Technical University of Denmark

DTU

Increasing Restorability for Local-to-Egress Restoration in GMPLS Controlled Networks with Limited Wavelength Conversion

Ruepp, Sarah Renée; Buron, Jakob Due; Andriolli, N.

DTU Library
Technical Information Center of Denmark

Increasing Restorability for Local-to-Egress Restoration in GMPLS Controlled Networks with Limited Wavelength Conversion

Sarah Ruepp, Jakob Buron
COM•DTU
Technical University of Denmark
Kgs. Lyngby, Denmark
{sr, jbu}@com.dtu.dk

Nicola Andriolli
CEIRC
Scuola Superiore Sant'Anna
Pisa, Italy
nick@sssup.it

Abstract— **We study the performance of the local-to-egress restoration method in GMPLS controlled optical networks, when a limited number of wavelength converters are available. We evaluate the recovery percentage for a converter-saving label assignment scheme and compare its performance to a simple scheme with and without releasing resources that are not used after the failure (stub release). Our simulations show that the converter-saving label assignment scheme without stub release outperforms the simple label assignment scheme with stub release, to the extent that stub release can be avoided without sacrificing restorability.**

Keywords – local-to-egress restoration; stub release; GMPLS; wavelength conversion; suggested vector.

I. INTRODUCTION

Resiliency to single span failures in mesh networks has traditionally been provided by using either path or span restoration [1]. In path restoration, the affected connection is restored between its end nodes, where the long notification time is counterweighted by high resource efficiency and the possibility of identifying a path of equal length compared to the original path. In span restoration, the traffic is restored between the failure-adjacent nodes by bypassing the failed span, resulting in a short notification time but also causing longer (i.e. more resource consuming) end-to-end paths compared to path restoration. Another method is local-to-egress restoration [2][3], where affected connections are restored between the upstream failure-adjacent node and the destination node. It combines the advantages of span and path restoration, resulting in short notification time, possibility of short total path length and high resource efficiency (between that of path and span restoration [2]), which makes it a strong candidate for solving restoration tasks. An overview of the three restoration methods is illustrated in Fig. 1; the focus of this work is local-to-egress restoration.

When setting up connections in Wavelength Division Multiplexed (WDM) networks, each connection requires a route and a wavelength, which are usually obtained in a
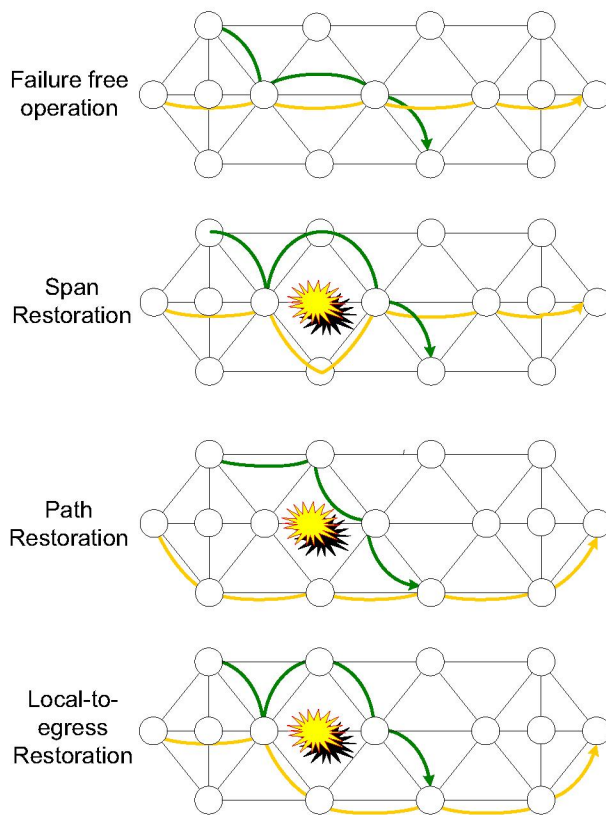
Figure 1. Overview of mesh restoration methods.
Local-to-egress restoration is analysed in this study.

decoupled process. Generalized Multi-Protocol Label Switching (GMPLS) [4] has evolved from the MPLS control plane and has shown to be well-suited for controlling next generation optical networks, such as WDM networks. GMPLS is a protocol suite consisting of the Open Shortest Path First Protocol (OSPF), the Resource Reservation Protocol (RSVP), both with Traffic Engineering Extensions (TE), and the Link Management Protocol (LMP). The aforementioned routing and wavelength (i.e. label in the GMPLS paradigm) assignment

process is split between the OPSF-TE protocol, which is responsible for collecting topology information used when finding a route for each connection; and the RSVP-TE protocol, which takes care of finding an appropriate label. A label has to be continuously available between the source and destination if conversion between wavelengths is unavailable. Current emerging technologies allow for conversion between wavelengths, so the end-to-end label continuity constraint is removed. Nevertheless, these wavelength converters (WCs) are expensive, which means that typically only a limited number of WCs are available at each node. If all WCs are occupied, a label-continuous path has to be available, which greatly reduces the probability of successful connection setup or restoration.

In this study, we investigate how the restorability in case of single span failures can be increased by using a converter-saving label assignment scheme in local-to-egress restoration. We compare the restoration percentages of this scheme to a simple label assignment scheme. In particular, we examine if the new scheme still compares favorably when the simple scheme uses the performance-enhancing stub release technique. The remainder of the paper is organized as follows: section 2 explains converter-saving label assignment and section 3 illustrates local-to-egress restoration path setup. Section 4 presents the simulation study and the results are shown in section 5. Section 6 concludes the paper.

II. CONVERTER-SAVING LABEL ASSIGNMENT

The simplest way to assign a label for a Label Switched Path (LSP) is to use a No Preference (NP) scheme. When a setup request reaches the destination node, a label is chosen first-fit amongst the free wavelengths on the previous hop. A reservation request for this label is then propagated upstream. The chosen label will be used as long as possible, but when a node cannot use it any longer because the specific label is already occupied on its previous hop, a WC is required. If a WC is available, a new label is chosen amongst the free upstream wavelengths and this label is then propagated upstream for as long as it can be used.

In order to ensure that the use of WCs is minimized during LSP setup, we use a novel signaling extension to the RSVP-TE protocol called Suggested Vector (SV) first proposed in [5]. It is an optional object used together with the label set, which is a standard protocol extension [6]. The label set allows an upstream node to control the label selection of a downstream node. The SV collects information on the number of WCs a specific label requires when provisioning a connection. This allows the destination node to select the label that requires fewest WCs. If several labels have equal SV-values, ties are broken first-fit. The label set and the SV are propagated downstream along an envisaged route in the RSVP-TE Path message. Whenever a WC is required to use a specific label, the corresponding SV element is increased by one. Once the downstream node is reached, the label with the lowest SV value is chosen and then propagated upstream in the Resv message. The principle is illustrated in Fig. 2, where the SV scheme selects label 3, requiring 0 WC. The NP scheme, using first-fit amongst free labels chooses label 1, requiring 1 WC.
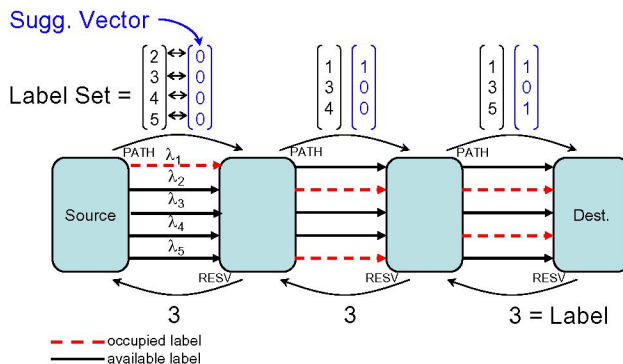

Figure 2. Suggested Vector (SV) operation for connection setup

In the event of a span failure, the upstream failure-adjacent node at the same time initiates the request for a restoration path between itself and each affected connection's destination. Once the route is found, the upstream failure-adjacent node sends a Path message towards the destination node. No resources are reserved in the Path message and due to the tight timeframe, several connections' restoration requests often compete for the same resources. These contentions can be solved by wavelength conversion, but especially when WCs are limited it is necessary to minimize their use. The SV scheme has shown to minimize WC usage in the provisioning phase [5]. The SV scheme for local-to-egress restoration is shown in Fig. 3. From the upstream failure-adjacent node, the label set and the SV are propagated towards the connection's destination. Along the restoration path, the SV is updated when WCs are required and the destination node can then choose the label requiring fewest WCs along the route in its Resv messages.
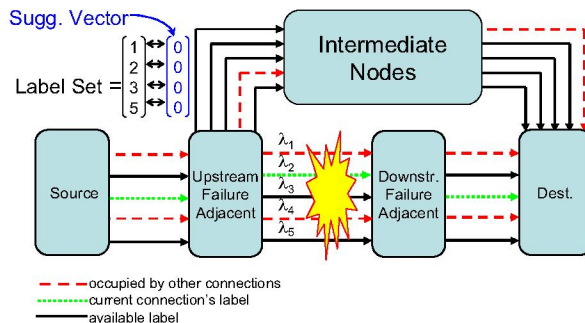

Figure 3. Suggested Vector (SV) used in local-to-egress restoration

III. RESTORATION PATH SETUP

Before actually restoring affected connections, stub release can be performed, meaning that the span and node resources that a connection used before the failure (i.e. the stubs) can be released. Stub release is performed by letting the downstream failure-adjacent node send teardown messages towards each affected connection's destination node. Upon reception of such a message each node frees the stub's labels and WCs, if used. These resources are now available for establishing restoration paths. The principle of stub release is illustrated in Fig. 4. Keeping the stubs occupied simplifies the control, but also

means that resources can be unnecessarily occupied. Releasing the stub resources requires additional signaling actions and complicates reversion after the span failure is repaired, but having more resources for recovery available increases the possibility of successful connection restoration. Especially the complicated reversion process is considered a drawback. Since during failure-free operation connections in a network are routed in the best possible way, they are diverted to a sub-optimal path in case of a failure. Hence, the traffic should be redirected to its pre-failure path as fast and uncomplicated as possible after the failure has been repaired, but if stub release was preformed, the reversion process becomes more complex.

Due to the freeing of resources, we expect a better restoration success rate when stub release is used, especially when available labels and WCs are limited. But the question remains: Can a more sophisticated label assignment scheme (i.e. the SV) achieve the same performance without using stub release as a simple label assignment scheme that uses stub release in the restoration phase? If yes, then the drawbacks of stub release can be avoided without sacrificing restorability.
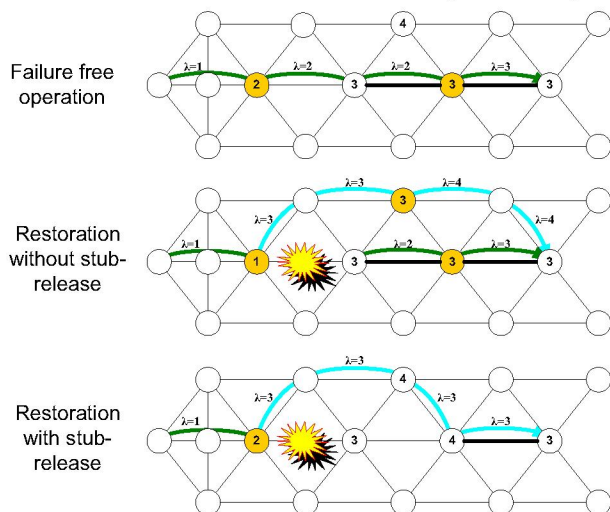

Figure 4. Restoration of the connection shown at the top with and without stub release. Bold spans are fully occupied. Dark green is original path. Light blue is restoration path. The numbers in some nodes state the number of free WC's. A WC is required at the nodes with shaded background.

To minimize the delay and the resource usage, the failed connections should be restored over the shortest possible path. This is obtained by using the Dijkstra shortest path algorithm to find a least-hop path between the upstream failure-adjacent node and the destination node. A span is included in the route computation, as long as it has at least one label available. Due to the separation of the routing and wavelength assignment process, a situation may arise where a route over a set of spans with available capacity can be found, but the desired label is already occupied and the required WC is unavailable. In this case, the connection setup request experiences blocking, receiving proper RSVP-TE error messages. Blocking will also occur if a span had available resources during Path message flow, but has been fully occupied before the actual resource reservation (i.e. Resv message flow) was completed. When the upstream failure-adjacent node receives such an error message

it calculates a new restoration route, excluding the span that caused the blocking. By using the experimental GMPLS crank-back feature [7], the upstream failure-adjacent node can be notified of which span or node is causing the blocking. After a backoff period, restoration is re-attempted on the shortest path excluding spans that previously caused blocking in the same connection's restoration request (the k-shortest path). The restoration procedure is illustrated in Fig. 5, which is repeated until the connection is either successfully restored or the Dijkstra algorithm cannot find a route, which means that the connection is non-restorable.
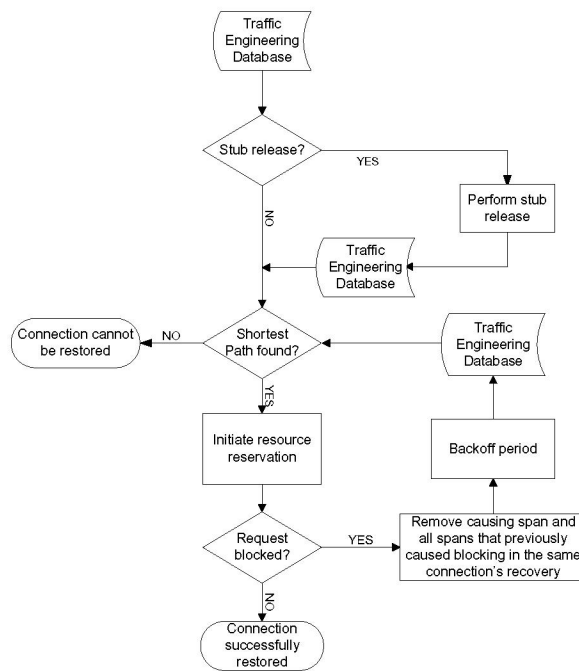

Figure 5. Connection restoration process

IV. SIMULATION STUDY

In this study, we evaluate the restoration efficiency of the SV scheme compared to the NP scheme with and without stub release applied to local-to-egress restoration. The simulations are carried out in OPNET Modeler [8] and evaluated in the NSFNET [9] and the Pan-European triangular topology network [10]. The network topologies are illustrated in Fig. 6 and 7 respectively.
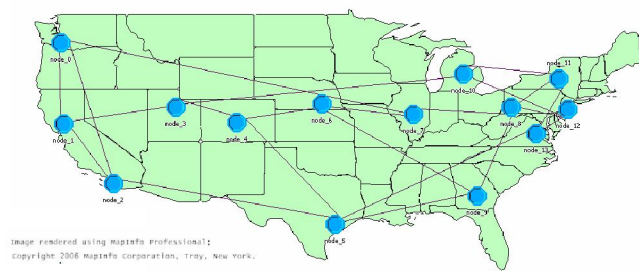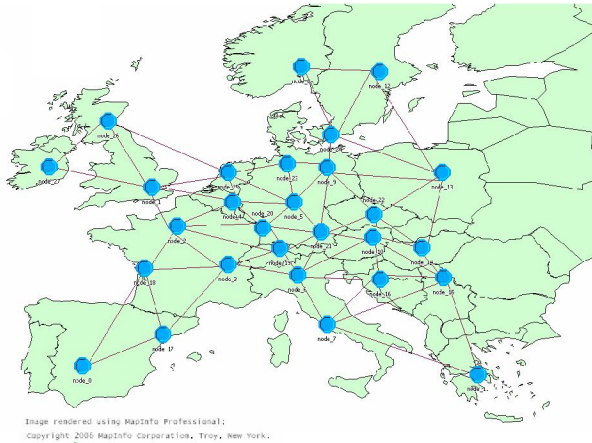

Figure 6. NSFNET network topology

Figure 7.   Pan-European triangular topology network

The number of nodes, spans and the nodal degree for both network topologies are illustrated in Table I.

TABLE I.   EVALUATED NETWORK TOPOLOGIES

Topology	Nodes	Spans	Nodal Degree
NSFNET	14	22	3.14
Pan-European	28	61	4.36

Each span has a capacity of 10 wavelengths (labels). The network is incrementally populated with unidirectional connections up to a given average span load (0.4 to 0.6). Each connection occupies one wavelength and is assigned a label. The connections are uniformly distributed over all source-destination pairs in the network. When the desired load is reached, the spans are failed consecutively, meaning that the network is reversed to its pre-failure state before the next span failure is simulated. We evaluate three different scenarios:

NP: Connections are setup and restored with the NP label assignment scheme. Stub labels and WCs are <u>not</u> released.

NP Stub Release: Connections are setup and restored with the NP scheme. When a failure occurs and before initiating restoration requests, stub release is performed (i.e. freeing labels and WC that are no longer used).

SV: Connections are setup and restored with the SV label assignment scheme. Stub labels and WCs are <u>not</u> released.

V.   RESULTS

The results in Fig. 8 and Fig. 9 are calculated with a confidence interval of 95%, averaged over all single span failures, averaged over 10 repetitions of connection setups with different random seeds. The graphs for the NSFNET (Fig. 8) and Pan-European network (Fig. 9) respectively, as expected show that the NP scheme allows a higher restoration percentage when stub release is enabled, compared to when no stub resources are freed. The NP scheme can assign labels in an
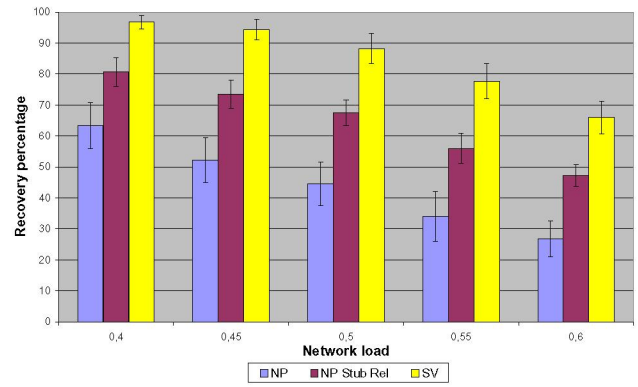

Figure 8.   NSFNET recovery percentages for varying loads. 5 WCs per node.
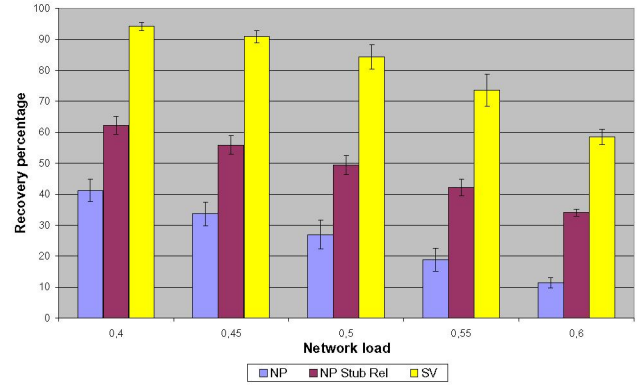

Figure 9.   Pan-European network recovery percentages for varying loads. 5 WCs per node.

inefficient manner, resulting in a high WC usage per connection. This leaves few WCs for the restoration process, meaning that the possibility for successful connection restoration is diminished especially when WCs are sparse. If the stubs are released, these resources can be used during the restoration process and more connections can be restored. More importantly, we see that the SV scheme without using stub release gives a better restoration success rate than both NP and NP with stub release. This is due to the fact that the SV scheme assigns labels in a converter-saving manner, providing more WCs for the restoration process, hence allowing more connections to be restored. The results show that a more efficient label assignment scheme (performed through proper control plane enhancements, i.e. adding the optional SV object) can significantly increase restoration percentages, even without resorting to stub release actions.

When comparing Fig. 8 and Fig. 9, we notice that the SV scheme achieves different improvements in recovery percentage depending on the network topology. For the NSFNET, the recovery percentage for SV compared to NP with stub release is relatively increased by approximately 30 percent (calculated as: (mean recovery percentage (SV) - mean recovery percentage (NPwithStubRel)) / mean recovery percentage (NPwithStubRel)), averaged over all loads). For the Pan-European topology, the increase is in the order of 65 percent.
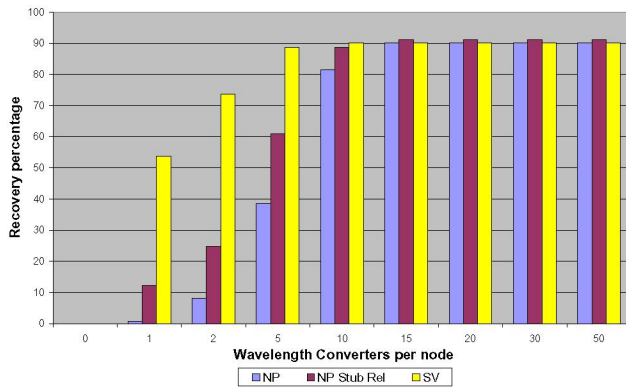
Figure 10. NSFNET recovery percentages for varying WCs per node. The network load is fixed at 0.5.
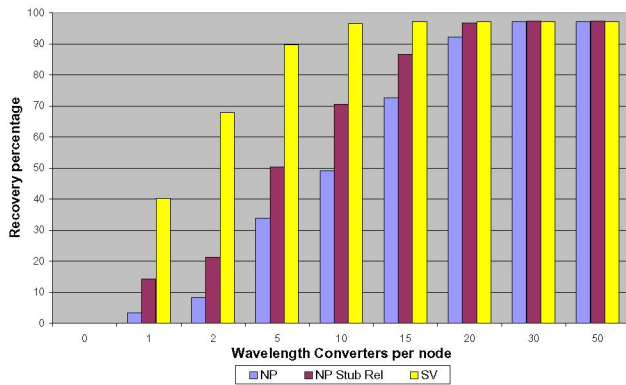

Figure 11. Pan-European triangular topology network recovery percentages for varying WCs per node. The network load is fixed at 0.5.

Fig. 10 and Fig. 11 illustrate the recovery percentages of the different schemes when varying numbers of WCs are available in the nodes. When no WCs are available, none of the affected connections can be restored. It can be seen that the SV scheme is especially useful in a situation where WCs are limited; the advantage diminishes if many WCs are available. This is due to the fact that the WCs no longer are the limiting factor; the success of a restoration now depends on whether there are labels available on the desired locations in the network. This also explains the slightly higher restoration percentage of the NP with stub release scheme when many WCs are available. Comparing Fig. 10 and Fig. 11, we notice that the Pan-European topology requires more WC per node to reach its highest restoration percentage. This is due to the fact that it has a dense core, which many connections need to traverse.

VI. CONCLUSION

Based on the simulation results, we observe that the increase in restoration percentage for SV compared to NP with stub release depends on the number of WCs per node and the network topology. Furthermore, the difference is only weakly dependent on the load.

We have shown that the SV scheme is particularly useful for label assignment in local-to-egress restoration when WCs are sparse. Simulations show that using the optional SV object surpasses the NP scheme even with stub release in terms of restorability. Most importantly, our results clearly show that the SV label assignment scheme is ideal for provisioning and local-to-egress restoring connections when the drawbacks of stub release must be avoided without sacrificing restorability.

REFERENCES

[1] Vasseur, J.-P., Pickavet, M., Demeester, P. "Network Recovery – Protection and Restoration in Optical, SONET-SDH, IP, and MPLS", Elsevier – Morgan Kaufmann Publishers, 2004, ISBN: 0-12-715051-x

[2] Gruber, C.,G. "A Comparison of Bandwidth Requirements of Path Protection Mechanisms", ICN 2005, LNCS 3420, pp. 133-143, Springer-Verlag Berlin Heidelberg 2005.

[3] Wang, J., Sahasrabuddhe, L., Mukherjee, B. "Path vs. Subpath vs. Link Restoration for Fault Management in IP-over-WDM Networks: Performance Comparisons Using GMPLS Control Signaling", IEEE Communications Magazine, 2002

[4] Mannie, E. (ed.), "Generalized Multi-Protocol Label Switching (GMPLS) architecture", RFC 3945, October 2004.

[5] Andriolli, N., Buron, J., Ruepp, S., Cugini, F., Valcarenghi, L. and Castoldi, P., "Signalling protocol extensions for converter-saving wavelength assignment in GMPLS optical networks". HPSR 2006, Poznan, Poland, 2006.

[6] Berger, L. (ed.), "Generalized Multi-Protocol Label Switching (GMPLS), Signaling Functional Description", RFC 3471, January 2003.

[7] Farrel, A. (ed.), "Crankback Signaling Extensions for MPLS and GMPLS RSVP-TE". <draft-ietf-ccamp-crankback-05.txt>, IETF Draft, work in progress

[8] OPNET Technologies, Inc., http://www.opnet.com.

[9] Miyao, M. and Saito, H., "Optimal Design and Evaluation of Survivable WDM Transport Networks", IEEE. J. Sel. Areas in Comm., vol. 16, no. 7, pp. 1190-99, September 1998.

[10] Inkret, R., Kuchar, A. and Mikac, B. (eds.), "Extended Final Report of COST Action 266 – Advanced Infrastructure for Photonic Networks", 2003