Applying Data Visualization for Failure Localization

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

Abstract: Data visualization is applied to BER measures in a bigdata repository. Bubble charts are produced to identify lightpaths with increasing BER and spectrum color maps are then used to identify the most likely degraded link.

OCIS codes: (060.0060) Fiber optics and optical communications, (060.1155) All optical networks

1. Introduction

Recent advances on data analytics architectures for optical networks (e.g., [1]) are facilitating the introduction of intelligence and cognition towards autonomous network operation. In this new context, the role of human operators in the control and management of the network cannot be put aside but the opposite, it should be reinforced by the availability and accessibility of rich and accurate monitoring data. Such large amount of monitoring data, however, needs to be adequately presented by means of advanced *operation-oriented* data visualization methods that guide operators through myriads of data. In fact, insightful visualization cannot simply consist in periodically plotting a set of charts trying to statistically summarize the current status of the network. Instead, operation-oriented charts that are plotted according to a *visualization process* need to be specified for each desired use case, in the same way as different use cases require from different algorithms.

A challenging use case is the localization of soft failures affecting optical systems, e.g. optical amplifiers, that might degrade the Quality of Transmission (QoT) of optical connections (*lightpaths*) that are supported by such systems. What makes soft failures difficult to detect is that the produced degradation can be initially very subtle and thus, very difficult to detect before lightpaths' degradation exceed some threshold. In our previous work in [2], we presented an algorithm to detect Bit Error Rate (BER) degradation in lightpaths that runs directly in the network nodes assuming a scenario with nodes with computing capabilities [3] together with a data analytics algorithm [4] for failure identification running in the network controller. Certainly, another algorithm could be devised to correlate several degraded lightpaths and localize the common set of links supporting the lightpaths.

In this paper, we follow a different approach and explore the use of data visualization techniques to guide operators in failure localization tasks. The proposed visualization process computes first a *bubble chart* using all BER measures (e.g., one measure for every 15 minutes) for all lightpaths in the network for a given period of time (e.g., the last month) available in a centralized big data repository. The number of bubbles can be fixed to achieve the best data representation, whereas the color of each bubble is computed considering two different metrics: maximum BER and BER trend; the resulting bubble chart shows extreme usefulness to detect lightpaths with an increasing BER degradation within the considered time period. In a subsequent step, a *network spectrum color map* computed with the lightpaths in the bubble of interest allows detecting the link(s) that might be responsible for QoT degradation.

2. Optical Links' Soft Failure Visualization

Many systems for network management include some sort of data visualization to facilitate operators to localize failures in the network. Although these tools are really useful when the degradation is high, they fail to provide trend information, so the detection cannot be anticipated. For illustrative purposes, Fig. 1 shows an example, where two maps are presented for two different times t_1 and t_2; colors are used to give information about lightpaths' QoT thus highlighting those with poor values.

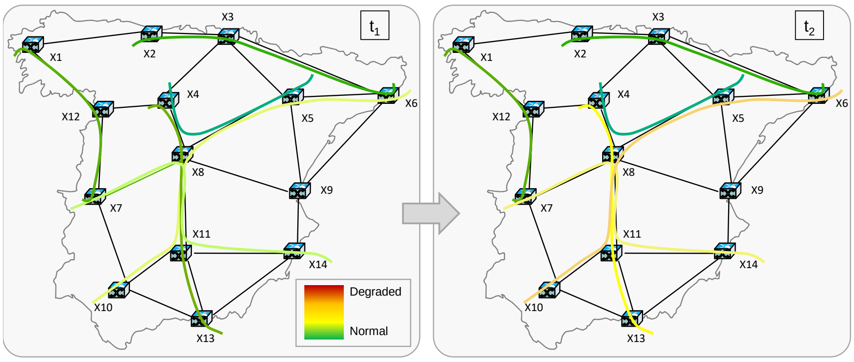

Fig. 1. Lightpath BER degradation evolution as a consequence of a soft failure in link X8-X11.

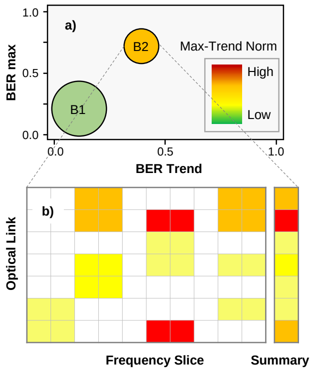

Fig. 2. Bubbles and spectrum map.

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).

In view of the above, we propose to use bubbles charts especially tailored to provide the specific information that network operators need to detect soft failures before they can degrade the QoT of established lightpaths. Fig. 2a illustrates the bubble chart resulting from the evolution from t_1 to t_2 depicted in Fig. 1. Bubble sizes give information about the number of lightpaths they include, whereas its position gives information of the relation between the maximum BER w.r.t. the BER change in the period; the metrics are relative to the expected BER computed using analytical formulae. Finally, bubble colors represent the L2-norm of the bubble position. Two bubbles are represented in Fig. 2a aggregating lightpaths with low BER and trend (B1), and those lightpaths with high BER and appreciable trend (B2). In view of bubble B2, the operator can decide to further analyze the cause of failure of the paths contained in such bubble; to this end, he/she selects bubble B2 and chooses to represent the lightpaths in a Network Spectrum Color Map (Fig. 2b). The spectrum color map shows the use of every frequency slice in the network link by link; each cell is colored according to the L2-norm of the lightpath using it. A *row-summary column* is additionally displayed to assist the operator in finding those links supporting the highest number of degraded lightpaths.

3. Data Visualization Techniques

Focused on obtaining a faithful representation of what is going on in the network, two visualization techniques are proposed. In general, any visualization process requires from pre-processing monitoring data records stored in one or more (big data) repositories to produce meaningful variables to be visualized in a chart. In the case of lightpaths' monitoring, data records contain, among others: *i*) time stamp (t); *ii*) lightpath identifier (p); and *iii*) measured BER, BER_{tp}. In addition, the lightpath operational database (LSP-DB) contains data about the lightpaths themselves, including their route and spectrum allocation, length, and estimated BER reference value, BER_Ref_p. Finally, a global BER threshold, BER_Thr, is configured. From such data, the pre-process phase for the considered visualization techniques transforms lightpaths' BER measurements producing a new variable, BER'_{tp}, representing the BER within the interval [BER_Ref_p, BER_Thr]; i.e., $BER'_{tp} = (BER\_Thr - BER_{tp})/(BER\_Thr - BER\_Ref_p)$, where BER_{tp} is previously forced to be confined in the defined interval. Next, two variables are computed by aggregating BER'_{tp} data from a selected time period (T): *i*) maximum BER in T, BER_Max_{Tp}, computed as the quartile with probability of 95% in order to avoid spurious values; and *ii*) BER trend in the period, BER_Trend_{Tp}, computed using the averaged first and last BER'_{tp} values. The visualization database, *visualizationDB*, is finally created combining these two variables together with useful data about the lightpaths; such database will be the input of the visualization algorithms that eventually will produce the charts.

Owing to the fact that visualization is fostered by colors, in this paper we use a color palette specifically designed to guide operators in finding problems in the network. The proposed color palette is defined as a set of concatenated non-overlapping segments of gradient color and threshold values in the continuous interval [a, b] ({[color_a, color_b], [a, b]>}, a, b in [0, 1]), where the color of a given data value in the interval [0,1] results from finding the segment representing the data and then computing the color in the defined gradient.

Now, visualization algorithms can produce operation-oriented charts. The bubble chart algorithm in Table I uses a k-means algorithm to find points in a 2D space (i.e., centroids), so that paths are grouped by assigning them to the nearest centroid. Each centroid is characterized by the coordinates BER_Trend_{Tp} (x-axis) and BER_Max_{Tp} (y-axis), and by the list of paths contained in the centroid (line 1 in Table I). Next, for each centroid, its color is computed according to the L2-norm of the vector representing its position, i.e., $\|(BER\_Max_{Tp}, BER\_Trend_{Tp})\|_2$, within the color palette (lines 2-4). Finally, the algorithm returns the set of bubbles B (line 5). According to bubbles' color, one can infer the severity of the paths enclosed; note that this will pilot the operator towards these suspicious paths.

Table I Bubble Chart Algorithm

```
INPUT visualizationDB, colorPalette, numBubbles     OUTPUT Bubbles
1:  C={<xPos, yPos, paths>} ← k-means(visualizationDB, numBubbles)
2:  $B \leftarrow \emptyset$
3:  for c in C do
4:      $B \leftarrow B \cup$ {<c.xPos, c.yPos, |c.paths|, getColor(colorPalette, ||(c.xPos, c.yPos||$_2$)>}
5:  return B
```

A subsequent analysis of these paths will derive in finding common causes leading lightpaths appreciable BER values or BER trend. To assist this task, we propose to use a network spectrum color map, a matrix representing the optical links as rows and the spectrum slices as columns; the color of each cell inherits the color of the lightpaths, computed likewise as for the bubbles, i.e., using the L2-norm of its vector.

4. Illustrative Results

To evaluate the proposed data visualization techniques, we carried out simulations on a realistic 30-node and 56-link Spanish Telefonica's optical network, where 800 100 Gb/s lightpaths using 3x12.5GHz frequency slices were set-up sequentially between randomly selected nodes. Next, lightpaths' BER was computed considering the expected OSNR in the links (used to compute BER_Ref_p) plus a randomly generated amount of errors; expected links' OSNR considered not only link's length but also its load [5]. A maximum pre-FEC BER that transponders can support before a lightpath is torn-down was set up as BER_Thr. Finally, we emulated a gradual degradation in link F08-F09, which decreases its OSNR and hence, increases BER of lighpaths using this link.

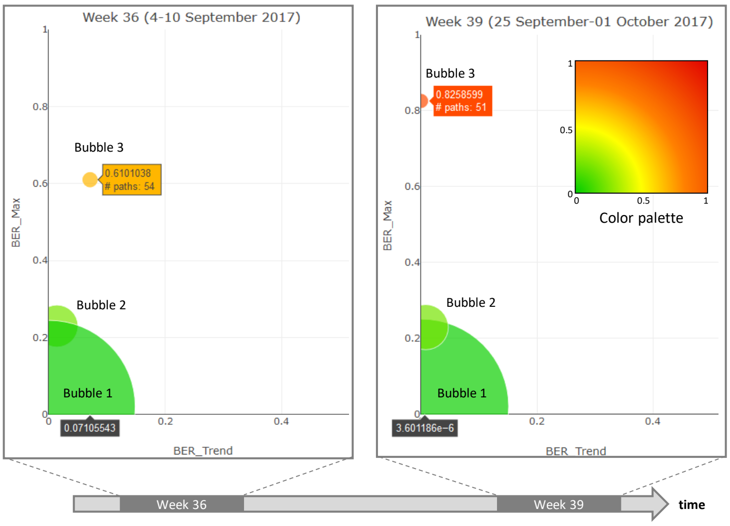

Fig. 3. Bubble charts for weeks 36 and 39.

Fig. 3 presents the results of applying the bubble chart algorithm previously defined on two very different scenarios: *i*) stable scenario (right), and *ii*) gradual degradation scenario (left). In both scenarios, lightpath BER measurements of the selected week (36 and 39) are visualized using three bubbles. The defined color palette is also presented.

Table II Bubble Charts Summary

	Bubble	**Paths**	**BER Trend**	**BER max**
Week 36	1	604	0.0	0.02
	2	142	0.01	0.23
	3	54	0.07	0.61
Week 39	1	595	0.0	0.008
	2	154	0.0	0.23
	3	51	0.0	0.826

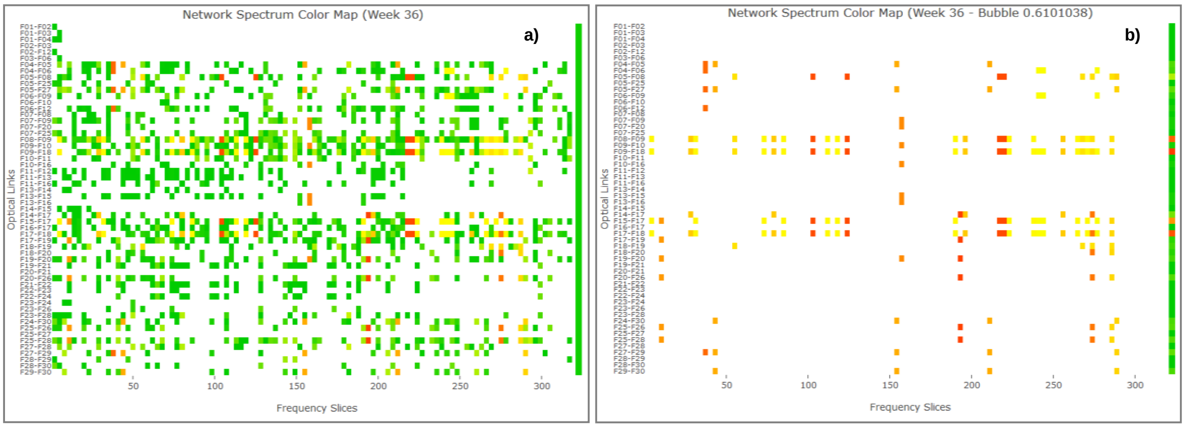

Fig. 4. Network Spectrum Color Map computed with all lightpaths (a) and with the lightpaths in bubble 3 (b).

In the first scenario (Fig. 3 right), one can observe that although one bubble appears with high BER, there is no trend, i.e., the cause of the high BER in the lightpaths is now stable. The operator could request to visualize previous weeks to find the period where the degradation happened. This leads to our second scenario, which is presented in Fig. 3 left; in this bubble chart for week 36, the operator clearly identifies one bubble with significant BER trend (bubble 3). A summary of the two bubble charts is presented in Table II.

Once the time when the degradation appeared has been identified, the operator might decide to find whether the cause of failure is in an optical link; to that end, he/she can select another operation-oriented chart to visualize a network spectrum color map. To clearly appreciate the goodness of the proposed visualization process, let us assume that no previous filtering is performed, so Fig. 4a presents the network spectrum color map when all the paths in the network (800) are selected. Although a trained eye could perceive that few links might be the responsible for the degradation, such conclusion is not obvious. In fact, the row-summary column in the spectrum color map, which is intended to highlight the most likely degraded links, do not show any clear identification. Conversely, computing the spectrum color map with only the paths in bubble 3 that summarizes 54 paths, results is a much clear map (see Fig. 4b). In this case, the row-summary column identifies four links (out of 56) to likely be the responsible of causing degradation on the lightpaths. It is easy now for the operator to inspect one by one each of these optical links to find whether there is a clear responsible for the degradation.

5. Conclusions

A data visualization process based on advanced graphical representation has been proposed for the localization of soft failures affecting lightpaths. In the first step, bubble charts using specific metrics and color palette has been proposed to identify, if any, those lightpaths deserving deep inspection because of unexpected high and/or increasing BER. Secondly, network spectrum color maps have been proposed as an *ad hoc* technique for accurate localization of the failing optical fiber link.

References

[1] L. Velasco et al. “An Architecture to Support Autonomic Slice Networking,” *IEEE/OSA JLT*, 2018.
[2] A. P. Vela et al., “BER Degradation Detection and Failure Identification in Elastic Optical Networks,” *IEEE/OSA JLT*, 2017.
[3] A. P. Vela et al., “Distributing Data Analytics for Efficient Multiple Traffic Anomalies Detection,” *Elsevier Comp. Comm.*, 2017.
[4] C. Bishop, *Pattern Recognition and Machine Learning*, Springer-Verlag, 2006.
[5] P. Poggiolini, “The GN Model of Non-Linear Propagation in Uncompensated Coherent Optical Systems,” *IEEE/OSA JLT*, 2012.