Service Margins for Wide-Band Optical Spectrum Services Implemented in Long-Haul Raman-Enabled Networks

Kaida Kaeval[1,2], Sai Kireet Patri[1], Jose-Juan Pedreno-Manresa[1], Klaus Grobe[1], Jörg-Peter Elbers[1], Helmut Griesser[1], Marko Tikas[2], Gert Jervan[3]
[1] ADVA Optical Networking SE, Martinsried, Germany
[2] Tallinn University of Technology, Tallinn, Estonia
[3] Tele2 Estonia, Tallinn, Estonia
kkaeval@adva.com

Abstract: Operators depend on reliable Q-margin estimates to reliably operate wide-band Optical Spectrum Services. We use long-term performance tracking in live Raman-enabled links to derive service margins for a 375-GHz Optical Spectrum Service.

1. Introduction

Dense Wavelength Division Multiplexing (DWDM) is the staple technology of every operator core network. To achieve higher efficiency in the networks and avoid unnecessary 3R regeneration (Re-amplification, Re-shaping, Re-timing) at domain borders, operators are open to consider new service models. One of the promising service models is Optical Spectrum as a Service (OSaaS) [1], which started to gain momentum with the network disaggregation trend. Unique for this service model, the end user is granted direct access to the spectral resources in the provider's DWDM network. With this, the end-user can freely choose their transceivers and operate them over the spectral resources provided. However, each underlying network infrastructure and DWDM system is unique, and system details are considered business-critical information. This complicates the achievable performance calculations and possible performance fluctuation estimations in the networks. To overcome this, statistical approach in maintaining service margins has been proposed by the cloud providers [2].

This paper utilizes telemetry data collection over a 25-days to derive service margins for long-haul Raman-enabled networks implementing wide-band OSaaS. For this task, a single 375-GHz wide-band OSaaS lightpath was configured on links with 1792 km, 3752 km, and 5728 km. Then, 30-second interval was used to collect instant Q-value readings from five transceivers operated within the OSaaS, configured to transmit 200-Gbit/s DP-QPSK 69-GBd. A two-hour averaging cycle is used to estimate the margins to compensate for the slow performance drifts, and raw values from an hour-long tracking are used to derive the margins for short-term power fluctuation compensation.

2. System and data collection set-up

In essence, an OSaaS is a transparent lightpath service connecting two end-points in a single or multi-domain optical network. Differentiated from an alien wavelength service, it can transport multiple carriers over a predetermined spectrum slot. This means the wavelength-dependent performance fluctuations are emphasized with the increasing length of the service. This is specifically evident for wide-band OSaaS operated in Raman-enabled long-haul networks.

For this work, a 375-GHz wide-band OSaaS was provisioned on Tele2 Estonia's long-haul Raman-enabled production network between Tallinn and Frankfurt, which provided three different link lengths via loopbacks. Following the definitions in ITU-T Recommendation G.807 [3], we define the service envelope for the OSaaS as a media channel (MC) and individual carriers within it as optical tributary signals (OTSi).

The test set-up with spectral assignments and geographical layout of the network is illustrated in Fig. 1, left. The optical user spectrum was generated with five ADVA TeraFlex™ transceivers configured for200-Gbit/s DP-QPSK 69-Bd transmission. The spectrum was added and dropped at a free terminal Reconfigurable Optical Add/Drop Multiplexer (ROADM) port using a set of 8:1 splitter/combiner module and an optical amplifier at the test site in Tallinn. The user signals were pre-emphasized between -3.0 dBm and +3.0 dBm to partially compensate for the Raman ripple effect within the allocated spectrum. To extend the transmission distance and to allow single-ended measurements from the test site, the MC-s were looped back in the far-end ROADMs. The three provided links with a looped length of 1792 km, 3752 km, and 5837 km utilized a standard single-mode fiber and used aerial Optical Power Ground Wires on the Baltic part of the network. To capture the maximum changes in the link performance, automatic power equalization on the link compensation was disabled.

The telemetry data collection arrangement is illustrated in Fig.1, Right. In the OSaaS service model, the system-related monitoring data, although collected by the provider's Network Management System (NMS), is not accessible to the OSaaS end-user. Hence, only the telemetry data collected from the transceiver equipment operated by the end-

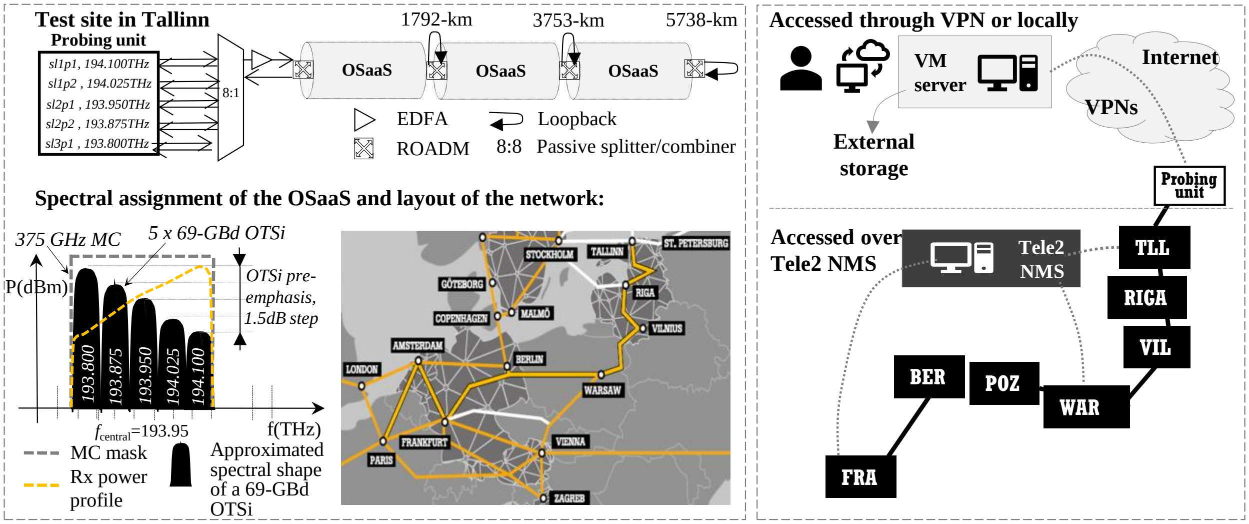

Fig. 1 Left: Test set-up, spectral assignment with signal pre-emphasis and layout of the network, Right: Remote data collection set-up for the telemetry collection

user, can be used for margin derivation. The telemetry data was automatically retrieved from the transceivers using automated telemetry scripts based on vendor-specific NETCONF commands. The data was retrieved every 30 seconds and included forty parameters per OTSi, including current Q-value and optical receive (Rx) power. The retrieved data was stored in a time-series database (InfluxDB) and visualized using a Grafana dashboard with an option to export as comma-separated values (CSV) for further processing.

3. Results and interpretation

The primary noticeable pattern in the recorded data was the slow drift with a 24-hour/daily pattern, which was surrounded by fast performance fluctuations. The longer the link, the more pronounced was the daily pattern. To analyze the magnitude of the variations and hence, derive the service margins for slow performance drift, averaging over a two-hour time window was performed to eliminate the impact from fast performance variations, limiting the dataset to 301 unique data points per transceiver, as presented on Fig. 2 left. As visible from the figure, the worst performance on the first link was on the OTSi operated at the best possible spectral position at 194.100 THz within the OSaaS according to Fig.1 left. This is caused by the signal pre-emphasis at the transmit end, where 194.100 THz is suppressed by -3 dBm and 193.800 THz increased by +3 dB, with other signals adjusted in between. By the end of second link, the performance over the spectrum has equaled out and varies only ±0.5 dB around the mean Q-value. On the longest link, however, the amount of pre-emphasis at the transmit end is not enough. This leaves three transceivers operated at the lower performing area of the OSaaS performing below the Q-value threshold. Transceivers operating below the Q-threshold were discarded from the analysis. Then, the standard deviation (σ) was calculated based on the averaged values to derive service margins based on 6×σ which would cover 99.7% of the performance variations caused by the slow drift. These margins were compared to the maximum and minimum detected Q-values in dB from the link. Assuming the ideal normal distribution of the 301 data points, 0.3 % would lead to less than one data point being outside of the found margins. The derived average service margins to compensate for the slow performance drift based on a 6×σ for the Q-value were 0.87 dB, 1.45 dB, and 1.56 dB for the 1792-km, 3752-km, and

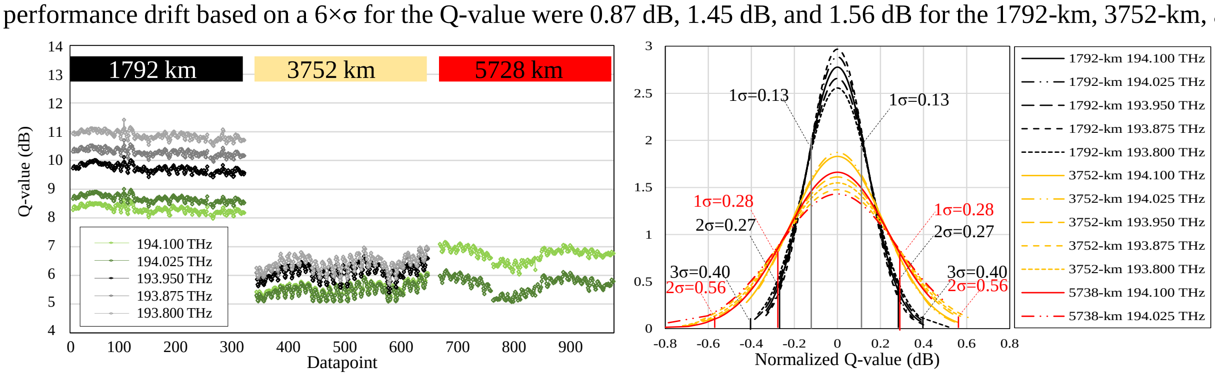

Fig. 2 Left: Averaged slow drift in Q-value for all links, Right: Normalized Q distribution on all links

5738-km links respectively and are able to compensate any operation point within the daily slow drift. Differences between the links and transceivers are presented in Fig. 2 right. The colors of the curves refer to the monitored link and style the used transceiver frequencies. The black and red standard deviation values on the figure stand for the best-performing transceiver on the 1792-km link and worst-performing transceiver on the 5738-km link.

Additionally, an hour-long raw data consisting of 124 data points per transceiver was used to calculate the standard deviation for fast performance fluctuations in the network for each of the transceivers. Similarly, to the slow drift, the standard deviations were calculated per transceiver and per link length. Based on all five ports, the per-link average service margins to compensate for the fast power fluctuations based on 6×σ were 0.32 dB, 0.345 dB, and 0.44 dB for the 1792-km, 3752-km and 5738-km links respectively. Comparing the per transceiver performance, the derived margins were higher for the transceivers working with lower Rx power levels. This emphasizes that wide-band OSaaS should be operated together with power spectral density-based spectrum equalization schemas when possible.

Fig. 3 summarizes the results and presents the total required margins to compensate for slow and fast performance variations in the tested long-haul network. The results per link length are marked with bar colors. As the signals operated in the lowest-performing area of the OSaaS spectrum were performing under the Q-threshold, the margins on the longest link could not be derived for transceivers working at 193.950, 193.875 and 193.800 THz. These are marked with N/A bars on the figure. The total margin consisting of two components is marked on the top of the bars.

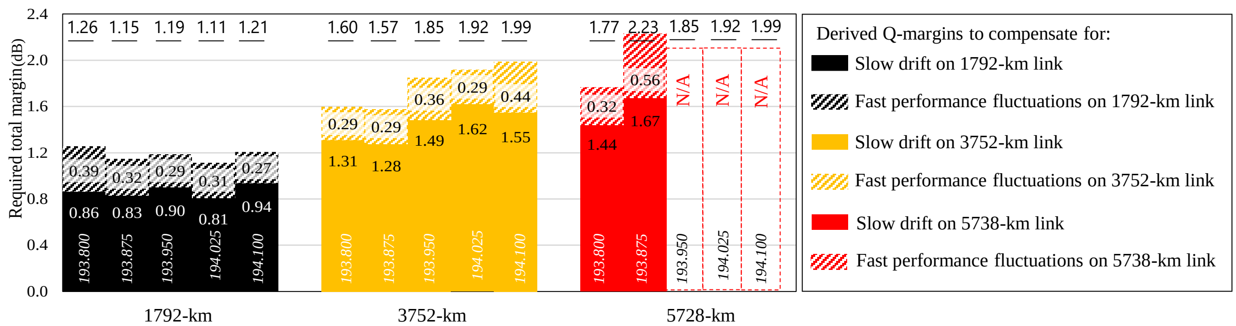

Fig. 3 Total margin requirements based on 6×σ for link lengths and individual transceivers

Fig. 3 shows, that the margins required by the transceivers on the shortest link are within close range regardless of the high variance in the performance. This may be caused by almost equal Rx powers on the link due to pre-emphasis. However, as the impact from pre-emphasis is lessened with the growing link length, the deviation in margin requirements is also increasing. This means, OSaaS operation in a single wide-band MC configuration over Raman-enabled long-haul networks requires margin adjustments to compensate for wavelength-dependent variations.

4. Conclusions

Optical Spectrum as a Service has become an attractive service model for customers demanding high capacity and flexible network resources. To achieve high service availability together with low-margins operation [4], the services operated within OSaaS should be continuously monitored. This enables end-users to verify that the deployed margins are feasible and adjust them to network aging or other impairments that may impact the end-to-end service continuity. Furthermore, if used together with Machine Learning, it can equip the end-user with an independent warning system for anomaly detection in the network [5].

In this work, we have collected long-term telemetry data from a 375-GHz wide-band OSaaS consisting of five 200-Gbit/s DP-QPSK 69-GBd signals operated on a Raman-enabled pan-European network. We have derived the standard deviation-based service margins for slow performance drift and fast performance fluctuations for each link. Based on the worst-performing transceiver, the maximum total margin of 1.26, 1.99 and 2.23 dB is required to facilitate stable operation of the commissioned services within a wide-band OSaaS over 1792-km, 3752-km and 5738-km links, respectively.

5. Acknowledgement. The work has been partially funded by the German Ministry of Education and Research (BMBF), in the project AI-NET-PROTECT (CELTIC-NEXT Grant No. 16KIS1279K).

6. References

[1] K. Kaeval et al., "Characterization of the optical spectrum as a service," in JOCN, 14(5), pp. 398-410.
[2] M. Filer et al., "Low-margin optical networking at cloud scale [Invited]," in JOCN, 11(10), pp. C94-C108
[3] ITU-T Rec. G.807, [Online, last accessed : October 2022] https://www.itu.int/rec/T-REC-G.807/en
[4] Y. Pointurier, "Design of low-margin optical networks," JOCN, vol. 9, no. 1, pp. A9-A17, Jan. 2017
[5] I. Dick, "Machine Learning based Fault Detection Algorithms or Long-Haul Elastic Optical Networks", Master Thesis @ TU Munich