

# Experimental Evaluation of Path Restoration for a Centralised Impairment-Aware GMPLS-Controlled All-Optical Network

Fernando Agraz<sup>(1)</sup>, Jordi Perelló<sup>(1)</sup>, Marianna Angelou<sup>(1,2)</sup>, Siamak Azodolmolky<sup>(1,2)</sup>, Luis Velasco<sup>(1)</sup>, Salvatore Spadaro<sup>(1)</sup>, Panagiotis Kokkinos<sup>(3)</sup>, Emmanouel Varvarigos<sup>(3)</sup> and Ioannis Tomkos<sup>(2)</sup>

<sup>(1)</sup> Universitat Politècnica de Catalunya (UPC), Barcelona (Spain), ✉ [agraz@tsc.upc.edu](mailto:agraz@tsc.upc.edu)

<sup>(2)</sup> Athens Information Technology (AIT), Athens (Greece)

<sup>(3)</sup> Research Academic Computer Technology Institute (RACTI), Patras (Greece)

**Abstract** This paper evaluates a centralised impairment-aware path restoration approach for GMPLS-controlled transparent optical networks. Experimental results on a 14-node network test-bed show successful QoT compliant path restoration of around 3.6 seconds.

## Introduction

Recent research activities are concentrated on solutions that will enable core networks to better cope with the ever-growing traffic demand. Next generation optical core networks may rely on the state of the Quality of Transmission (QoT) as a way to ensure unhindered operation in view of the increased forecasted traffic. In order to realise the vision of dynamic impairment-aware networking, an integrated scheme spanning from the optical transport plane up to the management plane is considered essential<sup>1</sup>. Therefore, an impairment-aware control plane is incorporated into the network, featuring dynamic connection provisioning and restoration capabilities.

The dynamic Network Planning and Operation Tool (NPOT) developed in the DICONET project<sup>1</sup>, integrated in a unified extended Generalized Multi-Protocol Label Switching (GMPLS) control plane, implements such a scheme. In a previous work<sup>2</sup>, the performance of this integrated platform was demonstrated over an experimental test-bed under dynamic traffic and using real-time QoT assessment.

This paper focuses on the capability of the centralised scheme to detect failures and restore the affected connections while ensuring acceptable QoT. To this end, a 14-node experimental setup is utilised to evaluate the failure handling capabilities and the inter-working of the different modules in the proposed centralised architecture for lightpath restoration. From the obtained experimental results, the average lightpath restoration proves to be quite fast, although the centralised and therefore sequential nature of the backup path computation induces some extra delay. The computed restored paths are longer but still satisfy the QoT constraint.

## Centralised Implementation Approach

In what follows, we describe the deployed centralised integrated scheme and the failure han-

dling process. As depicted in Fig. 1, the NPOT resides on top of the proposed scheme, which in the centralised operation mode is an engine common to all network entities that performs the lightpath computation. The control plane is realised via a set of Optical Connection Controllers (OCCs) managing one optical device each. In particular, OCCs implement the full GMPLS protocol set (i.e., RSVP-TE<sup>3</sup>, OSPF-TE<sup>4</sup> and LMP<sup>5</sup> protocols) in order to provide automated lightpath establishment, maintenance and release over the all-optical network. For the communication between OCCs and NPOT an ad-hoc TCP-based messaging protocol has been developed.




Fig. 1: The centralised integrated scheme.

When a failure occurs, the optical transparency leads to a propagation of Lost of Light (LoL) alarms, detecting LoL in the incoming ports of all downstream nodes from the failure point. The LMP protocol is responsible for providing the GMPLS control plane with failure localisation capabilities<sup>5</sup>. For its operation, LMP requires an active control channel between the failed link endpoints. This restricts its use to only *out-of-fiber* control plane configurations, where control channels remain alive in spite of data plane failures.

As soon as LoL alarms for one or multiple input wavelengths are received, the GMPLS control node has to determine whether the failure is localised in the local link connected to the adjacent node, that is, the upstream node of the affected lightpaths, or in any further upstream link. To this end, the node sends a *Channel-*

*Status* message<sup>5</sup> to the adjacent node containing the list of individual failed wavelengths (if no individual wavelength is specified, this indicates that the whole link is failed). Upon receiving this message, the adjacent node correlates the failure checking whether it is also detected locally for the affected lightpaths. If the failure is clear on its input wavelengths, the failure is localised on the link connecting both nodes. Otherwise, the failure is located on a further upstream link. Once the failure is correlated, the upstream node sends a *ChannelStatus* message back to the downstream node, indicating whether the link is failed or not.

When an upstream node localises the failure, it can start the restoration of the affected lightpaths. In our centralised implementation approach, a database update message is firstly sent to the centralised NPOT informing it about the current status of the failed link. In this way, the NPOT can update its global Physical Parameters and Traffic Engineering Databases (gPPD and gTED), which describe the current network status both at the topology and physical layer levels. Next, the upstream node sends an RSVP-TE *Notify* message to the source node of each affected lightpath.

Upon receiving the RSVP-TE *Notify* message for a specific failed lightpath, the source OCC requests a backup lightpath from the NPOT for restoration purposes. In this case, the NPOT forwards the lightpath request to its online Impairment Aware Routing and Wavelength Assignment (IA-RWA) module<sup>6</sup>, which uses the QoT estimator of the NPOT in order to assess the feasibility of a given lightpath (reflected in the Q-factor). The QoT estimator that lies in the core of NPOT consults the two global repositories (gPPD and gTED), already updated, to retrieve information about the current network status. Therefore, the NPOT estimates the new backup lightpaths avoiding those failed network components.

As soon as the NPOT finds a backup lightpath to serve the failed connection, it communicates it back to the source OCC node. This OCC is then responsible for initiating the lightpath establishment process using standard RSVP-TE. Being this one successful, the gPPD and gTED databases are updated using an extended version of the OSPF-TE protocol that enables the advertisement of wavelength and physical parameters information. Besides, the OCC informs the Network Management System (NMS) either about the successful backup path establishment or its failure due to lack of resources or poor QoT.

## Experimental results

The performance of the proposed centralised lightpath restoration approach has been validated


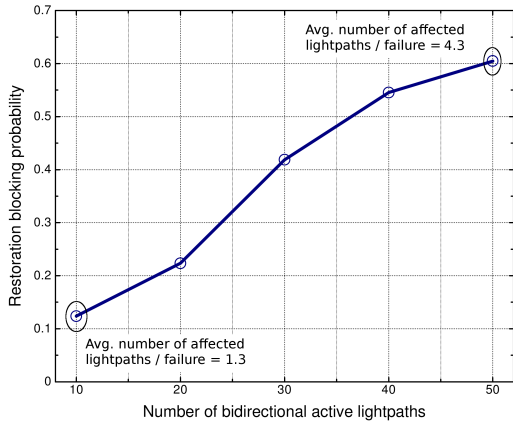


Fig. 2: The DICONET experimental test-bed.

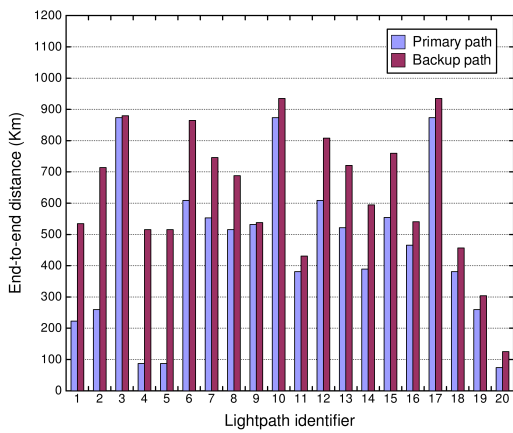
on the DICONET test-bed, located at the UPC premises (Fig. 2). The test-bed describes the same topology as the 14-node Deutsche Telekom (DT) network<sup>2</sup>, where 10 bidirectional wavelengths per link have been assumed. Each network node is composed of an OCC and a WSS-based OXC emulator interconnected through the CCI interface. In turn, all OCCs are interconnected to the NMS and the NPOT through the NMI-A and OCC-NPOT interfaces, respectively. The connectivity between OCCs is supported over 100 Mbps point-to-point Ethernet links, which describe an *out-of-fiber* control plane with the same topology as the emulated all-optical data plane. OCCs implement the whole GMPLS protocol set.

For the evaluation, we initially load the network with a certain number of bidirectional connections (10, 20, 30, 40, 50) between randomly selected node pairs. These connections can be either 1+1 protected, establishing working and backup paths for them, or restorable. We assumed a 30%-70% protected-restorable ratio. To achieve this goal, valid QoT compliant routes are requested to the centralised NPOT. Then, on each deployed network scenario, 10 independent failures are caused in randomly selected links (only the links carrying restorable traffic are considered), which makes restoration actions for each affected restorable lightpath to be triggered.

Fig. 3 shows the Restoration Blocking Probability (RBP) on each deployed network scenario. These values are defined as the total number of rejected restorations over the total number of restorable lightpaths affected by the 10 independent failures. As expected, the RBP increases with the number of active lightpaths in the network. This can be explained both by the greater



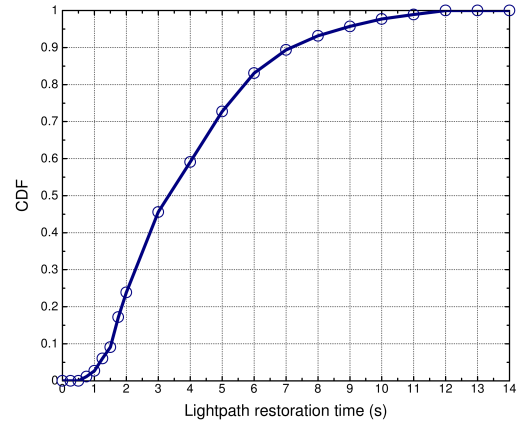
**Fig. 3:** Restoration blocking probability increases with the number of active lightpaths.



**Fig. 4:** End-to-end distance of primary and backup paths: backup paths still satisfy the QoT constraint.

number of active and affected lightpaths per failure, which leads to lack of resources, and the longer backup lightpaths that are rejected due to the QoT constraint. As shown, for 10 and 50 active lightpaths, the average number of affected ones per failure is 1.3 and 4.3, and the resulting RBP is around 0.1 and 0.6, respectively. In particular, Fig. 4 details the physical end-to-end distances of all the restored lightpaths in the scenario with 30 active ones. Note that the backup paths are always longer than the primary ones. In fact, the average physical distance of the primary lightpaths is 452 Km, whereas for the backup ones it increases to 630 Km. Nonetheless, the NPOT always assures the required QoT for them, which is crucial for the successful restoration.

The same set of experiments has been also used to measure the lightpath restoration time in the network, obtaining 3.6 s in average. Fig. 5 also plots the Cumulative Distribution Function (CDF) of the measured lightpath restoration times, that is, the probability that a lightpath restoration does not exceed a certain number of seconds. As seen, 72% of the lightpath restora-



**Fig. 5:** CDF function of the measured lightpath restoration times. 72% of the restorations take less than 5 s.

tions are performed within promising 5 s, half of the connection setup times reported in<sup>7</sup>. Note that the sequential processing behaviour of NPOT may occasionally lead to increased restoration times (e.g., 10 s in 2% of the cases), especially when a high number of lightpaths are affected. In fact, all lightpath restorations are equally treated in this work. Although not considered here, different restoration priority classes could be defined in the centralised NPOT scheduler. This would allow to serve the lightpaths with the highest priority first, thus ensuring low restoration times for them.

## Conclusions

This paper reported the experimental evaluation of a centralised impairment-aware lightpath restoration involving the NPOT and the extended GMPLS protocol stack developed within the EU DICONET Project. Future work within DICONET will be devoted to further reduce lightpath restoration times to milliseconds' time-scales by FPGA hardware acceleration of the QoT estimator.

## Acknowledgments

This work has been supported by the European Commission through the FP7 DICONET Project and the Spanish Science Ministry through the project ENGINE (TEC2008-02634).

## References

- 1 S. Azodolmolky et al., IEEE Comm. Mag. 47, 38 (2009).
- 2 F. Agraz et. al., Proc. OFC'10, PDPD5 (2010).
- 3 L. Berger (ed.), IETF RFC 3473 (2003).
- 4 D. Katz et. al., IETF RFC 3630 (2003).
- 5 J. Lang (ed.), IETF RFC 4204 (2005).
- 6 P. Kokkinos et al., Proc. GLOBECOM'09, pp.1-7, (2009).
- 7 T. Tsuritani et. al., Proc. OFC'08, NWF5 (2008).