European Conference on Optical Communication (ECOC) 2022 © Optica Publishing Group 2022

Investigating Q-drops and Their Probable Causes

Camille Delezoide, Petros Ramantanis, Patricia Layec

Nokia Bell Labs France, camille.delezoide@nokia-bell-labs.com

Abstract *We leverage data science to investigate the Q-drop phenomenon on a public optical network monitoring dataset. We show that Q-drops above 1 dB are common at network scale and correlated in at least 86% of cases with signs of packet loss ©2022 The Author(s)*

Introduction

Modern optical transponders have the ability to easily adjust their transmission rates[1]. Through elasticity[2], it becomes possible to optimize network capacity by leveraging the large margins inherited from traditional design methods[3],[4]. We previously investigated various solutions to best adjust the rate of any network connection based on prior observations of the quality of transmission (QoT)[5],[6]. We particularly observed that field QoTs can show major variations, sometimes exceeding 6 dB within an hour[7]. Thus, any rate increase may damage the quality of service[8],[9].

More recently, we leveraged field data monitored from multiple connections to localize hard failures[10], i.e. consecutive periods during which packet losses were observed. The monitoring data, collected every second for 10 months, suggested that the large QoT variations in[11] termed Q-drops by Ghobadi *et al*[12] could actually be hard failures with short durations, below the 15-minute-resolution of their dataset. Importantly, the events behind these observations may lead to packet loss whatever the margin. Thus, their impact on availability could only be mitigated through maintenance or optical protection[13].

In this paper, we leverage the data in[11] to investigate Q-drops with a focus on determining what they reveal. First, we discuss the potential QoT dynamics upon Q-drops and outline the possibility that Q-drops correlate with packet loss. Then, we propose a robust method to detect Q-drops as QoT outliers to study their frequency and severity. Finally, we correlate the Q-drops with outlier observations for the other reported metrics to discuss their potential causes.

QoT dynamics behind Q-drops

Q-drops were first defined[12] as events where the average Q^2 drops from its stable value while remaining above the FEC threshold. The implicit hypothesis was that the instantaneous QoT remained close to the 15-minute average during

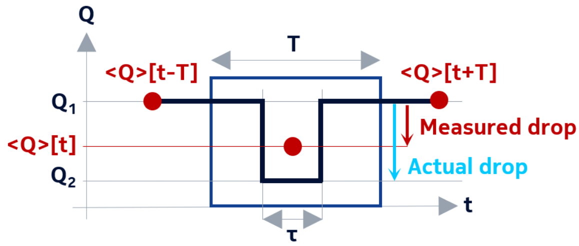

Fig. 1: Hypothesis for QoT variations upon Q-drops

Q-drops. Yet, from our own field data reported in[10], we observed that significant drops of the 15-minute-average QoT were always correlated with non-zero post-FEC BER values. These observations lead to us to an alternative hypothesis depicted in Fig. 1 describing the QoT dynamics behind the Q-drops observed in[11]. Naturally, this is a simplification of a more general view where QoT fluctuations may be considered arbitrary. Here, the duration of the drop τ is assumed smaller than the averaging period T. Consequently, the measured drop from QoT averages is necessarily smaller than the actual drop of the instantaneous QoT. From average conservation, this leads to:

$$\text{Measured drop} = \left(\frac{\tau}{T}\right) \text{Actual drop} \qquad (1)$$

From eq. (1), a measured drop of 5 dB could be explained by an actual drop of 15 dB during 5 minutes. This underlines the possibility of packet losses upon Q-drops, as observed in[10].

Method for investigating Q-drops

Henceforth, we focus on the data published in[11]. In a study focused on predicting outages from Q-drops[14], Hasegawa *et al* proposed to detect Q-drops as Q^2 values below 95% of its moving median over one day up to the studied value. Yet, this definition can lead to label Q-drops QoT values that are normal, i.e. not consistent with any sudden issue. This particularly occurs for connections with large daily variations of the trend.

In contrast, we choose here to define Q-drops as localized outlier QoT values, most likely due to severe issues. To do so, we first evaluated the

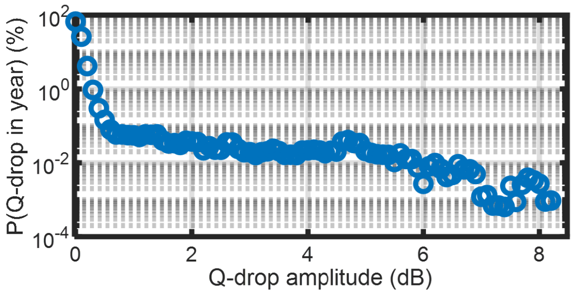

Fig. 2: Q-drop probability density function from field data

trends of the 4000 QoT time series of the dataset, one for each monitored connection. This was made in between service interruptions, occurring as missing data points. Amongst the many possible solutions to reconstruct the trend, we chose a robust weighted quadratic regression method[15]. Similarly to the moving median, this method minimizes the influence of outliers on the resulting trend. Once the trend is removed from the raw time series, outliers become much easier to detect. Yet, what constitutes an outlier remains ultimately subjective, as dependent on the regression window for the trend and on the selectivity of the subsequent outlier detection method. We chose a 6-hour regression window as what seemed to be an optimal choice to focus on short events, i.e. up to 30 minutes. Finally, we chose the quartiles method[16] to detect outliers in the detrended QoT time series, with the following Q-drop criterion:

$$Q^2 - \text{trend}(Q^2) < \Theta_1 - k(\Theta_3 - \Theta_1) \qquad (2)$$

where Θ_1 and Θ_3 are respectively the first and third quartiles of the detrended Q^2. We set the threshold factor k to 5. This was designed to favor false negatives over false positives, the latter being more harmful to subsequent interpretations.

Frequency and severity of Q-drops

We first apply the method described above to study how frequent Q-drops are depending on their amplitude, i.e. the absolute value of the detrended Q^2 outlier. In Fig. 2, we plot a Q-drop probability density function (PDF) as the probability to observe a Q-drop of amplitude in [X, X+0.1dB] for any given channel within a year as a function of X.

Below 1dB, the Q-drop PDF decays exponentially, which can be expected. Surprisingly however, the PDF does not significantly vary between 1 and 5 dB of Q-drop amplitude. This suggests that the underlying causes for Q-drops below and over 1 dB are completely different. Beyond 5 dB,

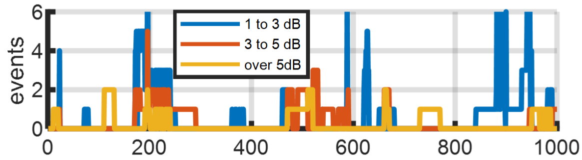

Fig. 3: Distribution of Q-drops for the first 1000 connections

we observe a progressive decay of the Q-drop PDF. However, this decay may only be due to the progressive rarity of connections that can support a drop of 5dB and more while maintaining above the 6.5dB FEC limit, below which data has been censored prior to public release. To study the disparities amongst the various connections, we plot in Fig. 3 the number of detected Q-drops according to amplitude for the first 1000 connections, other connections showing comparable results. We observe that many channels do not exhibit Q-drops above 1 dB, and more generally that Q-drops amplitudes strongly depend on the connection. This indicates that underlying causes for Q-drops are localized issues rather than uniformly distributed, which is consistent with the conclusions of our previous report[10].

Probable causes through correlations

Here, we finally investigate the correlations between Q-drops and anomalies detected on the other metrics available in the dataset: the received power (P), the cumulated chromatic dispersion (CD) and the polarization mode dispersion (PMD). Similarly to what is described in our Q-drop investigation method, we detrended the associated time series to focus on how much these metrics deviated from their six-hour-trends each time a Q-drop is detected. Note that we only made correlations between Q-drops and other metrics collected from the same connection. Then, for each metric, we flagged significant trend deviations by applying the following thresholds: 1 dB for Q^2, 0.5 dBm for P, 10 ps/nm for CD and 5 ps for PMD. Those threshold values stem from statistical studies similar to what is represented in Fig. 2. We chose them to filter out commonly observed variations, therefore most likely corresponding to nominal operation and not failure modes. The last step has been to count the occurrences for each of the 8 combinations of flags corresponding to the observations of Q-drops of 1dB and above. Note that we verified that alternative sets of thresholds only had a minor impact on the results discussed below.

In Fig. 4, we plot a pie chart summarizing the most commonly observed cases. The results

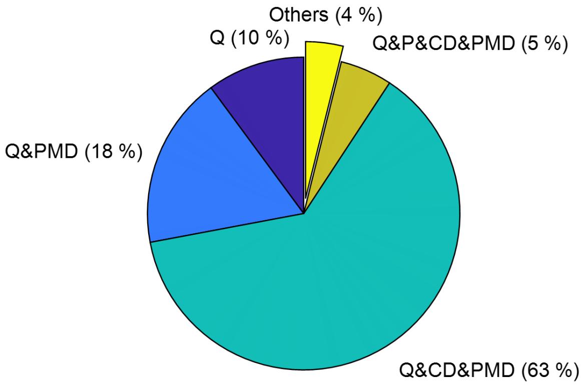

Fig. 4: Correlations of Q-drops with flagged deviations of other reported metrics

show that 63% of Q-drops above 1dB are correlated with outlier deviations of both CD and PMD, up to -20000 ps/nm and -80 ps respectively, and generally leading to absurd levels for such metrics. This suggests abnormal convergence of both CD compensation and constant modulus algorithm (CMA) blocks during digital signal processing (DSP) within the receiver, which should logically lead to packet loss and would explain the simultaneous Q-drop. A plausible explanation would be a major attenuation of the signal prior to at least one amplifier. Indeed, any subsequent amplification would fill the spectral slot with noise, explaining why the received power does not significantly vary in such cases. In comparison, cases where all metrics strongly deviates including P represent only 5% of cases. Those are most likely due to issues occurring between the last amplifier and the receiver, hence the lesser probability. Representing 18% of occurrences, the second most frequent cases are Q-drops correlated with outlier PMD deviations, pointing to a CMA failure. Indeed, the fact that the CD remains nominal suggests that the underlying issues are polarization-related[17]. In 10% of cases, Q-drops above 1dB are not accompanied by outlier deviations of the three other metrics. These issues may occur because of an issue only revealed beyond the CMA, such as the carrier frequency estimation (CFE) or phase estimation (CPE). Finally, the rest of the cases represents 4%, representing situations that are more complex to diagnose.

For a deeper analysis of the most frequent cases, we plot in Fig. 4 a heatmap where each point represents a Q-drop above 1dB. The coordinates are the CD and PMD deviations at the Q-drop and the colors represent the Q-drop amplitudes. The Q-drops found on the x-axis corresponds to events where Q-drops are only correlated with CD-drops. Such cases only occur be-

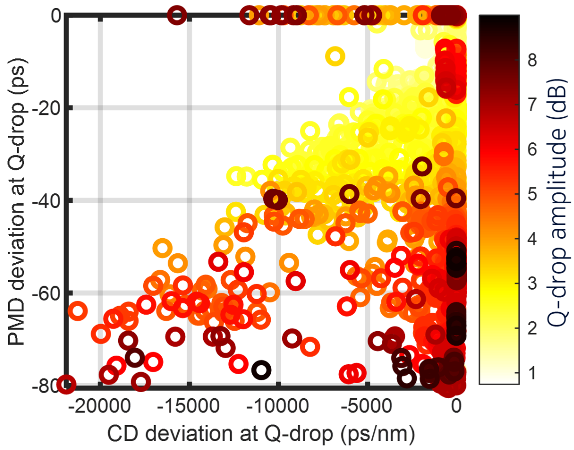

Fig. 5: Heatmap of Q-drops above 1dB where coordinates stands for the simultaneous deviations of CD and PMD

cause the PMD values for the related channels were not available. Note that because they represent incomplete observations, those cases were removed from the statistics in Fig. 4. In contrast, the Q-drops on the y-axis are real cases where the PMD drops, but not the CD. Yet, the majority of displayed Q-drops correspond to correlated and significant deviations of Q^2, CD and PMD. Interestingly, we observe that the Q-drop amplitude strongly correlates with the PMD deviation, but not with the CD deviation. A possible explanation is that when the signal is missing, the output of the CD measurement tends to be random while the PMD output reaches a constant negative value below -80ps, similarly to what is depicted in Fig. 1 for Q^2. Thus, amplitudes of Q-drops and PMD-drops may be correlated by the actual duration τ of the underlying failure. The fact that most Q-drops lie in the same half of the quadrant suggests that while CD output is mostly random upon Q-drops, there should be a low saturation value below -20000 ps/nm under which the instantaneous CD never goes.

Conclusion

We leveraged field data to investigate Q-drops and their probable causes. We uncovered that Q-drops below and above 1dB seem to be generated by completely different phenomenons, and that issues behind the larger Q-drops appear localized, confirming prior observations. We further found that 86% of Q-drops above 1dB are correlated with outlier PMD values pointing to critical failures, thus consistent with our hypothesis that Q-drops generally hide packet losses.

Acknowledgements

We thank AINET Fog-Antillas for their financial contribution.

References

[1] J. Cho and P. J. Winzer, “Probabilistic constellation shaping for optical fiber communications”, *Journal of Lightwave Technology*, vol. 37, no. 6, pp. 1590–1607, 2019.

[2] M. Jinno, “Elastic optical networking: Roles and benefits in beyond 100-gb/s era”, *Journal of Lightwave Technology*, vol. 35, no. 5, pp. 1116–1124, 2016.

[3] J.-L. Augé, “Can we use flexible transponders to reduce margins?”, in *Optical Fiber Communication Conference*, Optical Society of America, 2013, OTu2A–1.

[4] C. Delezoide, “Method for a comprehensive evaluation of margins in optical networks”, in *45th European Conference on Optical Communication (ECOC 2019)*, IET, 2019, pp. 1–4.

[5] C. Delezoide, K. Christodoulopoulos, A. Kretsis, *et al.*, “Marginless operation of optical networks”, *Journal of Lightwave Technology*, vol. 37, no. 7, pp. 1698–1705, 2019.

[6] C. Delezoide, P. Ramantanis, and P. Layec, “Leveraging field data for the joint optimization of capacity and availability in low-margin optical networks”, *Journal of Lightwave Technology*, vol. 38, no. 24, pp. 6709–6718, 2020. DOI: 10.1109/JLT.2020.3022107.

[7] C. Delezoide, P. Ramantanis, and P. Layec, “Machine learning and data science for low-margin optical networks: The ins and outs of margin optimization”, in *Machine Learning for Future Fiber-Optic Communication Systems*, Elsevier, 2022, pp. 281–315.

[8] M. Tornatore, G. Maier, and A. Pattavina, “Capacity versus availability trade-offs for availability-based routing”, en, *Journal of Optical Networking*, vol. 5, no. 11, p. 858, 2006, ISSN: 1536-5379. DOI: 10.1364/JON.5.000858. [Online]. Available: https://www.osapublishing.org/jon/abstract.cfm?uri=jon-5-11-858 (visited on 06/11/2019).

[9] C. Delezoide, *Moneybaud: Using data and automation to get the most from optical networks*, 2021. [Online]. Available: https://www.nokia.com/blog/moneybaud-using-data-and-automation-to-get-the-most-from-optical-networks/.

[10] C. Delezoide, P. Ramantanis, L. Gifre, F. Boitier, and P. Layec, “Field trial of failure localization in a backbone optical network”, in *2021 European Conference on Optical Communication (ECOC)*, 2021, pp. 1–4. DOI: 10.1109/ECOC52684.2021.9606152.

[11] Microsoft, *Wide-area optical backbone performance*, 2017 (accessed December 9, 2020). [Online]. Available: https://www.microsoft.com/en-us/research/project/microsofts-wide-area-optical-backbone/.

[12] M. Ghobadi and R. Mahajan, “Optical Layer Failures in a Large Backbone”, en, in *IMC proceedings*, ACM Press, 2016, pp. 461–467. DOI: 10.1145/2987443.2987483.

[13] M. Tornatore, G. Maier, and A. Pattavina, “Availability design of optical transport networks”, en, *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 8, pp. 1520–1532, Aug. 2005, ISSN: 0733-8716. DOI: 10.1109/JSAC.2005.851774. [Online]. Available: http://ieeexplore.ieee.org/document/1490640/ (visited on 06/11/2019).

[14] Y. Hasegawa and M. Uchida, “Predicting network outages based on q-drop in optical network”, in *2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 1, 2019, pp. 258–263. DOI: 10.1109/COMPSAC.2019.00045.

[15] R. Andersen, *Modern methods for robust regression*. Sage, 2008.

[16] P. J. Rousseeuw and M. Hubert, “Robust statistics for outlier detection”, *Wiley interdisciplinary reviews: Data mining and knowledge discovery*, vol. 1, no. 1, pp. 73–79, 2011.

[17] S. J. Savory, “Digital coherent optical receivers: Algorithms and subsystems”, *IEEE Journal of selected topics in quantum electronics*, vol. 16, no. 5, pp. 1164–1179, 2010.