Chatty ROADMs: Streaming Telemetry with Open Source Software and Open Hardware

Jan Kundrát(1)(2), Václav Kubernát(1), Tomáš Pecka(1)(3), Ondřej Havliš(1), Martin Šlapák(1), Jaroslav Jedlinský(1), Josef Vojtěch(1)

(1) CESNET, Prague, Czechia,
(2) Charles University, Faculty of Mathematics and Physics, Prague, Czechia,
(3) Czech Technical University in Prague, Faculty of Information Technology, Department of Theoretical Computer Science, Prague, Czechia,
jan.kundrat@cesnet.cz

Abstract *An interactive demo presents live, real-time optical performance monitoring of the optical spectrum from an Open Optical Line System. Built with open source software, open hardware, and using standard YANG tooling, we achieve a sub-GHz resolution and a sub-second refresh rate over the C-band.*

Introduction

In the past years, we have witnessed an ongoing shift towards more open, disaggregated optical networks and Optical Line Systems (OLS)[1]. Organizations such as the Telecom Infra Project[2], OpenROADM[3], Open Networking Foundation[4] and OpenConfig[5] have contributed towards this trend with standardization effort, interoperability guidelines[6], as well as APIs[7],[8], tooling improvements[9],[10], and demonstrations[11]. Despite these efforts, many networks are still being designed with the OLS provided by a single vendor, often as a turnkey solution.

One of the challenges in operating a successful multi-vendor network is *monitoring*. In absence of a third-party integrator, operators must rely on their in-house expertise when debugging operational issues. Once the network does not quite run as smoothly as it was designed to, a trove of continuously updated performance records becomes a priceless resource[12],[13].

In this paper we build upon our recently published work[14] on efficient streaming telemetry. Achieving a truly real-time telemetry has been demonstrated before for transponders[15], however, our work presents a significant improvement in latency of OLS-based spectrum measurement. We present an architecture where the OLS Network Elements (NEs) provide a continuous stream of data, measuring the optical spectrum at least once a second, at a high resolution, and using a common, off-the-shelf software for telemetry storage and analysis.

978-1-6654-3868-1/21/$31.00 ©2021 IEEE

Demo Description

In the past, DWDM monitoring systems often provided 15-minute rolling averages, or the telemetry capture was otherwise throttled[16]. To improve upon this and deliver $\mathcal{O}(1s)$ telemetry refresh rates, we used open source software stack with roots in the datacenter world, cloud, and Infrastructure as a Service (IaaS) and Software as a Service (SaaS) offerings.

The telemetry data flow is shown on Figure 1. The physical measurements are performed using a commercially available OCM (Optical Channel Monitor) built into the ROADMs, and from the embedded amplifiers. After scanning, the data are retrieved over a relatively high-speed SPI bus and processed by the application code. Once read via software, the data are transformed into an in-memory database holding YANG-formatted data in sysrepo[17]. A RESTCONF-ish server[18] running on the network elements *subscribes* to changes in the data that are managed in sysrepo, and upon a change of the data, it generates an event that is compliant with the YANG Push[19] standard. The notification is made available over a RESTCONF channel, using HTTP's *Server-Sent Events*, to any subscribed client (Listing 1).

The telemetry updates are processed by a small proxy[20] written in Python using the `asyncio` HTTP client library. Once a new `push-update` is issued by the RESTCONF server, the data are processed and pushed to a Time-Series Database (TSDB) via the OpenMetrics[21] protocol. The TSDB is implemented via VictoriaMetrics[22], a high-performance in-memory data store which

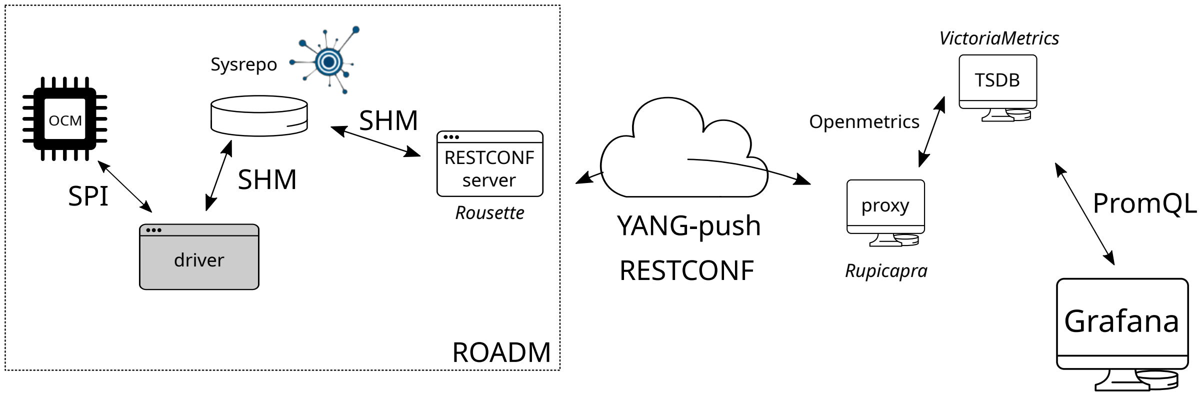

Fig. 1: Telemetry data from ROADMs to the visualization software

provides Prometheus-compatible querying API (PromQL[23]) and ingestion points. Final data visualization is implemented via Grafana using a standard PromQL connector.

The optical part of the demo is shown in Figure 2 and Figure 3. The full system consists of a pair of transponders provided by vendor A, and an OLS system[24] provided by vendor B. The OLS system comprises three ROADMs of degree two, along with a pair of Add/Drop modules[25], one per each transponder. A smaller-scale, portable setup uses a reduced set of hardware while still demonstrating the same rate of streamed telemetry updates.

The monitoring system focuses on L0 metrics. In order to demonstrate results that are applicable to the real world, the monitored parameters were set to those commonly available with commercial ROADMs and amplifiers. As such, the metrics track per-channel (Media Channel, MC) optical power, as well as the total C-band optical power in case of amplifiers. In addition, a high-resolution OCM which provides data at the resolution of 312.5 MHz was used in one device, yielding more than 15,000 data points per scan. Other ROADM modules were configured with an OCM which returns circa 2,000 data points per C-band.

The fine resolution of the OCM imposed a significant load on the monitoring infrastructure. A straightforward approach with YANG-level `list` and `leaf-list` statements identified bottlenecks in processing of YANG-level diffs. Parts of this problem were addressed via additional optimization of the YANG software stack, however, the root of the problem dates back to the YANG 1.0 times when the `leaf-list` was defined with the semantics of an *ordered set*, not a list. We solved this by using the `anydata` statement which effectively bypasses an $\mathcal{O}(n^2)$ data path.

The nature of the monitoring data, being *vector* measurements rather than simple scalars, and the related increase of the number of time series entries to process when rendering a spectrum scan visually, led to non-negligible latency at the Grafana level. The actual monitoring infrastructure (except the GUI layer), however, showed no symptoms of performance issues even when running on a virtual machine (VM) on an obsolete server hardware with spinning disks and no SSDs.

The transport channel we have chosen imposes a certain overhead to the communication. In future, the overhead can be reduced by switching to CBOR[26] serialization of YANG data and to the UDP transport[27] of notification messages. That said, HTTP compression is an effective tool for reducing the data bandwidth consumption.

ECOC Relevance and Conclusion

The technical part of this proposed demo has been previously described[14]. However, the format of a journal article did not allow interactive presentation and interaction with the audience.

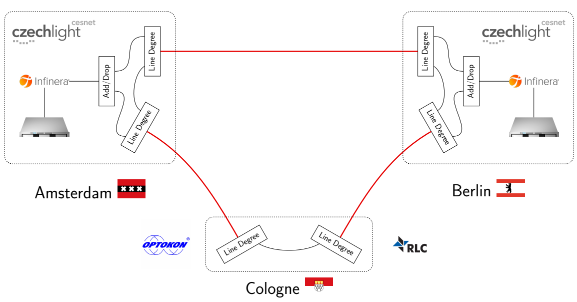

Fig. 2: Optical setup

Fig. 3: Portable demo

```
$ curl -v http://add-drop-spi/telemetry/optics
> GET /telemetry/optics HTTP/1.1
> Host: add-drop-spi
> User-Agent: curl/7.71.1
> Accept: */*
>
< HTTP/1.1 200 OK
< Date: Wed, 02 Jun 2021 00:55:18 GMT
< Access-Control-Allow-Origin: *
< Content-Type: text/event-stream
< Transfer-Encoding: chunked
< Server: nghttpx
< Via: 2 nghttpx
<
data: {
  "ietf-restconf:notification": {
    "eventTime":"2021-06-02T00:55:18.891074890-00:00",
    "ietf-yang-push:push-update": {
      "datastore-contents":{
        ...
        "czechlight-roadm-device:aggregate-power": {
          "common-in":"-34.51",
          "common-out":"-29.74"
        },
        ...
      }
    }
  }
}
```

Listing 1: YANG `push-update` example

Fig. 4: ROADM setup in the lab

For ECOC’2021, we propose a live demonstration to the attendees who will have the opportunity to explore the console of the monitoring tool (Grafana) and watch ad-hoc queries being set up on the fly. Examples of such queries are combining the parameters from the Optical Supervisory Channel (OSC) with the power level measurements at the ROADM’s booster and preamplifiers, and monitoring the attenuation of a fibre link that way. It will also be possible to watch the metrics change in real time in response to, e.g., fiber cuts.

Depending on the COVID-19 situation which we cannot predict at the demo submission time, we have prepared contingency options. Our preference is to bring a reduced-scale demo on site (two one-degree ROADMs, an Add/Drop, and a non-coherent signal), in total a 4U rack (Figure 3). As a backup option, we can demonstrate our work via a live connection to our lab in Prague with 12U worth of hardware, fibre spools, etc. (Figure 4).

Acknowledgements

We wish to thank Radek Krejčí and Michal Vaško for their contribution towards the presented demo.

References

[1] C. Xie, L. Wang, L. Dou, *et al.*, “Open and disaggregated optical transport networks for data center interconnects [Invited]”, *J. Opt. Commun. Netw.*, vol. 12, no. 6, pp. C12–C22, Jun. 2020. DOI: 10.1364/JOCN.380721.

[2] S. Barguil, V. Lopez, and J. P. Fernandez-Palacios Gimenez, “Towards an Open Networking Architecture”, in *2020 International Conference on Optical Network Design and Modeling (ONDM)*, 2020, pp. 1–3. DOI: 10.23919/ONDM48393.2020.9133038.

[3] M. Birk, O. Renais, G. Lambert, *et al.*, “The Open-ROADM initiative [Invited]”, *IEEE/OSA Journal of Optical Communications and Networking*, vol. 12, no. 6, pp. C58–C67, 2020. DOI: 10.1364/JOCN.380723.

[4] A. Campanella, B. Yan, R. Casellas, *et al.*, “Reliable Optical Networks With ODTN: Resiliency and Fail-Over in Data and Control Planes”, *J. Lightwave Technol.*, vol. 38, no. 10, pp. 2755–2764, May 2020. DOI: 10.1109/JLT.2020.2983587.

[5] A. Sgambelluri, F. Paolucci, and F. Cugini, “Telemetry-driven validation of operational modes in OpenConfig disaggregated networks”, in *45th European Conference on Optical Communication (ECOC 2019)*, 2019, pp. 1–4. DOI: 10.1049/cp.2019.0960.

[6] A. Triki, E. L. Rouzic, O. Renais, *et al.*, “Open-Source QoT Estimation for Impairment-Aware Path Computation in OpenROADM Compliant Network”, in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–3. DOI: 10.1109/ECOC48923.2020.9333351.

[7] R. Vilalta, C. Manso, N. Yoshikane, *et al.*, “Experimental evaluation of control and monitoring protocols for optical SDN networks and equipment [Invited Tutorial]”, *IEEE/OSA Journal of Optical Communications and Networking*, vol. 13, no. 8, pp. D1–D12, 2021. DOI: 10.1364/JOCN.424631.

[8] R. Vilalta, N. Yoshikane, R. Casellas, *et al.*, “Controlling and Monitoring Optical Network Equipment in Optical SDN Networks”, in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4. DOI: 10.1109/ECOC48923.2020.9333345.

[9] G. Bosco, A. Carena, M. R. Zefreh, P. Poggiolini, and F. Forghieri, “Advances in Modeling and Mitigation of Nonlinear Effects in Uncompensated Coherent Optical Transmission Systems”, in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4. DOI: 10.1109/ECOC48923.2020.9333258.

[10] A. Ferrari, K. Balasubramanian, M. Filer, *et al.*, “Softwarized Optical Transport QoT in Production Optical Network: a Brownfield Validation”, in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4. DOI: 10.1109/ECOC48923.2020.9333280.

[11] J. Kundrát, A. Campanella, E. L. Rouzic, *et al.*, "Physical-Layer Awareness: GNPy and ONOS for End-to-End Circuits in Disaggregated Networks", in *Optical Fiber Communication Conference (OFC) 2020*, Optical Society of America, 2020, M3Z.17. DOI: 10.1364/OFC.2020.M3Z.17.

[12] K. S. Mayer, J. A. Soares, R. P. Pinto, C. E. Rothenberg, D. S. Arantes, and D. A. A. Mello, "Soft Failure Localization Using Machine Learning with SDN-based Network-wide Telemetry", in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4. DOI: 10.1109/ECOC48923.2020.9333313.

[13] C. Natalino, A. di Giglio, M. Schiano, and M. Furdek, "Root Cause Analysis for Autonomous Optical Networks: A Physical Layer Security Use Case", in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4. DOI: 10.1109/ECOC48923.2020.9333141.

[14] J. Kundrát, V. Kubernát, T. Pecka, *et al.*, "Opening up ROADMs: Streaming Telemetry [Invited]", *IEEE/OSA Journal of Optical Communications and Networking*, vol. 13, no. 10, 2021. DOI: 10.1364/JOCN.425167.

[15] R. Vilalta, C. Manso, N. Yoshikane, *et al.*, "Telemetry-enabled Cloud-native Transport SDN Controller for Real-time Monitoring of Optical Transponders Using gNMI", in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4. DOI: 10.1109/ECOC48923.2020.9333143.

[16] S. Xu, Y. Hirota, M. Shiraiwa, *et al.*, "Emergency OPM Recreation and Telemetry for Disaster Recovery in Optical Networks", *J. Lightwave Technol.*, vol. 38, no. 9, pp. 2656–2668, May 2020. DOI: 10.1109/JLT.2020.2969471.

[17] CESNET, *Sysrepo: storing and managing yang-based configurations for unix/linux applications*, https://www.sysrepo.org/, Accessed: 2021-03-13.

[18] ——, *Rousette: an almost-restconf server*, https://github.com/CESNET/rousette, Accessed: 2021-03-13, 2021.

[19] A. Clemm and E. Voit, "Subscription to YANG notifications for datastore updates", RFC Editor, RFC 8641, Sep. 2019.

[20] CESNET, *Rupicapra: reading the yang push telemetry stream*, https://github.com/CESNET/rupicapra, Accessed: 2021-03-13, 2021.

[21] R. Hartmann, B. Kochie, B. Brazil, and R. Skillington, "OpenMetrics, a cloud-native, highly scalable metrics protocol", IETF Secretariat, Internet-Draft draft-richih-opsawg-openmetrics-00, Nov. 2020.

[22] *Victoria Metrics: The Aspiring open source Monitoring Solution*, https://victoriametrics.com/, Accessed: 2021-03-13, 2018.

[23] N. Sabharwal and P. Pandey, "Working with Prometheus Query Language (PromQL)", in *Monitoring Microservices and Containerized Applications: Deployment, Configuration, and Best Practices for Prometheus and Alert Manager*. Berkeley, CA: Apress, 2020, pp. 141–167, ISBN: 978-1-4842-6216-0. DOI: 10.1007/978-1-4842-6216-0_5.

[24] J. Kundrát, O. Havliš, J. Jedlinský, and J. Vojtěch, "Opening up ROADMs: Let Us Build a Disaggregated Open Optical Line System", *Journal of Lightwave Technology*, vol. 37, no. 16, pp. 4041–4051, Aug. 2019, ISSN: 0733-8724. DOI: 10.1109/JLT.2019.2906620.

[25] J. Kundrát, O. Havliš, J. Radil, J. Jedlinský, and J. Vojtěch, "Opening up ROADMs: a filterless add/drop module for coherent-detection signals", *IEEE/OSA Journal of Optical Communications and Networking*, vol. 12, no. 6, pp. C41–C49, 2020. DOI: 10.1364/JOCN.388893.

[26] M. Veillette, I. Petrov, and A. Pelov, "CBOR encoding of data modeled with YANG", IETF Secretariat, Internet-Draft draft-ietf-core-yang-cbor-15, Jan. 2021.

[27] G. Zheng, T. Zhou, T. Graf, P. Francois, and P. Lucente, "UDP-based transport for configured subscriptions", IETF Secretariat, Internet-Draft draft-ietf-netconf-udp-notif-01, Nov. 2020.