Anticipating BER Degradation in Optical Networks

Alba P. Vela, Marc Ruiz, and Luis Velasco *
Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
**e-mail: lvelasco@ac.upc.edu*

ABSTRACT
Optical connections support virtual links in MPLS-over-optical multilayer networks and therefore, errors in the optical layer impact on the quality of the services deployed on such networks. Monitoring the performance of the physical layer allows verifying the proper operation of optical connections, as well as detecting bit error rate (BER) degradations. Anticipating BER degradation facilitates self-decision making to keep committed service level. In this paper, we analyze several failure causes affecting the quality of optical connections and propose the BANDO algorithm focused on detecting significant BER changes in optical connections. BANDO runs inside the network nodes to accelerate degradation detection and sends a notification to the centralized controller. Results show significant improvement anticipating maximum BER violation.
Keywords: BER degradation, elastic optical networks, monitoring and data analytics.

1. INTRODUCTION

Service layer connections are usually set up on virtual network topologies, where virtual links are supported by lightpaths in the optical layer. Thus, errors in optical transmission translate into packet losses and retransmissions leading to unacceptable Quality of Service (QoS) and to Service Level Agreements (SLA) violations. Such violations represent money losses for the network operator. Although commercially available optical equipment is able to correct degraded optical signals by means of Forward Error Correction (FEC) algorithms, a value of pre-FEC Bit Error Rate (BER) over the pre-defined limit (max BER) would imply a non-error-free post-FEC transmission and, as a result, communication would be disrupted. Therefore, a prompt detection of optical connections with excessive pre-FEC BER can greatly reduce SLA violations.

Monitoring the physical layer is essential to verify the fulfillment of SLAs and, in the case of faults or degradations, e.g. transmitter laser drift or filters misconfiguration, to localize the failed elements, and to take actions for preserving the services. Information retrieved by commonly used power monitors measuring received power (P_{Rx}) can be combined with monitoring information accessible through emerging transponders based on coherent detection [1]. In particular, such transponders offer the possibility to monitor several parameters associated with connections or to the traversed links, e.g., pre-FEC BER or linear dispersion.

Quality of transmission estimation, computed as a function of the links and nodes traversed by each optical connection, is a useful information since it can be used to configure a BER threshold at connection set-up, which would help to detect BER degradation by comparing the actual measured BER against it. However, if the threshold value is set to a value too close to the actual BER, many threshold-crossing notifications would be raised because of small BER changes, which, in addition to add control overhead, do not give useful information. On the contrary, if the threshold value is relaxed, e.g., closed to the equipment max BER, degradation detection could not be anticipated early enough the transmission is totally disrupted.

Regarding failure localization, several works in the literature have proposed methods for localization of hard link failures that affect a number of established connections, focused on reducing restoration times. However, few works have been focused on soft-failure localization that might affect a single or a reduced set of optical connections. For this very reason, in this paper, we focus on BER degradation detection and report the main results from our previous paper [2]. First, we analyze four different failures affecting the signal quality of an optical connection and motivate the definition of the BER Anomaly Detection (BANDO) algorithm focused on detecting significant BER changes in optical connections. Next, the BANDO algorithm is defined in detail as a finite state machine (*fsm*) to follow the metered BER and to raise notifications in case of abrupt BER changes. Finally, based on measured values, realistic scenarios are generated, and exhaustive simulations are run, where obtained results show the performance of the proposed algorithm.

2. BER DEGRADATION DETECTION

Figure 1 illustrates four failures affecting the signal quality of an optical connection: a) *signal overlap* (Fig. 1a) happens when the spectrum allocation of an optical connection invades that of a neighboring one. This might be caused by the inaccuracy in the central frequency of the laser and/or the filters of one of the connections; b) *tight filtering* (Fig. 1b) appears when there exists a central frequency misalignment or a width inaccuracy in the filters along the route of an optical connection. Figure 2 presents four different causes of tight filtering, where filter F2 is misaligned in Fig. 2a, filter F2 width is narrower than the required frequency slot width in Fig. 2b, filters F2 and F3 are misaligned in Fig. 2c, and the central frequency of the signal is misaligned in Fig. 2d. Note that the consequence on the optical signal is similar for all four cases; c) *gradual drift* (Fig. 1c) appears when either the optical signal or the filter gradually deviate from the central frequency determined at set-up time; and d) *cyclic drift* (Fig. 1d) occurs when a gradual drift describes a cyclical movement with time.

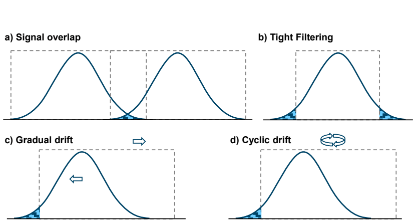

Figure 1. Four failures affecting the signal of an optical connection.

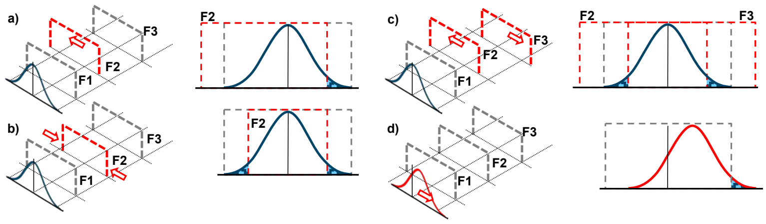

Figure 2. Causes of tight filtering.

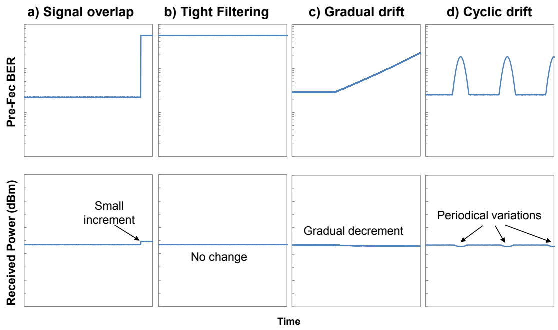

Figure 3. Example of pre-FEC BER and P_{Rx} monitoring time series for the considered BER degradation failures.

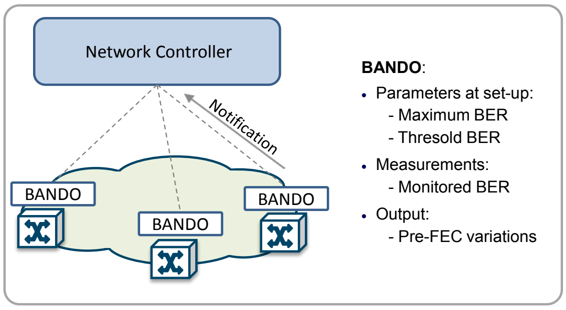

Figure 4. Proposed architecture and algorithm features.

For illustrative purposes, Fig. 3 plots the evolution with time of pre-FEC BER and P_{Rx} monitoring data metered at the receiver side of a connection affected by each of the failures above-described. In the case of signal overlap (Fig. 3a), the allocation of a neighboring optical connection results in a sudden increment in both, BER and P_{Rx}, of the previously established connection. In the case of the new connection, high pre-FEC BER and within limits P_{Rx} values can be measured just after its set-up. As for tight filtering (Fig. 3b), similarly as for newly established connection in the previous case, high pre-FEC BER and P_{Rx} values within limits can be measured in the receptor. In the case of gradual drift (Fig. 3c), pre-FEC BER shows a gradual deterioration with time, while measured P_{Rx} reduction is almost imperceptible. Finally, in the case of cyclic drift (Fig. 3d), high pre-FEC BER and slight P_{Rx} reduction periods when part of the signal is out filters' bandwidth are followed by normal values when the signal is inside them. Note that any combination of the previous failures might happen, e.g., a *gradual cyclic drift* would produce increasingly higher pre-FEC BER periods. It is hard to discern the real cause of the above failures since transmitter laser degradation, and filters misconfiguration could lead to similar evidence. In this paper, we concentrate in the prompt detection of pre-FEC BER degradation.

For the BER degradation detection, we propose the BANDO algorithm that can be placed inside network nodes, closer to the monitoring points, to reduce the amount of monitoring data to be conveyed to the control/management plane [3]. BANDO detects changes in the monitored BER measured in the receptor of an optical connection. Figure 4 presents the suggested architecture and placement to run BANDO algorithm inside the optical nodes, where it has access to fine-granular monitoring data in order to accelerate BER degradation detection. Once BER variation is detected, a notification is sent towards the controller for further analysis.

3. BER ANOMALY DETECTION (BANDO) ALGORITHM

We assume that metered pre-FEC BER and P_{Rx} data for every connection is received at a given rate (e.g., every minute) and stored in a vector M of fixed capacity n in the node. BANDO algorithm analyzes pre- FEC BER data to detect gradual changes with time that might derive into BER degradation and intolerable BER values, as well as sudden anomalous BER values.

Figure 5 illustrates three cases of BER evolution with time, where the dark continuous line represents monitored BER. Besides, two different limits are presented: *i*) *BER max* is the maximum pre-FEC BER that equipment can correct; and *ii*) a BER threshold for the current connection computed as a function of the estimated BER (e.g., 5 × estimated BER) and represents the maximum tolerable BER for such connection.

To follow BER evolution with time, an *outer boundary* is used to anticipate BER threshold violation and to detect sudden BER variations. In addition, two *inner boundaries*, named as a lower boundary (*lBound*) and upper boundary (*uBound*), are used to trigger boundary re-estimation when measured BER reaches, exceeds or falls below one of them. Inner and outer boundaries are estimated as $bound = \mu(M{\cdot}ber)$ +/- $k{\cdot}\sigma(M{\cdot}ber)$, where $\mu(M.ber)$ and $\sigma(M.ber)$ are the mean and the standard deviation computed on the last n BER measures and k is a multiplicative factor different per each boundary. Every time an event occurs, a notification is sent to the controller; defined events include: *i*) the *boundary* is re-estimated (*bCh*), *ii*) the boundary is exceeded (*bExc*), *iii*) BER exceeds the threshold (*thExc*) and *iv*) BER falls below the threshold (*thDec*).

Figure 5a presents an example of monitored BER evolution with time causing boundary changes. As soon as monitored BER crosses one of the inner bounds, a boundary re-estimation is triggered, and a notification is sent towards the controller. Figures 5b and 5c present two examples of sudden BER variation where the bound and the BER threshold is exceeded, respectively. In such cases, boundaries are reset, and notifications are sent to the controller. Note that, in case pre-FEC BER exceeds maximum BER, a special notification will be sent.

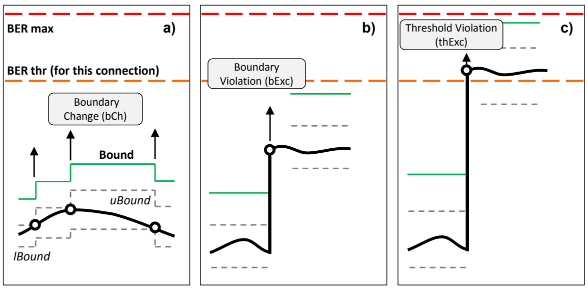

Figure 5. BER and boundaries evolution with time.

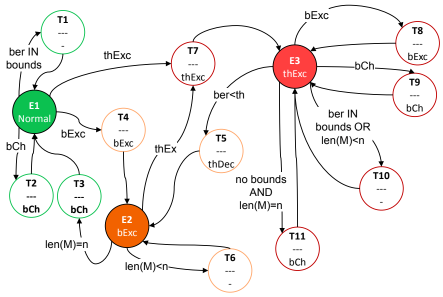

Figure 6. BANDO finite state machine.

The BANDO algorithm has been designed as an fsm with three main states and 11 transient states (Fig. 6); main states are used to store whether BER status is normal or has exceeded either the boundary or the threshold, whereas transient states are used to produce notifications and actions (i.e., boundaries re-estimation or reset). Every time a sample arrives, two fsm transitions are performed, one to obtain the output and action, and another to move to the new main state. State E1 (*normal BER*) is reached when the last BER value falls below the boundary and the threshold. Transitions to transient state T1 follow BER within boundaries, while transitions from transient states T2 and T3 re-estimate the boundaries (as in Fig. 5a). State E2 (*boundary exceeded*) is reached when the last BER value has exceeded the boundary, but it is still below the threshold (as in Fig. 5b). Transitions from transient states T4 and T5 reset boundaries, so n-1 new samples are needed to arrive to re-compute new boundaries. Finally, state E3 (*threshold exceeded*) is reached when the last monitored BER is above the threshold (as in Fig. 5c). Transitions from transient states T7 and T8 reset boundaries, whereas from transient states T9 and T11 re-estimate them.

4. RESULTS

First, experimental measurements to reproduce the BER degradation failures presented in Section 2 were carried out in CNIT laboratories in Pisa (Italy) aiming at retrieving data that will be used to generate synthetic data for the illustrative simulation results. Second, according to the experimental measurements, we generated synthetic monitoring time series at a rate of one sample per minute by means of a generator implemented in R. Each monitoring sample includes a synthetic measure of pre-FEC BER and P_{Rx}. The generator allows reproducing realistic monitoring activity of a set of optical connections with different characteristics, such as route, spectrum allocation, and slot width. Based on such characteristics and those of the underlying optical network topology, signal behavior in the absence of failures is generated. Besides, a per-connection BER threshold is computed based on an estimated BER value computed as a function of the links' OSNR in its route [4]. Besides, it allows reproducing any of the failures described in Section 2. According to the selected failure, one or more connections become affected at a given time, when some of their relevant physical properties are altered, e.g., filter bandwidth is narrowed; in the case of gradual changes, the magnitude of the alteration increases linearly with time following a predefined rate. Varying optical connection properties, failure class, failure magnitude, and gradual variation rate, we generated more than 100 distinct configurations. For each configuration, five 60-day instances (each generating 86,400 samples per optical connection) were randomly generated. Some of these configurations produced instances where BER never exceeded connection's BER threshold, whereas the rest contain at least one monitoring sample exceeding the connection's BER threshold.

The BANDO algorithm was implemented in R and integrated into a simulator following the architecture presented in Fig. 4. Aiming at finding the best configuration for BANDO parameters (to avoid an excessive number of notifications being sent to the controller while keeping it informed of meaningful BER changes), we set $n = 15$ and perform several tests with a wide range of k values for inner and outer boundaries; results are reported in Fig. 7 were values are normalized to those for the minimum k.

Starting with inner boundaries, Fig. 7a shows a number of *bCh* notifications for different values of k for connections affected by a failure and for those normal. Hence, configuring k equal to 3 allows keeping boundaries constant when normal BER behavior is monitored. In the event of connections with failure, less than 1% of all monitoring samples generate a *bCh* notification, which is enough to keep track of BER evolution with time as it will be shown in the following results.

Regarding the outer boundary, Fig. 7b shows the amount of *bExc* notifications as a function of k. Fixing k equal to 6 eliminates those notifications caused by atypical BER measures that do not entail failures, as well as keeps more than 90% of those notifications raised in the event of a failure. It is worth noting that *bExc* notifications are much less frequent than *bCh* ones and consequently, its impact on total notification overhead is negligible.

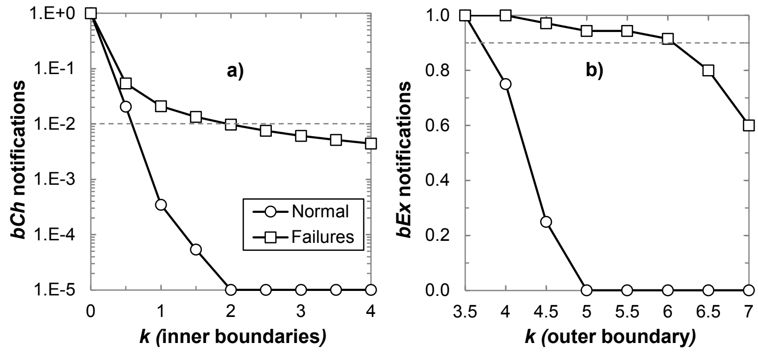

Figure 7. Tuning of BANDO parameters.

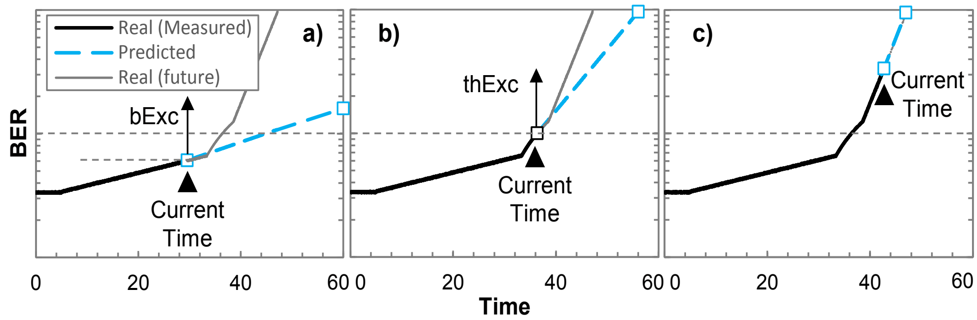

Figure 8. BER and boundaries evolution with time.

Finally, Fig. 8 illustrates the accuracy of the estimation for the time when max BER (1E-6) will be reached in case of a gradual drift failure. Prediction based on linear extrapolation is shown at three different time instants. Upon the reception of a *bExc* at day 30 (Fig. 8a), due to the lack of evidence of the actual future BER trend evolution, no max BER violation in the following 30 days is predicted. Later, upon the reception of a *thExc* at day 36 (Fig. 8b), max BER violation is predicted to happen in the near future. It is not until day 42, i.e., five days before the connection is disrupted, that prediction becomes steady to a constant value, which happens in Fig. 8c; hence, this method provides enough anticipation for an optimal reaction against the failure. Comparable results were obtained for the rest of gradual drift instances.

5. CONCLUSIONS

SLA violations entail money losses for the network operators and hence, minimizing such violations is of paramount importance to them. This paper focused on anticipating BER degradation detection at the optical layer, which typically supports many of the offered services.

In this regard, BANDO algorithm which works inside the optical nodes to take advantage of a fine monitoring granularity has been proposed to detect changes in the BER of optical connections. To evaluate its performance, different BER degradation failures were considered, including gradual and periodical degradation. Aiming at studying realistic scenarios, experimental measures were carried out on two different setups involving commercial equipment. The results of the experiments were used to generate synthetic data used to simulate the considered BER degradation failures.

Simulation results show that maximum BER violation was anticipated several days before the connection was disrupted, which allows planning a network reconfiguration to be performed on low activity hours. Interestingly, the BANDO algorithm demonstrated its advantage for failure detection compared to a centralized algorithm receiving notifications only after BER threshold violations.

ACKNOWLEDGEMENTS

The research leading to these results has received funding from the EC through the METRO-HAUL project (G.A. nº 761727), from the Spanish MINECO TWINS project (TEC2017-90097-R), and from the Catalan Institution for Research and Advanced Studies (ICREA).

REFERENCES

[1] A. Napoli *et al.*, "Next generation elastic optical networks: The vision of the European research project IDEALIST," *IEEE Communications Magazine*, vol. 53, pp. 152-162, 2015.

[2] A. P. Vela *et al.*, "BER degradation detection and failure identification in elastic optical networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 35, pp. 4595-4604, 2017.

[3] A. P. Vela, M. Ruiz, and L. Velasco, "Distributing data analytics for efficient multiple traffic anomalies detection," *Elsevier Computer Communications*, vol. 107, pp. 1-12, 2017.

[4] N. Sambo *et al.*, "Modeling and distributed provisioning in 10-40-100 Gb/s multirate wavelength switched optical networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 29, pp. 1248-1257, 2011.