Integrated OAM and Transport SDN Architecture for Automatic and Dynamic Restoration of Signal-Degraded Connections within Flexi-Grid Optical Networks

R. Martínez, R. Casellas, J. M. Fabrega, R. Vilalta, R. Muñoz, F. J. Vilchez, L. Nadal, M. Svaluto-Moreolo, A. Villafranca* and P. Sevillano *

Centre Tecnòlogic de Telecomunicacions de Catalunya (CTTC), Castelldefels (Barcelona), Spain

** Aragon Photonics Labs, S. L., Zaragoza, Spain*

Tel: (+34) 93 645 2900. e-mail: ricardo.martinez@cttc.es

ABSTRACT

The programmability of upcoming transport optical networks is handled by a centralized Transport SDN (T-SDN) controller. The T-SDN controller dynamically computes and selects the resources to automatically accommodate optical connection requests by configuring the involved network elements and devices (i.e., optical switches and transceivers). Additionally, the T-SDN controller should provide management functions to ensure high reliability and correct service status in terms of QoS, QoT, failure detection, etc. during the whole connection duration. These management functions are typically addressed by the operation, administration and maintenance (OAM) Handler. Thereby, the OAM Handler interacts with the monitoring system to gather notifications about connection failures or degradations and trigger the required actions to preserve / recover the services. This work first overviews the designed and implemented architecture integrating required OAM functions into a deployed T-SDN controller for flexi-grid optical networks. Next, two restoration routing approaches are presented, which are compared with respect to the blocked bandwidth ratio (BBR), restorability and average number of used sub-transceivers. The experimental performance is made varying both the traffic load and the generation of signal degraded events of existing connections, forcing their automatic restoration.

Keywords: T-SDN, PCECC, Flexi-Grid Optical Network, OAM, Dynamic Restoration.

1. INTRODUCTION

Data traffic in transport network segments is forecasted to continuously increase due to the proliferation of cloud services such as Content Delivery Networks (CDNs) [1]. This is challenging network operators to seek for feasible and cost-effective technological solutions for dynamically and flexibly transporting heterogeneous data services, by demanding different capacities and network requirements (e.g., latency, availability, etc.). To deal with such a challenge a candidate solution relies on adopting programmable elastic optical network (EON) [2]. EONs allow a finer granularity of the optical grid (*flexi-grid*) compared to traditional *fixed-grid* DWDM systems. This leads to improve the optical spectrum utilization taking advantage of the (sliceable) bandwidth-variable transponders (SBVTs). SBVTs accommodate heterogeneous data flow connections requiring different bit rates and distances, leveraging the inherited flexibleness for tailoring transceivers capabilities: modulation format (MF), number of subtransponders, spectrum width, etc. [3]. EON programmability is made by a centralized Transport SDN (T-SDN) controller, which dynamically computes feasible paths and selects the optical resources triggering a Routing, Spectrum and Modulation Assignment (RSMA) algorithm [3]. The resulting path and resources (i.e., optical switches and SBVTs) are then configured by the T-SDN controller.

The implementation of a T-SDN controller is based on the architecture of a Path Computation Element Central Controller (PCECC) [4][5]. The Southbound Interface (SBI) between the T-SDN controller and the network agents is realized via the PCE protocol (PCEP) with PCECC extensions [4] to configure the data plane. The T-SDN controller can be enhanced to support management functions to provide high reliability and correct service status (i.e., QoT, failure detection, etc.). Such functionalities are addressed by the OAM Handler [6]. The OAM interacts with the data plane devices (e.g., monitoring system) to receive notifications about connection failures or degradations, and accordingly triggers actions to preserve / restore the affected services. To this end, it is needed a tight control collaboration between both the OAM Handler and T-SDN controller.

It is assumed a distributed monitored system wherein each optical switch is equipped with a high-resolution optical spectrum analyzer allowing in-band OSNR measurements. The measurement element is a commercial product called BOSA [7]. The goal is to provide OSNR measurements for each existing connection, and to report them to the OAM Handler. By doing so, the OAM checks whether a connection's OSNR level is acceptable according to the bit-rate and MF. If not, the connection is declared as signal-degraded and needs to be restored by the T-SDN controller. This paper extends the work in [8] detailing an exhaustive performance evaluation of two proposed restoration routing approaches: i) restoring over the same spatial path but selecting a lower MF; ii) compute a complete link-disjoint path. The performance metrics for the comparison are: the blocked bandwidth ratio (BBR), the average number of used subtransponders and the restorability. The numerical results are obtained for different traffic loads as well as varying the generation of signal-degradation events.

2. TRANSPORT SDN CONTROLLER AND OAM HANDLER INTEGRATION

Fig. 1. a. shows the integration of the T-SDN controller and the deployed OAM Handler. Upon receiving a connection request (specifying the endpoints and data rate in Gb/s), the T-SDN controller uses the topology and network resource information, updated via BGP-LS and stored in the Traffic-Engineering Database (TED), to compute the spatial path (i.e., nodes and links) and select the SBVT subtransponders, frequency slot and MF. The applied RSMA algorithm [5] to do that is an iterative distance-adaptive mechanism, prioritizing the use of the most advance MFs as long as a feasible path (maximum permitted distance) is found. The computed path and resources are then allocated via a PCEP SBI [5]. Finally, all active connections are stored in the Label-switched Path Database (LSPDB).

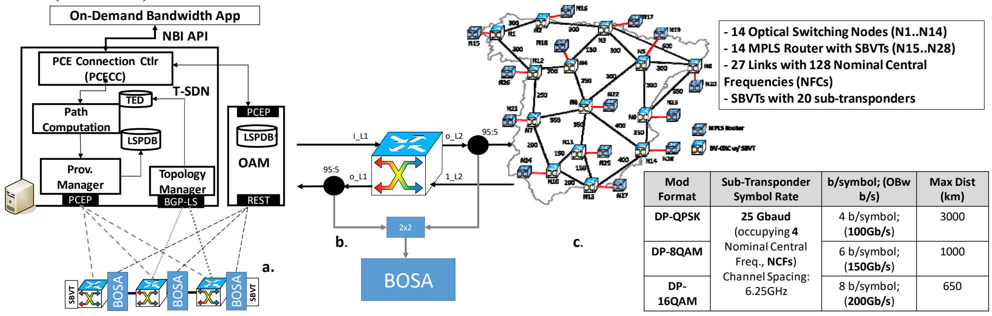

Mod Format	Sub-Transponder Symbol Rate	b/symbol; (OBw b/s)	Max Dist (km)
DP-QPSK	25 Gbaud (occupying 4 Nominal Central Freq., NCFs) Channel Spacing: 6.25GHz	4 b/symbol; (100Gb/s)	3000
DP-8QAM		6 b/symbol; (150Gb/s)	1000
DP-16QAM		8 b/symbol; (200Gb/s)	650

Figure 1. a. T-SDN Controller – OAM Handler Integration; b. Distributed Monitoring System; c. Flexi-Grid Network Topology (detailing link distances in km) and SBVT Capabilities

The OAM Handler may or may not be physically co-located with the T-SDN controller. Both elements interact via a PCEP interface providing: i) replication of LSPDB in the OAM Handler's LSPDB* to have connection details awareness; ii) degraded connection re-computation requests to T-SDN controller. The OAM Handler decision to restore a signal-degraded connection is triggered after processing the OSNR measurements sent by the distributed BOSA devices. As shown in Fig. 1.b., 5% of the optical signal on each node's outgoing link is processed by the BOSA attaining the aggregated in-band OSNR for each active flow. The measured OSNR is reported to the OAM Handler (e.g., REST interface) with the Node Id (IP address), Link Id and central frequency of the optical flow. With such BOSA notifications, the OAM Handler retrieves connection features from the LSPDB* such as the path and allocated resources as well as the MF used by the optical signal. This allows OAM resolving whether the received OSNR is below the expected threshold for that flow. If this occurs, the connection is declared as signal-degraded and needs to be restored. The OSNR threshold for each optical flow depends on a set of connection attributes such as the bit rate, MF, path distance and number of hops.

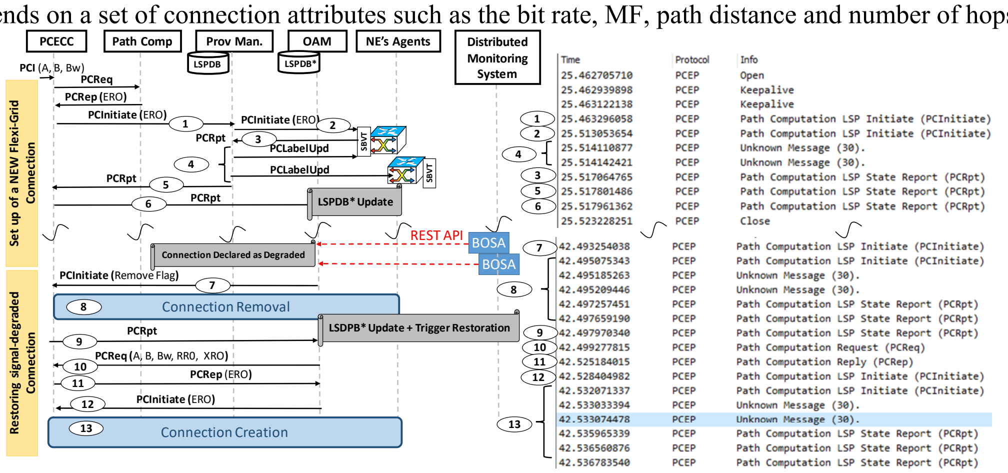

Figure 2. a. Workflow for restoring a signal-degraded connection; b. Captured PCEP messages

The PCEP interactions between the T-SDN controller and the OAM Handler are illustrated in a workflow (Fig.2.a.). This is divided into two parts. The first one (steps 1 – 6) regards to the setting up of a new flexi-grid connection. The T-SDN controller computes and configures the connection using the procedures of the PCECC architecture (T-SDN Controller). The second part addresses the restoration of a signal-degraded connection (steps 7 – 13). As aforementioned, BOSA devices report to the OAM Handler a connection measured OSNR level. If the connection restoration needs to be triggered, a break-before-make strategy is adopted. That is, the OAM Handler triggers the connection removal (via a PCEP PCInitiate message). Next, the T-SDN controller deallocates all the resources (aggregated in Step 8) and reports (using a PCRpt message) to the OAM Handler that such connection is to be removed from LSPDB* (step 9). After that, the OAM Handler requests a path computation (PCReq message) to the T-SDN controller for restoring the signal-degraded connection (Step 10). This message carries, besides the endpoints and the requested bandwidth, information (into the Record Route Object, RRO) about the route (i.e., spatial path, frequency slot and SBVT resources) of the degraded connection.

Additionally, the eXclude Route Object (XRO) is also added into the PCReq message to indicate the used MF. Both RRO and XRO information may be then used to constrain the restoration path computation according to the selected path computation approach. If a feasible route is found, this is sent back to the OAM Handler (Step 11) in the PCRep message with the new computed route and selected resources (formatted into the Explicit Route Object, ERO). The restored path is then instantiated by the OAM Handler sending a PCInitiate to the T-SDN controller with the new ERO (Step 12) which is set up as described in [5]. Fig. 2.b shows the validation of the above workflow showing the captured PCEP messages.

3. ON-LINE RSMA RESTORATION APPROACH

Two on-line RSMA restoration approaches (*A* and *B*) are proposed. *Approach A* tries to restore a signal-degraded connection using the same spatial path (specified in the PCReq RRO) as in the active connection but selecting a lower MF than that specified in the XRO information. The idea behind that is to exploit elastic capabilities of SBVTs and flexi-grid technologies where active connections are dynamically re-configured over the same path but occupying different amount of resources (i.e., optical spectrum and SBVT' subtransponders). If this fails (e.g., due to the lack of available resources or a lower MF), then *Approach A* computes a complete link-disjointed path. To do that, all links contained into the PCReq's RRO are pruned from the topology graph before triggering the RSMA algorithm. The RSMA algorithm computes paths satisfying both continuity and contiguity constraints and favours using the most advanced MFs when it is feasible. If both steps fail, the restoration cannot be completed, and the connection is blocked. *Approach B* focuses directly on computing a link-disjointed path using PCReq's RRO information. Of course, if link-disjointed path computation fails, the connection is blocked.

4. EXPERIMENTAL PERFORMANCE EVALUATION

The performance evaluation compares both *Approaches A* and *B* with respect to the BBR, the average number of used subtransponders and the restorability. The flexi-grid transport network formed by 14 optical switches is depicted in Fig. 1.c. It is also detailed the capabilities of both the flexi-grid links (i.e., number of NCFs and distances in km) and the SBVTs: number of subtransponders, supported MFs (i.e., DP-QPSK, DP-8QAM and DP-16-QAM), symbol rate (bit rate) and maximum permitted distance per MF. Each data point is obtained generating 5000 connection requests according to a Poisson process with mean inter-arrival time (*IAT_p*) set to 15s and the holding time (*HT*), exponentially modelled, is varied for different traffic loads (450, 600 and 950s). For each connection, source and destination nodes are randomly chosen requesting a bandwidth uniformly distributed as multiples of 100 Gb/s up to 500 Gb/s. The generation of a signal-degraded connection to be restored is randomly selected from those active LSPs contained in the LSPDB*. A failure (signal degradation) is modelled via a Poisson process with *IAT_r* varied from 50, 100, 150 and 200s.

Figure 3 shows the attained results. In general, regardless of the restoration approach, as expected, if the HT grows, the attained BBR and restorability are worsened. Higher HT values mean that more connections co-exist at the same time consuming more resources. This complicates the RSMA computation to find a feasible path for both serving a new incoming connection request or restoring a signal-degraded connection. This fact is reflected on the amount of used SBVT' subtransponders which is significantly increased as HT grows. It is also worth mentioning that at low *IAT_r* values (close to the *IAT_p*), e.g., 50s or 100s, the restoration RSMA algorithm for both approaches achieves better BBR compared to higher *IAT_r* values. This observation is also seen on the average used of subtransponders, where at low *IAT_r* less subtransponders tend to be occupied. For low *IAT_r*, the set of potential active connections to be degraded is smaller than at high *IAT_r*. Therefore, it is more likely that optical resources (e.g., SBVT' subtransponders) are less occupied. This does favour succeeding when computing either new or restoring optical connections. When comparing *Approaches A* and *B*, for the lowest HT (i.e., 450s), in all the performance metrics, both approaches perform similarly. However, as HT is increased, *Approach B* (i.e., directly computing link-disjoint restoration path) enhances *Approach A* in both BBR and restorability consuming less amount of SBVT's subtransponders. This behaviour is more noticeable at the highest HT value (i.e., 900s). The rationale behind this is that, in *Approach A,* the RSMA restoration mechanism first tries to restore signal-degraded connections through the same path, but selecting a lower MF. This may end up with a feasible restoration path where using a lower MF leads to allocate more SBVT' subtransponders compared to its associated active path. Consequently, *Appraoch A* tends to, in general, require more optical resources. This trend at high HT becomes more troublesome worsening both the BBR and restorability. In brief, *Approach B* attains better use of the network resources selecting the highest (and feasible) MF for the restored path. This, instead, is done at the expenses of increasing the re-configuration procedures and restoration time.

5. CONCLUSIONS

This work has presented the integration of a T-SDN controller and OAM Handler for the automatic restoration of signal-degraded connections within a flexi-grid optical transport infrastructure. To this end, a distributed monitoring system, using a high-resolution commercial BOSA, provides in-band OSNR measurements at every ongoing link traversed by an active optical connection. Measurements are retrieved and processed by the OAM Handler, which decides whether an active connection is underperforming with respect to the expected OSNR level. If this occurs, the restoration is triggered. Two restoration routing approaches are experimentally benchmarked with respect to the BBR, the average number of subtransponders and the restorability for different traffic loads and signal-degradation events. In light of the results, it is observed that directly restoring (*Approach B*) through a link-disjointed path attains a better use of the network resources (i.e., optical spectrum and SBVT' subtransponders) than trying to first restore through the same path but using a lower MF (*Approach A*). Consequently, *Approach B* improves both the attained BBR and the restorability compared to *Approach A* at the expenses of likely increasing the re-configuration procedures and restoration time.

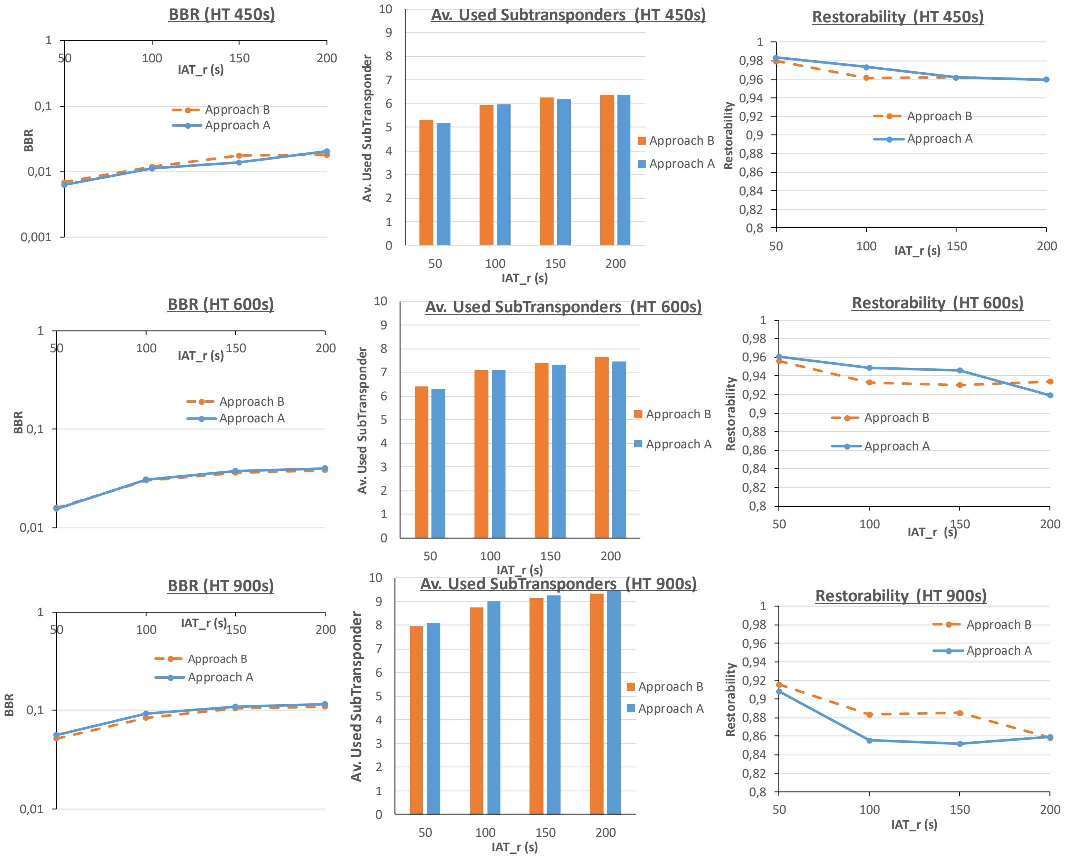

Figure 3. Performance results for the BBR, Av. Used Subtransponders and Restorability.

ACKNOWLEDGEMENTS

The research leading to these results has received funding from both Spanish MINECO DESTELLO (TEC2015-69256-R and the European Commission for the H2020-ICT-2016-2 METRO-HAUL project (G.A. 761727).

REFERENCES

[1] P. Pavon-Marino, et. al., "Evolution of Core Traffic for Growing CDNs: Is the growth rate of core network traffic overestimated?" in proc. of OFC 2017.

[2] N. Sambo, et. al., "Monitoring Plane Architecture and OAM Handler", IEEE/OSA J. of Lightwave Tech., vol. 34, April 2016.

[3] R. Martínez, et. al, "Control Plane Solutions for Sliceable Bandwidth Transceiver Configuration in Flexi-Grid Optical Networks", in IEEE Commun. Mag, vol. 54, Aug. 2016.

[4] A. Farell and Q. Zhao Ed., "An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control", IETF RFC 8283, Dec. 2017.

[5] R. Martínez, et. al., "Distributed vs. Centralized PCE-based Transport SDN Controller for Flexi-Grid Optical Networks", in Proc. of OFC 2017.

[6] D. King and A. Farell, "A PCE-Based Architecture for Application-Based Network Operations", IETF RFC 7491, March 2015.

[7] J. M. Fabrega, et. al., "Experimental Study of Adaptive Loading in IM/DD OFDM Using In-Band Optical Sub-Carrier SNR Monitoring", in proc. of OFC 2016.

[8] R. Martínez, et. al., "Experimental Validation of Transport SDN Restoration of Signal-Degraded Connections in Flexi-Grid Networks", in proc. of OFC 2018.