

Visualization Tools For Enhancing Failure Localization

Alba P. Vela, Marc Ruiz, and Luis Velasco

Universitat Politècnica de Catalunya, Barcelona, Spain

e-mail: apvela@ac.upc.edu

ABSTRACT

The use of data visualization techniques is explored to guide human operators in failure localization tasks. In this paper, we consider the case where a subset of lightpaths' experience gradual BER degradation. We aim at using visualization techniques to find out the most likely resource that is responsible for such degradation, as well as the start of the degradation. Note that the challenge is to analyze an amount of data that might be huge (e.g., 1 sample every 15 min for the last, let us say, 4 weeks, for every lightpath in the network) and display charts in seconds, while revealing meaningful information for human operators. The proposed task-oriented visualization tackles such challenge by using charts that first help finding the subset of affected lightpaths thus, reducing the number of lightpaths to be analyzed. In a second step, different charts are used to reduce the element in failure as much as possible. Finally, the few suspicious elements can be analyzed in detail using more traditional timeline graphs.

Keywords: data visualization, data analysis, failure localization and optical networks.

1. INTRODUCTION

Data analytics architectures for optical networks like the presented in [1], are facilitating the introduction of intelligence and cognition towards autonomous network operation. A challenging scenario, where data analytics helps, is the localization of soft failures affecting optical systems [2]. 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. It is important to highlight, in the context of autonomous networking, that 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.

In fact, insightful visualization cannot simply consist in periodically plotting a set of charts trying to statistically summarize the current status of the network. Indeed, typical data visualization tools available in many management systems include *timeline charts* to represent single time-series, e.g., the evolution of BER monitored on a chosen lightpath. Commercial data visualization tools are including charts fueled by statistical and ML algorithms that allow extending dashboards with information extracted from manually selected monitoring metrics, e.g., to highlight whether a lightpaths' BER is likely to be anomalous [BluePlanet].

Specifically for the case of failure localization, many network management systems include some sort of data visualization to facilitate such task to human operators. 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.

Instead the network maps in Fig. 1, *task-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. More sophisticated charts need to be integrated to facilitate human operators to understand the relationship among monitoring metrics, find meaningful events, and correlate observable consequences caused by the same event. In particular, we are considering the case of a gradual degradation in a link, where increasing BER values have been collected for those lightpaths traversing such link; we are assuming that BER monitoring values are available for all the established lightpaths. Some questions that cannot be answered looking at Fig. 1 at t_1 or t_2 , are: *i)* is lightpaths' QoT normal? *ii)* Might be BER is a bit high but, should we be concerned? *iii)* Are lightpaths quickly degrading?

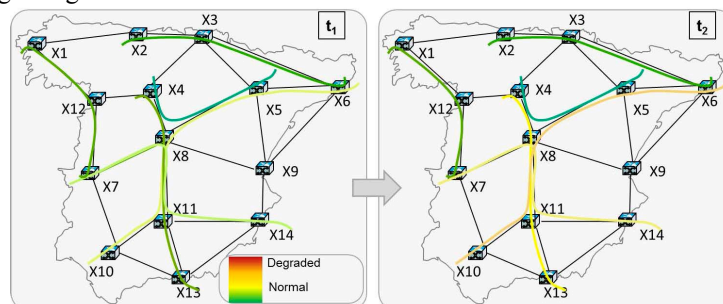


Figure 1. Lightpath BER degradation evolution as a consequence of a soft failure in link X8-X11.

2. TASK-ORIENTED VISUALIZATION

Aiming to answer the above questions and to localize the element responsible for the degradation, we propose a *visualization process* that consists of a set of *chained* task-oriented charts. When requested by the network operator, a huge amount of monitoring data belonging to all the established lightpaths collected during a period of time, need to be analyzed [left part of Fig. 2(a)]. It is clear that one cannot simply use a timeline chart to plot the BER evolution for all the lightpaths, so other charts should be used. The visualization process that we propose consists in iteratively present the operator with charts that help him/her to reduce the number of suspicious elements causing a failure that affects a subset of lightpaths [right Fig. 2(a)].

An example of this chained process is presented in Fig. 2; which uses a bubble chart at the first step, where we aim at first to have an insight of which lightpaths show a suspicious behavior. After this first step, one might want to know whether there is/are common resources responsible for such degradation, so other visualization charts are suggested for that end. Finally, once the number of suspicious elements is small, one-by-one analysis can be carried out. As it can be observed in the workflow in Fig. 2, we have analyzed a set of three chained task-oriented charts, however, other combinations of charts can be considered.

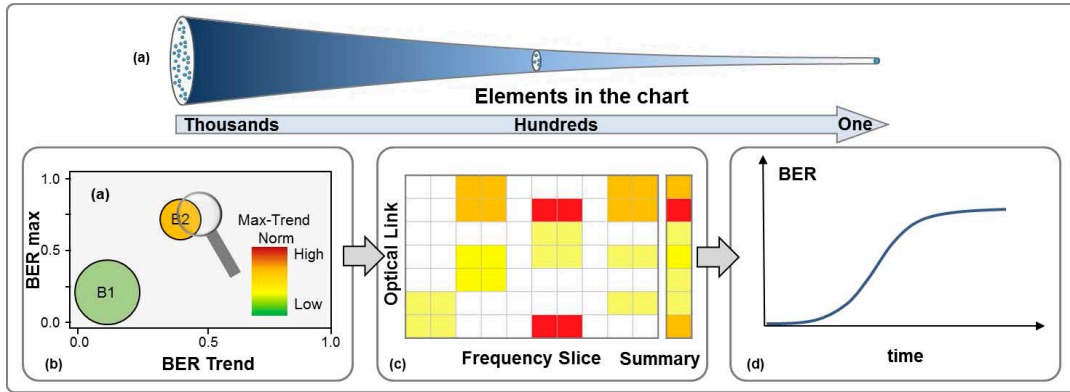


Figure 2. Example of chained task-oriented applying visualization-assisted data filtering (a), bubble chart (b), network color spectrum map (c), historical BER (d).

2.1 Data Pre-Processing

Task-oriented charts might require examining long monitoring time-series data sets from a large huge of elements; e.g., analyzing BER time-series for the last month for a set of 5,000 lightpaths entails retrieving and processing about 110 MB of data. Although processes behind task-oriented charts are designed to produce the essential chart configuration data, i.e., chart elements, color scheme, etc., to highlight the desired visualization effect, monitoring data should be aggregated at collecting time aiming at reducing processing at visualization time. Besides, in order to use visualization charts, pre-processing monitoring data records is mandatory to produce meaningful variables to be visualized. In the case of lightpaths' monitoring, data records contain, among others: *i*) time stamp (*t*); *ii*) lightpath identifier (*p*) and *iii*) measured BER, BER_p . In addition, the lightpath operational database contains data about the lightpaths themselves, including their route and spectrum allocation, length, and estimated BER reference value, BER_{Ref_p} . Finally, let us assume that a global BER threshold, BER_{Thr} , is configured as a limit of BER for all the lightpaths in the network. From such data, the pre-process phase transforms lightpaths' BER measurements producing a new variable, BER'_{tp} , representing the BER within the interval $[BER_{Ref_p}, BER_{Thr}]$; i.e.,

$$BER'_{tp} = \frac{(BER_{Thr} - BER_p)}{(BER_{Thr} - BER_{Ref_p})}, \tag{1}$$

where BER_p 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*):

- **Maximum BER** in *T*, $BER_{Max_{tp}}$, computed as the quartile with probability of 95% in order to avoid spurious values; and
- **BER trend** in the period, $BER_{Trend_{tp}}$, computed using the averaged first and last BER'_{tp} values.

A visualization database, *visualizationDB*, is 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, 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.

2.2 Different types of considered charts

Bubble Charts

Specially tailored *bubbles charts* can be used to provide the specific information that network operators need to detect soft failures before they can degrade the QoT of established lightpaths. Figure 2(b) illustrates the bubble chart resulting from the evolution from $t1$ to $t2$ depicted in Fig. 1. The proposed bubble chart uses BER measures for all lightpaths in the network for a given period of time (e.g., the last month) available in the big data repository in the MDA controller. The main features of the proposed bubble chart are summarized as follows:

- **Bubbles' position**, giving information of the relation between the BER value w.r.t. the BER change (trend) in the period; the metrics are relative to the expected BER for every lightpath.
- **Bubbles' size**, gives information about the number of lightpaths a bubble includes.
- **Color of each bubble**, computed with the L2-norm of BER and BER trend (bubble's position).

The resulting bubble chart shows extreme usefulness to detect lightpaths with an increasing BER degradation within the considered time period. As an illustrative example, two bubbles are represented in Fig. 2(b) 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 new task-oriented chart.

The bubble chart algorithm in Table 1 uses the *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 1). 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 1. Bubble Chart Algorithm.

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

Network Spectrum Color Map

A network spectrum color map (Fig. 2(c)) is 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. The intention is to find common causes leading lightpaths appreciable BER values or BER trend. A *row-summary column* is additionally displayed to assist the operator in finding those links supporting the highest number of degraded lightpaths.

Single Element Analysis

The previous charts can lead the operator to an accurate filtering, where coming from thousands of lightpaths [left part of Fig. 2(a)]. At the end of the visualization process, only a small amount of lightpaths need to be further analyzed [right part of Fig. 2(a)]. Figure 2(d) uses timeline plots to visualize the evolution of one single measurement with time.

3. ILLUSTRATIVE RESULTS

To evaluate the proposed data visualization techniques, we carried out simulations on a realistic 30-node and 56-link multilayer network based on the Spanish Telefonica's optical network, where a 14-node VNT was defined. Besides, 800 100 Gb/s lightpaths using 3×12.5 GHz frequency slices were set-up sequentially between randomly selected nodes; lightpaths' BER was computed considering the expected OSNR in the links (needed 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 [3]. 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 lightpaths using this link.

Figure 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. 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 2.

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. 4 presents the network spectrum color map when all the paths in the network (800 lightpaths) 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, does not show any clear identification.

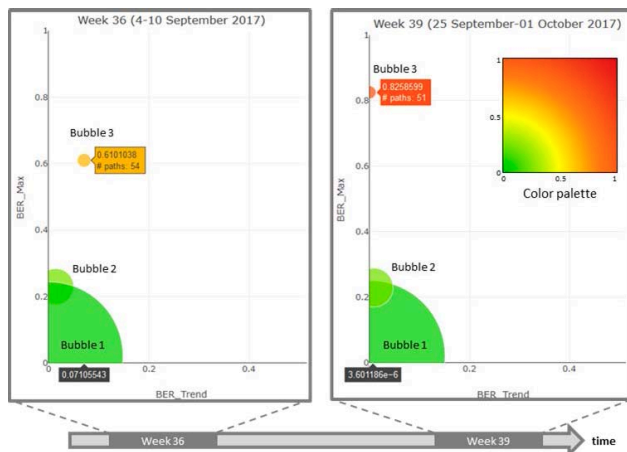


Figure 3. Bubble charts for weeks 36 and 39.

Table 2. Bubble Chart 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

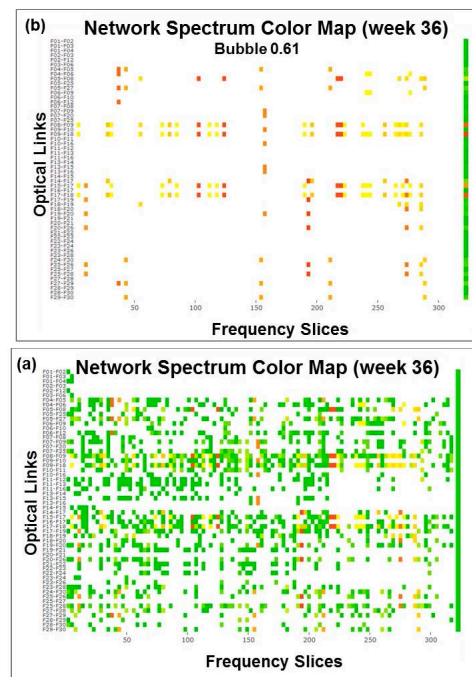


Figure 4. Network Spectrum Color Map computed with all lightpaths (a) and with the lightpaths in bubble 3 (b).

4. CONCLUSIONS

In this paper, we explored the use of data visualization techniques to guide operators in failure localization tasks. A data visualization process based on advanced graphical representation has been proposed for the localization of soft failures affecting lightpaths. In the first step, as a *visualization task-oriented chart*, the bubble chart using specific metrics has been proposed to identify, if any, those lightpaths deserving deep inspection because of unexpected high and/or increasing BER. Secondly, a network spectrum color map has been proposed as an *ad hoc* technique for accurate localization of the failing optical fiber link. As proven by the illustrative results, using the proposed chained charts, operators can easily track network performance and speeding up health diagnosis through a powerful and simple visualization process.

ACKNOWLEDGEMENTS

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

REFERENCES

- [1] L. Velasco *et al.*, "An architecture to support autonomic slice networking," *IEEE/OSA JLT*, 2018.
- [2] A. P. Vela, M. Ruiz, F. Fresi, N. Sambo, F. Cugini, G. Meloni, L. Poti, L. Velasco, and P. Castoldi, "BER degradation detection and failure identification in elastic optical networks," *IEEE/OSA JLT*, 2017.
- [3] P. Poggiolini, "The GN model of non-linear propagation in uncompensated coherent optical systems," *IEEE/OSA JLT*, 2012.