



Demonstration of DRL-based intelligent spectrum management over a T-API-enabled optical network digital twin

Downloaded from: <https://research.chalmers.se>, 2026-04-05 01:53 UTC

Citation for the original published paper (version of record):

Etezadi, E., Natalino Da Silva, C., Karunakaran, V. et al (2023). Demonstration of DRL-based intelligent spectrum management over a T-API-enabled optical network digital twin. IET Conference Proceedings, 2023(34): 1730-1733.
<http://dx.doi.org/10.1049/icp.2024.1801>

N.B. When citing this work, cite the original published paper.

Demonstration of DRL-based intelligent spectrum management over a T-API-enabled optical network digital twin

Ehsan Etezadi⁽¹⁾, Carlos Natalino⁽¹⁾, Vignesh Karunakaran⁽²⁾, Renzo Diaz⁽³⁾, Anders Lindgren⁽³⁾, Stefan Melin⁽³⁾, Achim Autenrieth⁽²⁾, Lena Wosinska⁽¹⁾, Paolo Monti⁽¹⁾, Marija Furdek⁽¹⁾

⁽¹⁾ Department of Electrical Engineering, Chalmers University of Technology, 412 96 Gothenburg, Sweden, ehsanet@chalmers.se

⁽²⁾ Adtran, Fraunhoferstr. 9a, 82152 Martinsried/Munich, Germany.

⁽³⁾ Telia Company AB, 169 94 Solna, Sweden.

Abstract

This demonstration showcases the applicability and benefits of a deep reinforcement learning (DRL) agent for spectrum defragmentation in a realistic deployment. This is achieved by integrating the DRL agent with the operations of a carrier-grade optical network digital twin via standard T-API messages. © 2023 The Author(s)

Introduction

One of the main challenges in dynamic elastic optical networks (EONs) is spectrum fragmentation (SF) which stems from the discrepancy between the incoming connection requests and the available spectral gaps. SF leads to inefficient use of the spectrum, degrading the performance of EON in terms of service blocking ratio (SBR)^[1]. Spectrum defragmentation (SD) is a way to consolidate the spectrum usage by reconfiguring a subset of connections, thus reducing gaps unsuitable for incoming connectivity services. Numerous SD approaches rely on, e.g., threshold-based heuristic algorithms^[2] or integer linear programming models^[3], typically aimed at SBR minimisation. Such methods may require bespoke threshold configuration or take long to find a solution, which limits their flexibility and applicability in dynamic service provisioning scenarios.

Intelligent and adaptable techniques, such as those based on machine learning (ML), are needed to meet the network operators' quest for efficient and automated network management. DeepDefrag^[4], a recently proposed SD framework based on deep reinforcement learning (DRL), has been shown to outperform existing deterministic algorithms in SBR minimisation. DeepDefrag performs proactive SD by deciding on the reconfiguration timing, the concerned subset of connections, and their target spectrum.

Integrating ML-based techniques in real-world optical network deployments is challenging due to, among other, potential mismatch between the data these techniques require and information made available by vendor-specific management

tools. Optical network disaggregation addresses this issue by defining, among other elements, a set of standard application programming interfaces (APIs) that allow operators to interact with network elements^[5]. Transport API (T-API), an example of such standards, supports a hierarchical software-defined networking (SDN) architecture that fits multi-vendor/multi-domain scenarios. T-API is regarded as a promising standard for different use cases, including connectivity service creation over dense wavelength-division multiplexing (DWDM) networks^[6] and enabling quantum encryption of optical end-to-end services^[7]. However, so far, the potential benefits of SD techniques have not been validated in T-API systems, especially those using carrier-grade T-API implementations. Such implementations can be provided by, e.g., a digital twin^[8], which mirrors the behaviour of real devices, allowing a realistic and real-time system performance evaluation without the prohibitive overhead of testing on real devices.

In this demonstration, we develop a new defragmentation module that uses standard T-API messages to realise SD decisions taken by a DRL agent, i.e., DeepDefrag. We demonstrate the module's capabilities through a dashboard that enables the audience to parameterise network operation settings, view the fragmented network state, observe the DRL-based SD decisions, and inspect their realisation over a carrier-grade digital twin. The DRL agent intelligently decides when to trigger defragmentation, selects the connections and the order of their reallocation, and finds the target spectrum slots.

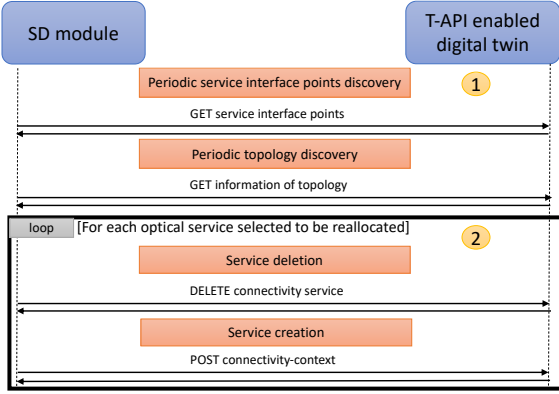


Fig. 1: Communication between the Spectrum defragmentation (SD) module and the digital twin.

Workflow

Figure 1 illustrates the workflow of the proposed demonstration, including the message exchange between the SD module and the T-API-enabled digital twin. In the following, we describe each message, highlighting the associated T-API use case^[9].

In phase 1, the SD module periodically requests information about existing connectivity services (including their unique identifiers) and topology from the digital twin (use cases 0a and 0b). The defragmentation module uses this information as input to the DRL agent, which, based on the current network state, decides whether to initiate an SD cycle or not.

Phase 2 starts when the DRL agent initiates an SD cycle. The DRL agent iteratively selects a connectivity service for reallocation and the target spectrum. Note that the path does not change during defragmentation. This process is repeated until the DRL agent stops the SD cycle.

Connectivity services are reallocated following a *break-before-make* approach. The service selected for reallocation is removed (use case 10), after which it is re-established by specifying the nodes, links and target spectrum slots. The process of establishing a connectivity service, which traverses specific nodes or links and occupies specific spectrum slots, is defined in use cases 2c, 3a, and 3b, respectively. The *break-before-make* approach allows our defragmentation module to take advantage of defragmentation solutions that overlap with the spectrum currently used by the service under reconfiguration.

Demonstration implementation

Figure 2 illustrates the deployment adopted in this demonstration. The defragmentation module is implemented specifically for this demonstration using Python. It uses the Optical RL-Gym^[10]

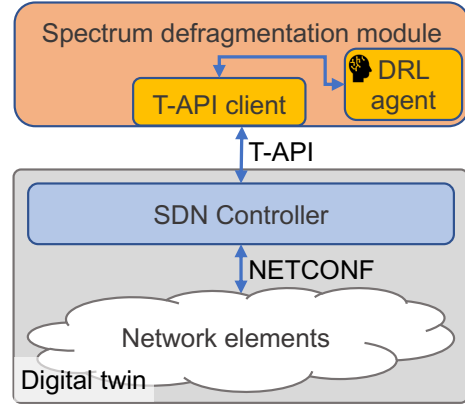


Fig. 2: Demonstrator architecture.




Fig. 3: Simple network example with 7 connectivity services.

to generate connectivity service requests following a Poisson process. The aim is to obtain a representation of the network state resulting from long-term operation (i.e., steady state representation). The DeepDefrag^[4] DRL agent making the defragmentation decisions during the demonstration is trained separately beforehand for practical purposes.

The digital twin of the optical network is implemented by mirroring each optical network element instance in its digital form, i.e., using the same operating system running over virtual machines. The deployed digital twin is controlled by a production-grade SDN domain controller, supporting T-API in the northbound and NETCONF in the southbound interface. The interaction between the defragmentation module and the digital twin uses the T-API specification version 2.1^[9]. The defragmentation module and dashboard run on the demonstrator computer connected to the digital twin located at a remote lab through a secure channel.

Demonstration storyline

The demonstration begins with an empty network and fully unassigned spectrum. In the first part, the demo focuses on simulating a fragmented network state. The audience can select the parameters for generating the connectivity service requests (e.g., inter-arrival time, holding time, bandwidth) in addition to the simulation

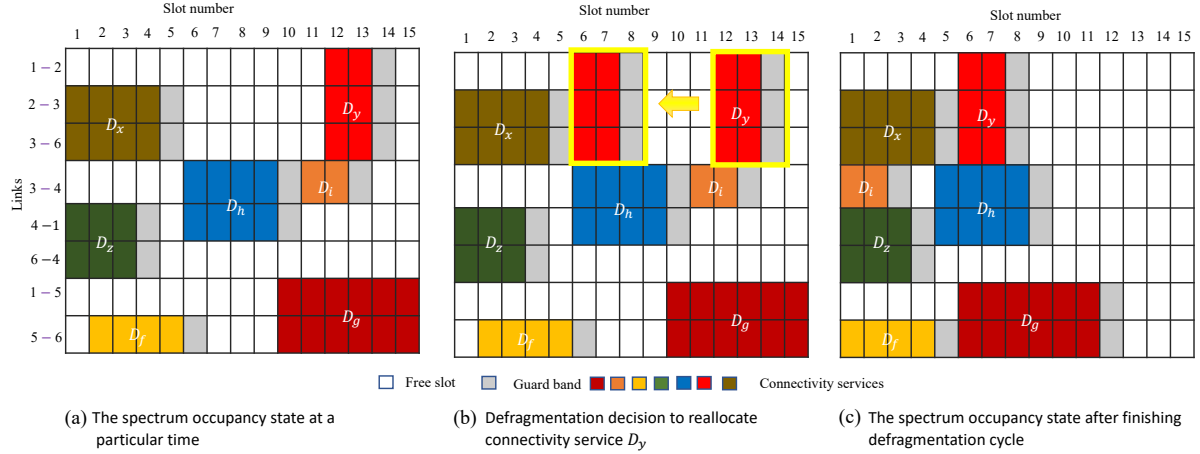


Fig. 4: The spectrum grid shown on the dashboard at different stages.

run time. The fragmented network state at the end of the simulation is consolidated into the digital twin. Fig. 3 illustrates a simple network topology and a snapshot of active connectivity services. Fig. 4(a) shows the fragmented spectrum resulting from the simulated arrivals and departures consolidated into the digital twin.

The second part of the demo consists in invoking the DeepDefrag agent and graphically showing its decisions about services selected for reallocation and their target spectrum. Fig. 4(b) illustrates a hypothetical decision where the agent, based on the spectrum state from Fig. 4(a), decides to reallocate connectivity service D_y from slot 12 to 6.

The audience can inspect the execution of individual decisions within a defragmentation cycle or advance to the end of the cycle. Fig. 4(c) illustrates the spectrum state at the end of a hypothetical defragmentation cycle that started with the state in Fig. 4(a). Note that 5 out of 7 connectivity services were reconfigured, representing an aggressive defragmentation cycle. For the demo, a more realistic number of slots is set up, and the number of active services will depend on the parameters set by the audience (e.g., inter-arrival time, holding time, and simulation time).

In addition to the spectrum visualisation illustrated in Fig. 4, various spectrum fragmentation metrics (e.g., number of cut^[11], and root of sum of squares^[12]) that vary with time are displayed to the audience, not included here due to space constraints. The impact of provisioning, departures, and defragmentation of connectivity services on the variations of these metrics is presented to the audience. The audience can also inspect the messages exchanged by the module and the digital twin, revealing how T-API can be leveraged to

implement the proposed approach.

Innovation

This demonstration is the first to take advantage of DRL to perform intelligent spectrum management over a T-API-enabled carrier-grade optical network deployment, exemplified here by a digital twin. The demonstration serves not only as a proof-of-concept of defragmentation operations over T-API, but also showcases a dashboard where the audience configures connectivity service parameters, and observes how the DRL agent intelligently defragments the spectrum. This work can foster discussions and spark interest in the ECOC community regarding the real-world implementation of intelligent spectrum management strategies in optical networks and their realisation using currently available interfaces.

Conclusions

This paper presents the first experimental and interactive demonstration of a proactive spectrum defragmentation module for elastic optical networks using the Open Networking Foundation (ONF) Transport API standard over a digital twin. The algorithm uses the merits of deep reinforcement learning to find the best set of actions based on the network condition.

Acknowledgements

Work partially supported by Sweden's innovation agency VINNOVA, within the framework of the EUREKA cluster CELTIC-NEXT project AI-NET-PROTECT (2020-03506), and Vetenskapsrådet (2019-05008).

References

- [1] W. Shi, Z. Zhu, M. Zhang, and N. Ansari, "On the effect of bandwidth fragmentation on blocking probability in elastic optical networks", *IEEE Transactions on Communications*, vol. 61, no. 7, pp. 2970–2978, 2013. DOI: [10.1109/TCOMM.2013.053013.120853](https://doi.org/10.1109/TCOMM.2013.053013.120853).
- [2] J. Comellas, L. Vicario, and G. Junyent, "Proactive defragmentation in elastic optical networks under dynamic load conditions", *Photonic Network Communications*, vol. 36, no. 1, pp. 26–34, 2018. DOI: [10.1007/s11107-018-0767-7](https://doi.org/10.1007/s11107-018-0767-7).
- [3] M. Zhang, C. You, and Z. Zhu, "On the parallelization of spectrum defragmentation reconfigurations in elastic optical networks", *IEEE/ACM Transactions on Networking*, vol. 24, no. 5, pp. 2819–2833, 2016. DOI: [10.1109/TNET.2015.2487366](https://doi.org/10.1109/TNET.2015.2487366).
- [4] E. Etezadi, C. Natalino, R. Diaz, *et al.*, "DeepDefrag: A deep reinforcement learning framework for spectrum defragmentation", in *GLOBECOM 2022-2022 IEEE Global Communications Conference*, IEEE, 2022, pp. 3694–3699.
- [5] F. Paolucci and A. Sgambelluri, "Telemetry in disaggregated optical networks", in *International Conference on Optical Network Design and Modeling (ONDM)*, 2020, pp. 1–3. DOI: [10.23919/ONDM48393.2020.9132992](https://doi.org/10.23919/ONDM48393.2020.9132992).
- [6] V. Lopez, R. Vilalta, V. Uceda, *et al.*, "Transport API: A solution for SDN in carriers networks", in *European Conference on Optical Communication (ECOC)*, 2016, pp. 1–3.
- [7] V. López, A. Gomez, A. Aguado, *et al.*, "Extension of the ONF transport API to enable quantum encryption in end-to-end services", in *European Conference on Optical Communication (ECOC)*, 2019.
- [8] D. Jones, C. Snider, A. Nassehi, J. Yon, and B. Hicks, "Characterising the digital twin: A systematic literature review", *CIRP Journal of Manufacturing Science and Technology*, vol. 29, pp. 36–52, 2020.
- [9] ONF, *Tr-547, functional requirements for transport api*, <https://opennetworking.org>, Jul. 2020.
- [10] C. Natalino and P. Monti, "The Optical RL-Gym: An open-source toolkit for applying reinforcement learning in optical networks", in *International Conference on Transparent Optical Networks (ICTON)*, 2020. DOI: [10.1109/ICTON51198.2020.9203239](https://doi.org/10.1109/ICTON51198.2020.9203239).
- [11] Y. Yin, M. Zhang, Z. Zhu, and S. B. Yoo, "Fragmentation-aware routing, modulation and spectrum assignment algorithms in elastic optical networks", in *2013 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC)*, IEEE, 2013, pp. 1–3.
- [12] P. Lechowicz, M. Tornatore, A. Włodarczyk, and K. Walkowiak, "Fragmentation metrics in spectrally-spatially flexible optical networks", in *2019 International Conference on Optical Network Design and Modeling (ONDM)*, Springer, 2020, pp. 235–247.