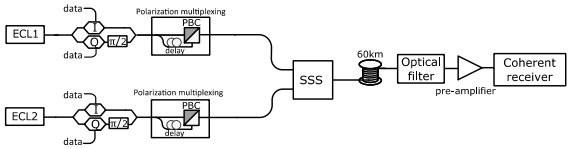

Fig. 3. Data plane experimental testbed

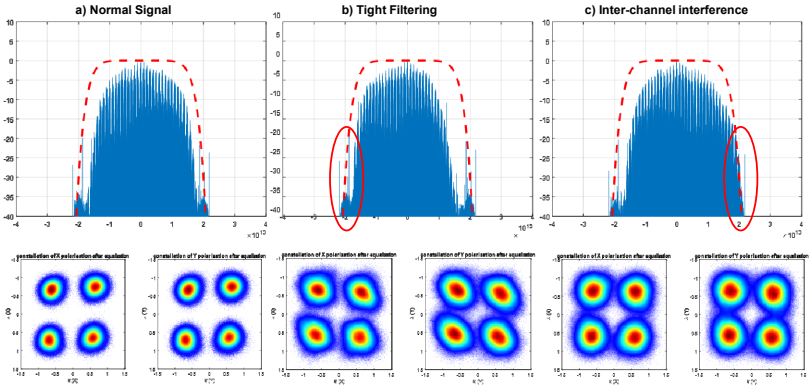

Fig. 4. Experimental results for the normal conditions and considered failures

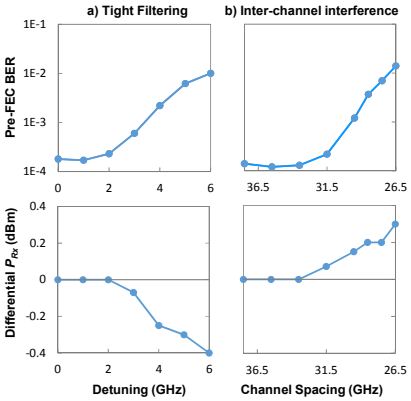

Fig. 5. Experimental BER and P_{Rx}

simulations we performed) in order to emulate the filtering effect of multiple 37.5 GHz filters.

Two experiments are carried out. Measurements are reported for *signal 1*. First, filtering effects are assumed upon *signal 1* drift. Second, inter-channel interference is induced, assuming the laser drift of *signal 2*. This way, the channel spacing among the signals decreases inducing an increase of interference. Spectrum related to *signal 1* is reported in Fig. 4 for both experiments. Fig. 4a shows signal 1 spectrum under normal conditions. In Fig. 4b, such signal experiences a slight shift in frequency due to the laser drift, while in Fig. 4c part of *signal 2* falls within the bandwidth of *signal 1*.

Fig. 5 reports BER and differential P_{Rx} measurements for both experiments (i.e. the difference between the received power and the received power in normal conditions). It is worth noting the differences between these plots and those in Fig. 1, where historical time series are plotted. Considering inter-channel interference, BER increases when the channel spacing between *signal 1* and *2* decreases, while P_{Rx} increases when channel spacing decreases. This is due to the fact that part of signal 2 enters in the *signal 1* bandwidth.

Assuming tight filtering, a similar behaviour on the BER is experienced. Indeed, BER increases with the laser drift since the impact of filtering effects becomes more relevant. The behaviour of the power with respect to the inter-channel interference case is different. Indeed, received power decreases while the laser drift increases since part of the power is cut by the filter. The different behaviours of P_{Rx} may drive the decision to discern between inter-channel interference and excessive filtering effects because laser drift. This is considered by the proposed decision algorithm.

According to the experimental values in Fig. 5, we generated synthetic monitoring time series for the normal signal and the considered failure cases. A set of 5,000 randomly generated time series were first used to train the BN and next,

Table 2. BN Goodness-of-Fit

		Real		
		Normal	Filtering	Interference
Pre-diction	Normal	**99.2%**	0.08%	0%
	Filtering	0%	**100%**	0%
	Interference	0%	0%	**100%**

500 additional ones used for testing. Table 2 reports the obtained goodness-of-fit computed as the probability that the BN predicts the actual failure cause as the first option. Note that only 0.8% error was observed in some tests where a normal signal was predicted instead of a tight filtering failure. In such cases, the second most probable cause of failure was tight filtering failure. This demonstrates the validity of the proposed procedure and BN to localize and identify failures in the optical layer.

Conclusions

Two different failure causes at the optical layer have been experimentally characterized and the obtained measurements used to generate time series to train a BN. When a service detects excessive errors, an algorithm uses the trained BN to localize and identify the most probable cause of the errors at the optical layer. Results showed the effectiveness of the algorithm.

Acknowledgements

This work was partially supported by the EC through the ORCHESTRA project (grant agreement 645360), from the Spanish MINECO SYNERGY project (TEC2014-59995-R) and from the Catalan Institution for Research and Advanced Studies (ICREA).

References

1 D. King et al., “A PCE-based architecture for application-based network operations,” IETF RFC 7491, 2015.
2 J. Tapolcai et al., "Neighborhood Failure Localization in All-Optical Networks via Monitoring Trails," in IEEE/ACM Transactions on Networking, vol. 23, 2015.
3 K. Christodoulopoulos et al. “Exploiting Network Kriging for Fault Localization," OFC 2016.
4 A. Napoli et al., "Next generation elastic optical networks: The vision of the European research project IDEALIST," IEEE Comm. Mag., vol. 53, 2015.
5 ORCHESTRA project. Deliverable D2.2, 2016.
6 Ch. Bishop, “Pattern Recognition and Machine Learning,” Springer-Verlag, 2006.