


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Restoration in all-optical GMPLS networks with limited wavelength conversion

Sarah Ruepp^{a,*}, Nicola Andriolli^b, Jakob Buron^a, Lars Dittmann^a, Lars Ellegaard^c

^a Networks Competence Area, DTU Fotonik, Technical University of Denmark, Ørstedss Plads, 2800 Kgs. Lyngby, Denmark

^b Scuola Superiore Sant'Anna, Pisa, Italy

^c Vitesse Semiconductor Corporation A/S, Herlev, Denmark

ARTICLE INFO

Article history:

Received 1 June 2007

Received in revised form 13 December 2007

Accepted 6 February 2008

Available online 25 March 2008

Keywords:

Restoration

Limited wavelength conversion

GMPLS

Simulation

ABSTRACT

A new generation of optical components and the advance of the Generalized Multi-Protocol Label Switching (GMPLS) control plane supporting dynamic provisioning and restoration of optical connections (i.e., lightpaths), brings the vision of the dynamic all-optical network closer to reality. An emerging technology is the conversion between wavelengths, which removes the wavelength continuity constraint, thus allowing an easier and more flexible connection allocation. A limitation in the number of wavelength converters impairs their benefits especially during the restoration phase, when many simultaneous recovery attempts must share residual resources.

This paper investigates the restoration performance of GMPLS-controlled all-optical networks with limited wavelength converter deployment. We investigate how different restoration methods, namely span restoration, segment restoration, and end-to-end restoration are affected by the availability of a limited number of wavelength converters at each node. For this purpose an enhanced wavelength assignment scheme compliant with GMPLS signaling is exploited, aiming at saving converters by assigning a higher preference to wavelengths not requiring conversion. An extensive simulation study has been conducted comparing the performance of this scheme to the most advanced scheme based on standard GMPLS signaling for the three restoration methods.

Simulation results show that the enhanced wavelength assignment scheme significantly reduces the number of wavelength converters (WCs) necessary to achieve good recovery performance. The enhanced scheme especially improves span restoration performance, where the matching between the stubs' and recovery segment wavelength may require a WC. End-to-end restoration is the least affected, due to a higher degree of freedom in the route choice, while segment restoration performance lies in between.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

The ever growing demand for bandwidth in fiber-based backbone networks has been addressed in the past years by increasing the span capacity using Wavelength Division Multiplexing (WDM). Aiming at scalability, flexibility and cost efficiency, commercially deployed network elements

have included static optical add-drop multiplexers (OADM) and signal transparent “optical” cross-connects with electrical regeneration. Only recently have reconfigurable optical add-drop multiplexers (ROADM) and optical cross-connects (OXC) been accepted commercially, having true optical cores allowing rearrangeability between ports [1], and possibly change of wavelength. The components represent a new generation of commercially available optical network elements seeking to gradually realize the all-optical network vision hitherto mostly considered in academia. Despite these recent progressions, wavelength

* Corresponding author. Tel.: +45 4525 3627; fax: +45 4593 6581.
E-mail address: sr@com.dtu.dk (S. Ruepp).

conversion is still a critical and expensive node capability, and it is within this context the present work should be seen.

In experimental wavelength routed optical networks (WRONs), lightpaths, i.e., all-optical end-to-end switched connections between node pairs, are typically established at wavelength granularity [2]. WRONs can be affected by failures. Since these networks carry large amounts of data, survivability is of paramount importance to cause as little disturbance to the traffic as possible. Survivability techniques are categorized into network protection and network restoration. Network protection refers to the situation of pre-calculated backup paths, which are provisioned simultaneously to the primary path; whereas connections are dynamically recovered after the failure occurs within a spare capacity pool in network restoration [3,4]. Both methods have their advantages and disadvantages. Network protection can provide guaranteed recovery for specified failure scenarios in a short timeframe (all the backup paths are pre-provisioned), but suffers from capacity inefficiency and low flexibility. Network restoration on the other hand must first provision the backup paths when the failure occurs, but the drawbacks of this approach are counterweighted by the increased flexibility towards the faced failure scenarios. The adaptation to a dynamic network environment makes network restoration a very suitable candidate to provide resilience towards failures. To overcome a failure in a network, different restoration methods are distinguished, mainly characterized by the end nodes of their recovery routes. In the literature, three main restoration methods are described, namely span, segment and end-to-end restoration [3–7]. All these methods have their advantages and challenges. In *span* restoration [8–10], the failed connection is recovered between the failure adjacent nodes, which means that the upstream failure adjacent node initiates the restoration, and the restoration path ends at the downstream failure adjacent node. Span restoration has the benefit of a short notification time (the node that detects the failure initiates the recovery), but also has a somewhat limited flexibility, since the downstream failure adjacent node must be visited. Furthermore, the wavelengths of the original path's stubs and the backup path must be matched at the failure adjacent nodes. In *end-to-end* restoration, the failed connection is restored between its source and destination nodes. While this means an increased notification time, a higher degree of freedom is achieved, where the connections' original path does not have to be resumed and the shortest possible path can be found on a network wide basis. A hybrid between span and end-to-end restoration is *segment* restoration. In segment restoration, the backup path can be established to circumvent any segment of the original path. It hence combines the advantages of span and end-to-end restoration. A particular case of segment restoration is *local-to-egress* restoration, where a connection is restored between the upstream failure adjacent node and the destination node, combining short notification time and high resource efficiency [11,12]. This form of segment restoration is used in this study.

Dynamic WRONs leverage a distributed control plane to automatically perform the required network functions,

such as lightpath provisioning, fault localization and recovery. A promising control plane standard is the Generalized Multi-Protocol Label Switching (GMPLS [13]) protocol suite, comprising extensions to routing (e.g., OSPF-TE) and signaling protocols (e.g., RSVP-TE), and a link management protocol (LMP). The routing protocol typically advertises only summarized link information, which is used during lightpath provisioning and restoration to find a feasible route for the incoming request. Then a signaling session is triggered along the found path in order to perform the wavelength assignment.

Wavelength assignment efficiency, and by extension provisioning and restoration efficiency, are significantly influenced by the wavelength conversion capability [14,15]. Wavelength-converters (WCs) overcome the wavelength continuity constraint, allowing setup of a lightpath even when no wavelength is continuously available on each hop of the route. The conversion between wavelengths can be achieved in different ways [16,17], but all these implementations are very costly. Furthermore, some WCs may only allow conversion between a limited range of wavelengths [18,19], increasing the complexity of routing and wavelength assignment [20,21]. Due to the high price of WCs, several studies focus on minimizing their deployment in a network, either through advanced routing and wavelength assignment heuristics or sharing schemes [15,21–27]. As a result of the high cost and the WC minimization achieved by the design studies, typically only a limited number of WCs are deployed throughout networks.

In this study, we aim at combining three emerging issues in optical networks. We combine (i) dynamic connection provisioning and restoration with (ii) the scenario of few available WCs in each node in (iii) a GMPLS controlled all-optical network, and evaluate the effect that limited wavelength conversion has on the restoration performance of span, segment and end-to-end restoration.

The remainder of this paper is organized as follows: in Section 2 the previous work in the field is reviewed and the novelties of this study are highlighted. Section 3 presents distributed wavelength assignment schemes leveraging GMPLS signaling, while Section 4 describes the network model used in this study. In Section 5 the simulation scenario is detailed, while results are shown and commented in Section 6. Section 7 concludes the paper.

2. Previous work

Network restoration techniques have been widely studied in the literature, highlighting their performance/complexity trade-offs in terms of resource efficiency, recovery time, and success rate [5,6,28–30]. In particular in [5], path (i.e., end-to-end) restoration is compared to span and sub-path (i.e., segment) restoration in a network with full wavelength conversion capability. Furthermore, end-to-end restoration with additional recovery retrials is shown to achieve high network availability at the expense of a restoration time slightly higher than span and segment restoration, which in turn are characterized by a lower success rate.

Restoration performance can be significantly improved by exploiting WCs, as shown in [31–34]. In [35,36], an

analytical model of a pre-planned segment restoration method is developed, showing that the benefits of full wavelength-conversion dramatically improve for increasing path length, but become less significant for increasing wavelengths per fiber. An online wavelength assignment algorithm is proposed in [37], where the WC usage and the connection blocking are evaluated, exploiting a multi-cast conversion model where the output of the WC is split. However only few protection and restoration papers considered limitations to wavelength-conversion capability. In [38], a protection method is presented where limited-range wavelength-converters are considered.

Different control and management approaches can be taken to control provisioning and restoration actions in an optical network. The desire to provide a unified control framework leads to the Generalized Multi-Protocol Label Switching (GMPLS) protocol suite [13,39,40], which has undergone standardization efforts during the past few years, and is emerging as a control plane for next generation networks. Despite this development, only few studies [41,42] consider the effect of GMPLS signaling when dealing with network restoration. However, these studies do not take the effect of limited wavelength conversion availability into account.

For this study we leverage the studies in [23,27], where we investigated wavelength assignment schemes in the provisioning phase. In particular, we introduced an enhanced signaling object called Suggested Vector (SV), which enables label ranking aiming at WC usage reduction. In [8–10] we studied specific techniques aimed at reducing WC usage and resource contention in span restoration. In [12], local-to-egress restoration was investigated, showing that a WC-saving wavelength assignment could achieve better recovery performance compared to a simple wavelength assignment scheme utilizing complicated stub-release mechanisms.

The focus of this paper is on the restoration performance of all-optical GMPLS controlled networks with limited WC placement. Given a fixed number of WCs deployed in each node, we perform a comparative study on the performance of span, segment and end-to-end restoration methods, in particular how these restoration methods are affected by the limitation of WCs and whether using standard or enhanced GMPLS signaling influences the results. To our knowledge, this is the first combined study of these important topics in optical network operation.

3. Network model

In order to evaluate the influence of WCs on the performance of span, segment and end-to-end restoration, we create a model of the network which serves as a basis for the conducted simulation study. The model incorporates realistic assumptions about the node and network architectures.

3.1. Node architecture

The used logical node architecture is illustrated in Fig. 1. Each node is connected to F_i input fibers and F_o output fibers, and each fiber span carries W wavelengths in each

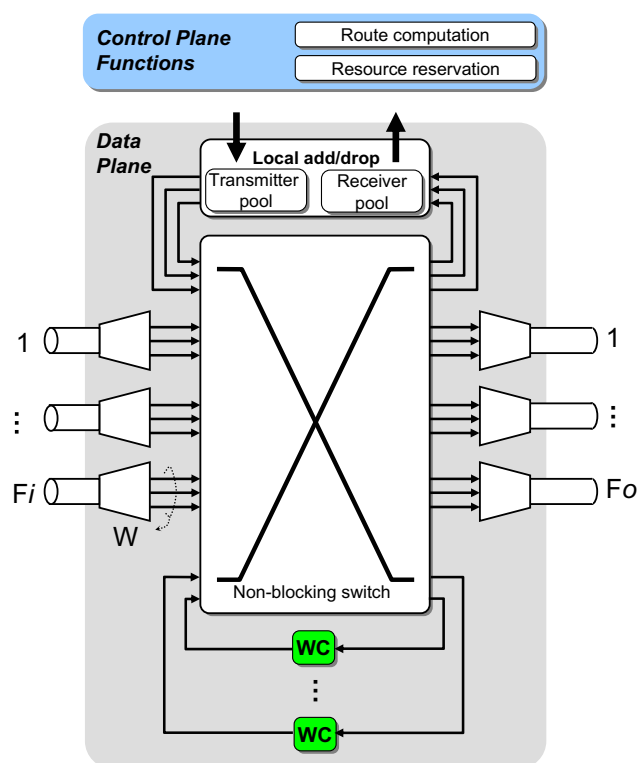


Fig. 1. Node architecture with shared converter bank.

direction. Due to distributed control, each node handles both control and data plane functions.

The control plane is responsible for route computation and resource reservation, and accomplishes these functions by exchanging GMPLS control messages. The data plane is responsible for switching data belonging to different connections, based on control plane instructions.

The switch architecture is assumed to be non-blocking, and to have enough add-drop interfaces to establish all locally terminated connections. With reference to pass-through connections, a WC is utilized if the same wavelength is not available at both the incoming and outgoing port. The WCs are assumed to be tunable over the whole wavelength range and shared by all ports as a per-node converter bank [14].

3.2. Network architecture

In order for the network to be independent of a given traffic matrix, all the spans in the network comprise the same number of wavelengths, and all the nodes have the same number of WCs in their converter bank. All the nodes are interconnected through bidirectional fibers spans. Each of these fiber spans comprises unidirectional fibers in each direction. This means that, e.g., given S bidirectional spans, the number of unidirectional spans is $2 \cdot S$. Since the investigated network is wavelength routed, the requested connections are set up as unidirectional connections at the granularity of one unidirectional wavelength channel.¹

¹ We consider only unidirectional connections here, since issues such as outdated wavelength status and forward reservation, which are out of the scope of this study, impact the performance of bidirectional lightpath provisioning.

The control messages are transmitted along the same route as the actual data. Each node can be the source or the destination of a connection request.

3.3. Route computation and signaling

The route and wavelength of a new connection are often found in decoupled steps [43]. To emulate the behavior of OSPF-TE in our model, the network topology is discovered at the beginning of a simulation run and the route for a connection is computed by running the Dijkstra shortest path algorithm. The simplified modelling of OSPF-TE is justified since routing is not in the scope of this paper.

The RSVP-TE protocol, which is used to reserve the necessary resources on the identified route, is modelled in detail. RSVP-TE uses a two way reservation mechanism, where a resource is requested by propagating a so-called *Path* message between the source and destination node. Upon reception of a *Path* message, the destination node chooses a wavelength (i.e., a label) and propagates a *Resv* message upstream towards the source. In the involved nodes, the resources are only reserved upon reception of such a *Resv* message. The concept of exchanging *Path* and *Resv* messages is illustrated in Fig. 2. The actual wavelength can be chosen based on different wavelength assignment schemes, which are further discussed in Section 4.

When the failure of a span is initiated, the upstream failure adjacent nodes is notified. It then checks which connections are affected by the failure and initiates recovery actions, either through notification of the end nodes or by itself depending on the chosen restoration method.

4. Wavelength assignment schemes

WCs play an important role in simplifying the wavelength assignment in all-optical networks, both during the provisioning and the restoration phase. Their availability avoids the wavelength continuity constraint and consequently increases the provisioning and the recovery probability, because also non-wavelength continuous

paths may be chosen. Moreover, WC availability avoids reservation collisions when two or more concurrent signaling sessions contend for the same wavelength channels. This benefit is particularly important in the restoration phase, when many recovery attempts are almost simultaneously triggered and must share a limited amount of available network resources. Additionally, if span or segment restoration is exploited, WCs are very useful to match the recovery path wavelength with the stub wavelength.

As WCs are a highly limited resource, their usage must be rigorously controlled both during the provisioning and the restoration phase. However the standard wavelength assignment performed with GMPLS signaling is not suited to achieve this goal. To perform a WC-saving wavelength assignment, an additional signaling object, called Suggested Vector, has been introduced in [23,27]. In the following sections we describe the standard GMPLS wavelength assignment, exploiting the Label Set object alone, and the enhanced wavelength assignment, exploiting the Suggested Vector object together with the Label Set.

4.1. Standard wavelength assignment – the Label Set

The Label Set (LS) is a standard protocol object defined in [39] with the purpose of simplifying the wavelength assignment in dynamic WRONs. The resulting scheme enforces the best possible wavelength assignment with current protocols and is considered as a benchmark in this study.

The LS object is propagated within the *Path* message and consists of an array of labels, containing only the wavelengths acceptable by the upstream node. The LS is typically used by an upstream node to control the selection of labels by downstream nodes. If no WCs are available, the LS is restricted to contain labels ensuring a wavelength-continuous path. If WCs are available, the LS contains all the labels that are available on a given span. Once the signaling session arrives at the destination node, a label (i.e., a wavelength) within the LS is reserved according to a tie-breaking policy (e.g., first-fit or random). The chosen label is propagated as far as possible along the reverse path. If a




Fig. 2. Exchange of GMPLS signaling messages.

node cannot reserve the current wavelength and a WC is available, a new wavelength is chosen within the node's received LS; otherwise the connection is blocked.

In the example illustrated in Fig. 3, the LS is propagated within the *Path* message from source to destination. At the destination, first-fit tie-breaking policy is enforced. The lowest-indexed label in the received LS (i.e., λ_1) is chosen and propagated upstream within the *Resv* message. However, we can notice that a wavelength conversion becomes necessary at the second node (highlighted in Fig. 3) because λ_1 is not available on the first hop. On the contrary, no WC would be needed if wavelength λ_3 had been chosen at destination. From this example it is clear that, using only the Label Set, the destination node lacks information to choose the wavelength that minimizes WC usage.

4.2. Enhanced wavelength assignment – the Suggested Vector

A WC-saving wavelength assignment can be ensured by using a novel signaling protocol extension, called Suggested Vector (SV). The SV can be used in the *Path* message together with the LS. The SV is an array of the same size as the LS, containing information allowing ranking of the labels within the LS.

In this study, we rank the labels according to the number of WCs each label requires to set up a connection. The source node fills the SV with zeros, because all available labels can be reached without WCs. Then the SV propagates along the route. If a label was available on the previous hop, its SV values are propagated without modification. If a label was not available on the previous hop, its SV value is calculated by adding one to the minimum SV value of the previous hop (since the choice of this label implies one conversion at the current node). When the destination node is reached, the label with the lowest SV value (i.e., the one requiring fewest WCs) is chosen and propagated towards the source. If two labels have equal SV values, a tie-breaking policy needs to be applied. If a node cannot

further use the current wavelength and a WC is available, a new wavelength with minimum value of the previous hop SV is chosen; otherwise the connection is blocked.

In the example of the SV scheme operation shown in Fig. 4, the SV is updated at each node during *Path* message propagation. At the destination the SV contains, for each label in the LS, the minimum number of WCs necessary to set up the request using that label: λ_3 is chosen because it requires zero WCs.

5. Simulation study

In this study, we investigate how the recovery performance of span, segment and end-to-end restoration is affected by a standard and an enhanced label assignment scheme in a network, when the number of available WCs per node is varied. The simulations are carried out in OPNET Modeler [44], and evaluated in the Pan-European triangular topology network [45], consisting of 28 nodes and 61 spans, yielding a nodal degree of 4.36, as illustrated in Fig. 5.

5.1. Case studies

In order to evaluate the performance of the three restoration methods with standard and enhanced GMPLS signaling, we consider two fundamentally different cases:

1. comparison of the standard and enhanced scheme exclusively in the restoration phase;
2. comparison of the two schemes in both the provisioning and restoration phase (real network scenario).

Case 1. Isolated restoration phase: The connections in the network are provisioned using unlimited conversion capability, and the wavelength assignment scheme in the provisioning phase is fixed to the standard scheme to




Fig. 3. Standard wavelength assignment.

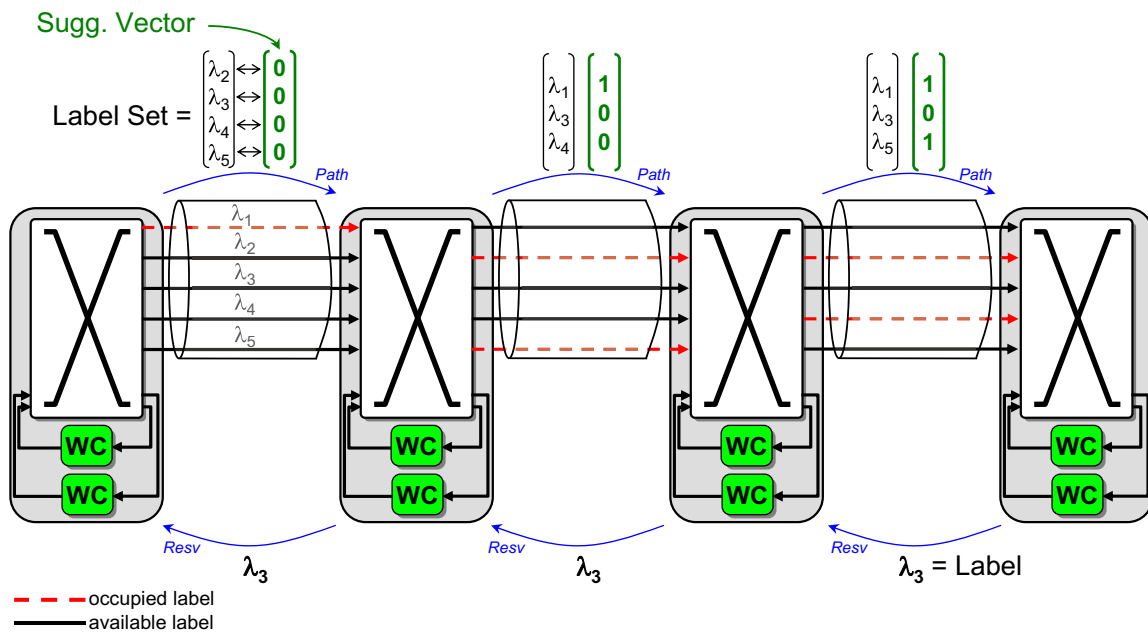


Fig. 4. Enhanced wavelength assignment.

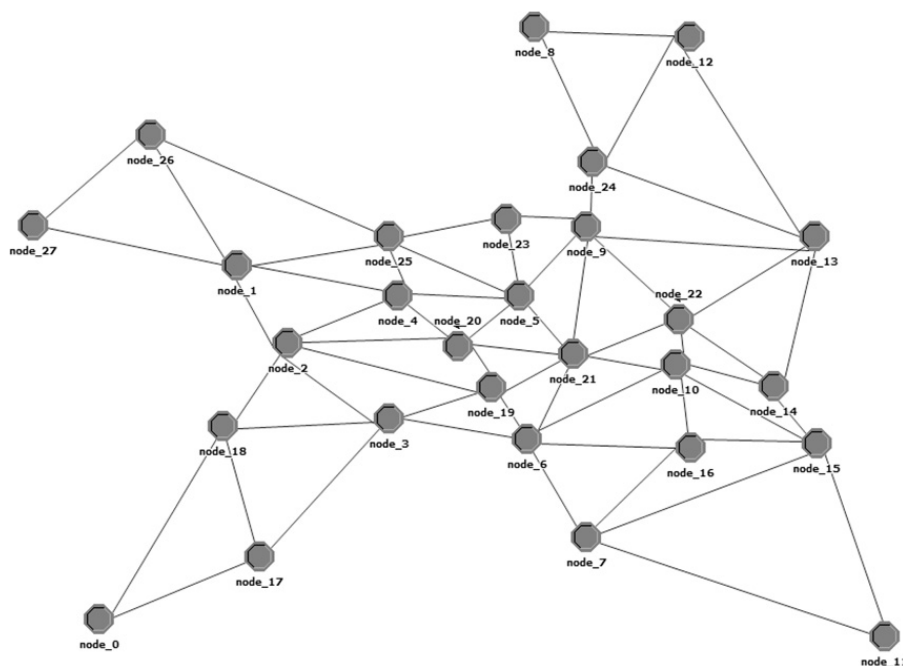


Fig. 5. Pan-European network topology.

ensure the same connection setup prior to the restoration phase. Once the desired load is reached, the number of WCs available for the restoration phase becomes limited, and the recovery of the span failures is attempted, using either wavelength assignment scheme.

Case 2. Correlated provisioning and restoration phase: In this case we investigate the restoration performance of a network which uses the same wavelength assignment scheme in both the provisioning and the restoration phase. This scenario is also the most realistic one from a real-life network operation point of view. In this case, all nodes are

given a specific number of WCs at the beginning of the provisioning phase, and no separate WCs will be assigned for the restoration phase. This means that the converter-saving property of the enhanced scheme can potentially help to save WCs already in the provisioning phase, which will keep the WCs available for the restoration phase.

5.2. Simulation parameters

We evaluate the three investigated restoration methods (span, segment and end-to-end), using either standard or

Table 1
Simulation parameters

Parameter	Value
Network topology	Pan-European network
Wavelengths per span	10
Wavelength converters per node	0–20
Connection interarrival time	10 s
Connection end nodes distribution	Uniform
Wavelength assignment scheme	Standard (LS) or enhanced (SV)
Setup reattempts provisioning	None
Setup reattempts restoration	k-Shortest paths with crankback until no route can be found
Connection holding time	Infinite
Tie breaking policy	First-fit
Average network load A	0.5
Span failures	All spans consecutively failed
Restoration method	Span, segment, end-to-end

enhanced wavelength assignment. The number of WCs available per node is varied and the average network load is kept constant at 0.5, since previous studies have shown that results are only weakly dependant on the load [12]. Table 1 contains an overview of the parameter values used in both simulation cases.

5.3. Simulation execution

In the provisioning phase of the simulation, the network is populated with connections of indefinite holding time, meaning that once provisioned, the connections are not torn down or changed. Each connection is unidirectional, occupying one wavelength/label. Labels are assigned using either the standard LS or enhanced SV wavelength assignment scheme. Ties are broken using the first-fit policy. The source and destination pairs are uniformly distributed over the entire network. If a connection request experiences blocking in the setup phase, it is dropped and a new source-destination pair is chosen. The provisioning phase continues until the desired average network load is reached. The average network load A is defined as the number of occupied unidirectional spans divided by the total number of spans in the network. That is

$$A = \frac{N_{\text{Occupied}}}{2 \cdot S \cdot W}, \quad (1)$$

where N_{Occupied} is the number of occupied unidirectional spans in the network, S is the number of spans in the network and W the number of wavelengths on each span.

When the desired load is achieved, the restoration phase begins and one span failure is simulated at a time. Typically, the restoration time is much shorter than the provisioning interarrival time. Hence, we assume that no new connection requests arrive during failure recovery, which allows us to make a clear distinction between the provisioning and restoration phases of the simulation. The stub resources are kept occupied (i.e., no stub-release is executed) and recovery with span, segment or end-to-end restoration is attempted, using either the standard or enhanced wavelength assignment scheme.

The setup of a restoration path can be blocked due to lack of free WCs or wavelength channels on the route. In this case, the restoration is retried, excluding the block-

ing-causing span from the route computation. Information about the blocking span is sent back to the node performing the route calculation using the crankback mechanism described in [46], where the RSVP-TE *PathErr* message carries the information. The source node caches all blocking locations for a particular connection until it has been successfully restored or has been deemed unrecoverable. When all affected connections are either restored or found to be unrecoverable, the network is reverted to its pre-failure state before the next span failure is simulated. This procedure is repeated for all spans in the network.

The following performance metrics are adopted:

- Recovery percentage:

$$RP = \frac{\text{Conn}_{\text{Recovered}}}{\text{Conn}_{\text{Failed}}}; \quad (2)$$

- WC usage per recovered connection:

$$\text{WCRC} = \frac{\text{WC}_{\text{Restoration phase}}}{\text{Conn}_{\text{Recovered}}}; \quad (3)$$

- Wavelength channel usage in the restoration phase:

$$\text{WLR} = \sum \text{WL}_{\text{Restoration phase}}; \quad (4)$$

- Hop count of recovery paths:

$$\text{HCR} = \frac{\text{WLR}}{\text{Conn}_{\text{Recovered}}}. \quad (5)$$

6. Results

All results presented are averaged over a number of simulation runs according to Table 1 with different random seeds for each data point, and confidence intervals at 95% confidence level are given.

6.1. Case 1. Isolated restoration phase

In this section, results from case study 1 are presented. As described in Section 5.1, an identical set of connections is created using the standard signaling scheme by providing an unlimited number of WCs in the provisioning phase. In this way, the restoration methods can be compared across different signaling schemes given the same starting point. In the restoration phase, all restoration methods are evaluated using both the standard and enhanced wavelength assignment schemes. Note that the number of WCs is limited in the restoration phase.

Fig. 6 shows that end-to-end restoration achieves the best recovery percentage followed by segment and span restoration in that order. Since end-to-end restoration has a higher degree of freedom when computing restoration paths on a network-wide basis, whereas the recovery paths of span and segment restoration emerge from the upstream failure adjacent node, this behavior is expected. Above 5 WCs per node available in Restoration Phase (WCiRP), this order is interestingly broken by segment slightly outperforming end-to-end restoration. At this point, the network moves away from being WC-limited to being span-resource-limited (i.e., available wavelengths




Fig. 6. Recovery percentages for case study 1.

on the spans are the scarcest resource). Hence, the potentially shorter restoration path of segment restoration is more useful than the freedom of route selection offered by end-to-end restoration. The span restoration method achieves the highest profit from a higher number of WCiRP, since the recovery segment of the connection has to be merged to the stubs of the pre-failure path at the upstream and downstream failure adjacent nodes, thus requiring a WC with high probability. The enhanced signaling scheme gives a better recovery percentage in all cases except the

degenerate case of 0 WCiRP, where both schemes achieve the same performance, as expected. The largest gain of using enhanced signaling is obtained at 1 and 2 WCiRP, resulting in an increase of approximately 5% in recovery percentage. The advantage of enhanced signaling diminishes with increasing WCiRP (since the network is no longer WC-limited), but is still visible at 5 WCiRP.

The WC usage per recovered connection (WCRC) is shown in Fig. 7. A general tendency to observe is that the usage of WCs increases when they are widely available

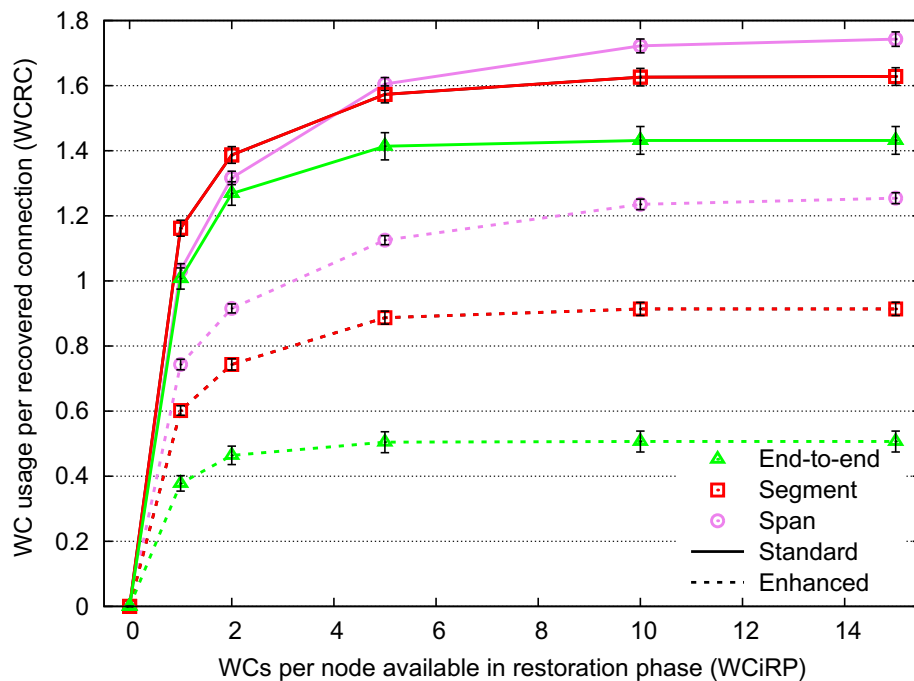


Fig. 7. WC usage per recovered connection (WCRC) for case study 1.

throughout the network. More WCs allow more connections to be restored, as can be seen in Fig. 6. It is clearly shown that the enhanced signaling schemes uses significantly fewer WCs than the standard scheme regardless of restoration method. The curves flatten from 5 to 10 WCiRP, indicating that the network is saturated with WCs and is no longer WC-limited. Using the enhanced scheme, the WC usage of the restoration methods can be ranked as end-to-end (fewest WCs), segment and span (most WCs). Using standard signaling schemes, the same ranking is seen above 10 WCiRP. But for fewer WCiRP span restoration behaves differently, performing in between end-to-end and segment restoration. This is due to the fact that the recovery percentage for span restoration is low at this point, because even though converters may be available in the network, they are not available at the needed locations (often the failure adjacent nodes). This is also illustrated in Fig. 6 by the fact that span restoration achieves the highest benefit from more WCiRP, since the recovery percentage steeply increases when they become available.

Fig. 8 illustrates the hop count of the recovery path (HCR) for the different restoration methods. End-to-end restoration has the longest restoration path, since it is restored between the connection's end nodes. Segment restoration has medium length, as paths are restored between the upstream failure adjacent node and the destination node. Span restoration, where connections are restored between the failure adjacent nodes, has the shortest restoration segment. For all methods, the hop count increases at 1 WCiRP. At this point, we also observe a large increase in recovery percentage, which means that also longer paths are recovered. When more WCiRP become available, the hop count decreases because WCs allow the use of shorter routes otherwise unavailable. The availability of wavelength converters also has an influence

on the hop count for the enhanced and the standard scheme. At few WCiRP, the enhanced scheme achieves a lower hop count, leveraging the saved converters, while the difference diminishes when WCs are no longer the limiting factor.

Fig. 9 shows the wavelength usage in the restoration phase (WLR) as a function of the WCs per node. Even though the enhanced scheme achieves a higher recovery percentage than the standard scheme, the WLR increase of the enhanced scheme compared to the standard scheme is very limited, which highlights the fact that the difference in recovery percentage is achieved by more efficient WC usage and not longer restoration paths. Comparing the different restoration methods, end-to-end restoration obtains the highest WLR, because it has the longest restoration paths and the highest recovery percentage, while in span restoration WLR is the lowest because the recovery paths just circumvent the failed span and the recovery percentage is the lowest. As expected, segment restoration lies in between. All restoration methods show a saturation behavior for increasing WCiRP, meaning that WCs are no longer a critical resource, but the network becomes span capacity limited.

6.2. Case 2. Correlated provisioning and restoration phase

For the results presented in this section, the same wavelength assignment scheme is used both in the provisioning and restoration phase, as would be the most realistic choice for real-life network operation. The influence of the different schemes on the RP, WCRC, HCR and WLR is presented.

Fig. 10 illustrates the recovery percentage obtained for the different restoration methods when the number of WCs per node given at the start of the provisioning phase

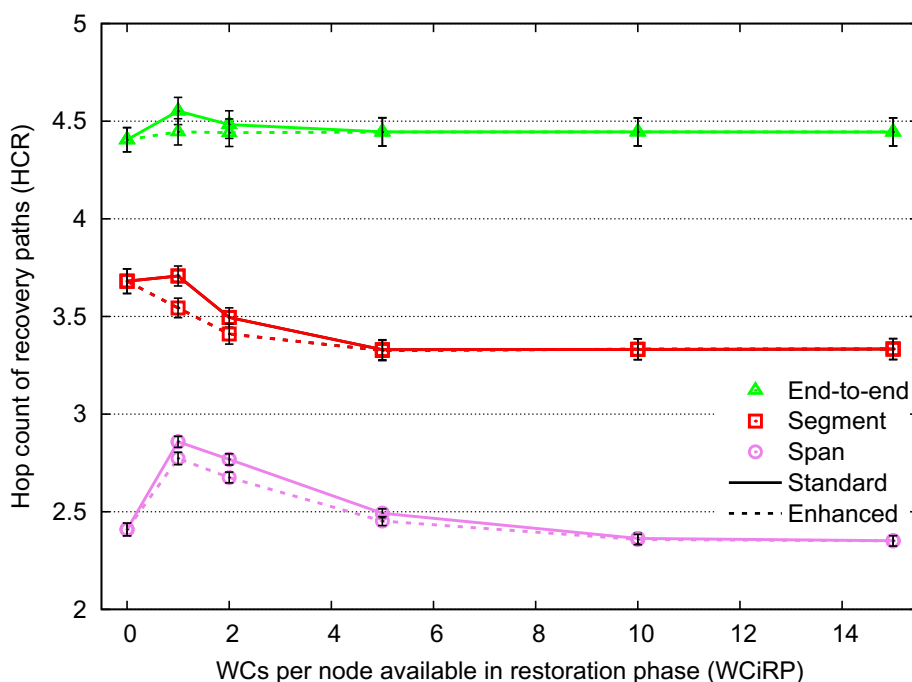


Fig. 8. Hop count of recovered connections (HCR) for case study 1.

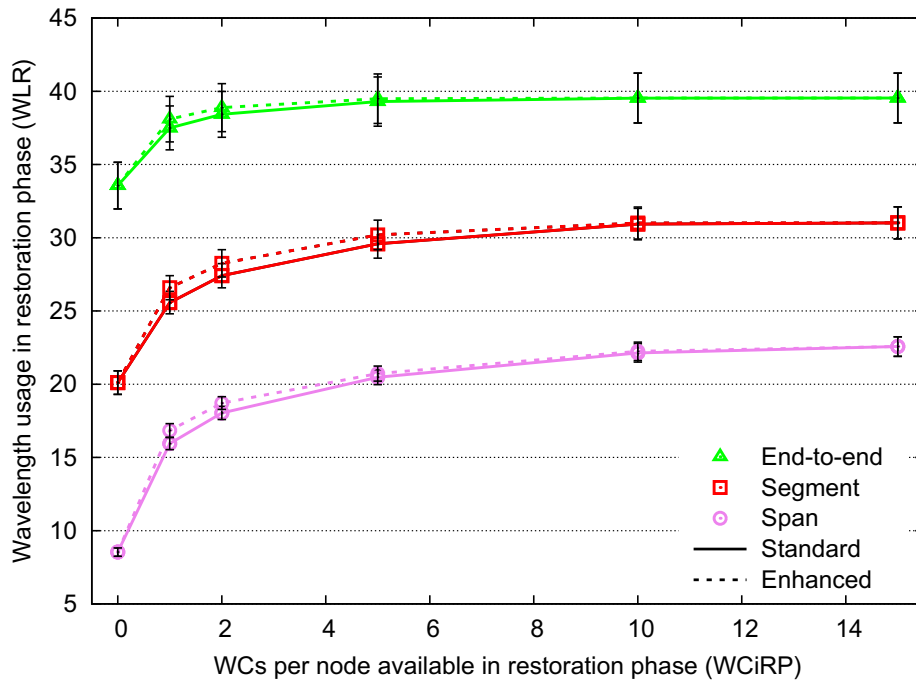


Fig. 9. Wavelength usage in the recovery phase for case study 1.




Fig. 10. Recovery percentages for case study 2.

is varied. A general observation is that the enhanced wavelength assignment scheme outperforms the standard wavelength assignment scheme for all restoration methods, except at 0 converters where both schemes perform the same as expected. Furthermore, we also observe a general tendency that the recovery percentage increases if more WCs are added. However, at 2 WCs using span restoration with the standard scheme, this trend is interestingly

broken. This is due to the fact that if no WCs are available, only wavelength continuous connections can be allocated, which also ensures order in the wavelength assignment on the different spans. With very few WCs however, the first few requests may use WCs and hence create disorder in the wavelength assignment on the different spans, without there being enough WCs to use the wavelengths that are available between groups of used wavelengths later on,

which results in a lower recovery percentage. Span restoration is particularly vulnerable to disorder in the wavelength assignment, since the restoration path must be merged with the remaining parts (i.e., stubs) of the pre-failure path at the two failure-adjacent nodes. The reason why the difference in recovery percentage between the standard and enhanced schemes is bigger compared to case 1 is that the enhanced wavelength assignment scheme can save WCs already in the provisioning phase,

which in turn will be available in the restoration phase. Span restoration achieves the largest performance increase, going from an average restoration percentage of 63–88% at 5 WCs per node. For segment and end-to-end restoration, the increase in recovery percentage is smaller, but still significant. When 20 WCs are available per node, the difference diminishes not only between the different wavelength assignment schemes, but also between the different restoration methods.




Fig. 11. WC usage per recovered connection (WCRC) for case study 2.

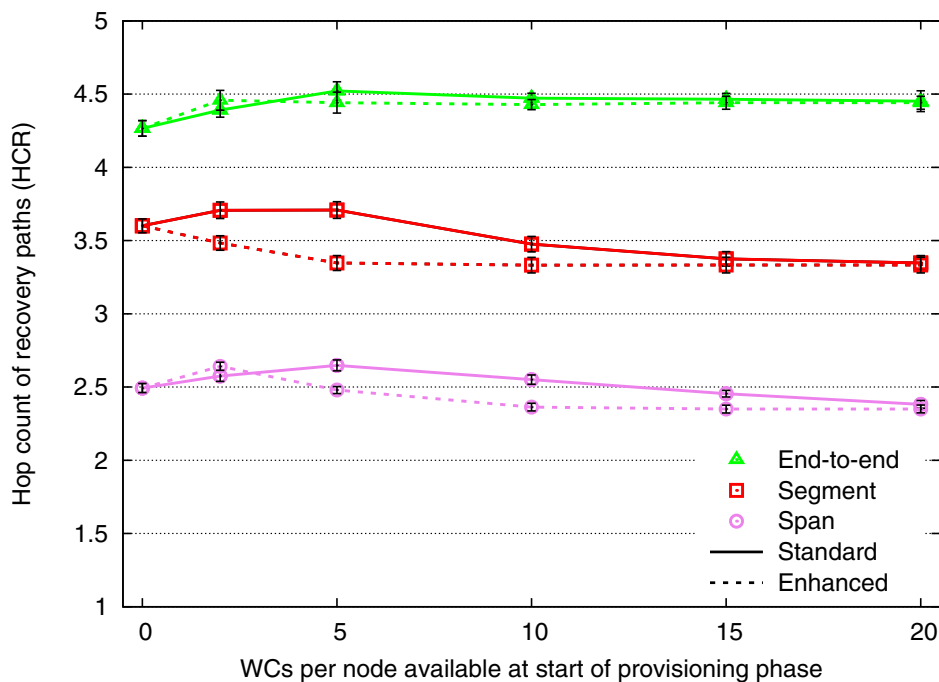


Fig. 12. Hop count of recovery path (HCR) for case study 2.

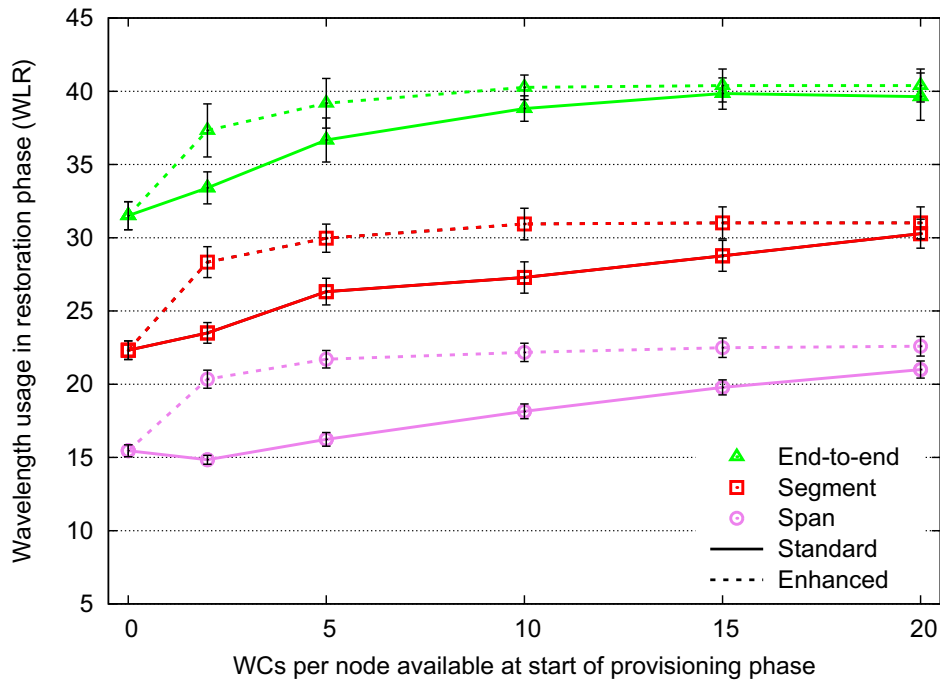


Fig. 13. Wavelength usage in the recovery phase for case study 2.

The WC usage per recovered connection (WCRC) is shown in Fig. 11. A general trend is that the usage of WCs increases when they are widely available throughout the network. The higher WC usage of the enhanced scheme at few WCs allows for more connections to be restored, as can be seen in Fig. 10. Furthermore, we observe a saturation behavior above 10 WCs, where neither the WCRC nor the RP increases. This is due to the fact that the network is becoming capacity limited and adding more WCs does not remedy this. For the standard scheme, which has a more wasteful wavelength assignment, more WCs are needed to achieve the saturation behavior. Using the enhanced scheme, the WC usage of the restoration methods can be ranked as end-to-end (fewest WCs), segment and span (most WCs). Using standard signaling schemes, the same ranking appears at high WC numbers. The reason for this ranking is that span restoration must merge the restoration path to the stubs at the two failure adjacent nodes, which is potentially WC consuming. Segment restoration only has one merger node, while end-to-end restoration entirely avoids merging.

Fig. 12 shows the hop count of the recovery path (HCR) for the different restoration methods. Span restoration, where connections are restored between the failure adjacent nodes, has the shortest restoration path. Segment restoration has a medium path length, since paths are restored between the upstream failure adjacent node and the destination node. End-to-end restoration has the longest restoration path, since the recovery path is found between the end nodes of a connection. The HCR has a similar performance for the two schemes, with slight variation at few WCs, which is due to the fact that a varying WC count allows for different connections to be restored. When WCs are abundant, both schemes achieve the same HCR.

Fig. 13 shows the wavelength usage in the restoration phase (WLR) as a function of the WCs per node. End-to-end restoration uses the most wavelength resources for its restoration paths, followed by segment and span restoration, which is due to the choice of merging nodes. We also observe that the enhanced scheme uses more wavelengths than the standard scheme, which is due to the fact that leveraging the available WCs in the restoration phase results in a higher recovery percentage, as shown in Fig. 10. Once WCs no longer are a critical resource and the network becomes span capacity limited, both signaling schemes show equal wavelength usage in each restoration method.

7. Conclusion

In this study, we investigated the effect of limited wavelength converter availability on the span, segment and end-to-end restoration methods. We simulated the performance of an all-optical network in a GMPLS controlled environment, evaluating the effect of both standard and enhanced (aiming at minimizing converter usage) wavelength assignment schemes for the different restoration methods. In two case studies, we investigated the effect of limited wavelength conversion (1) exclusively in the restoration phase, and (2) in a network operation scenario with correlated provisioning and restoration phase.

Our studies show that limited wavelength conversion availability has a different effect on the three restoration methods. The recovery percentage of span restoration is most severely affected, due to its nature of having to merge the restoration segment to the pre-recovery path at the failure adjacent nodes, which often requires a conversion between wavelengths. This is consistent with the number of converters that are used per recovered connection.

End-to-end restoration is the least affected, because it restores the affected connection between the its source and destination nodes, and hence profits from a high degree of freedom and the ability to avoid potential bottlenecks when setting up restoration paths. The performance of segment restoration lies between the previous two methods.

Another interesting observation is that if a certain number of WCs are available, all three restoration methods achieve very similar restoration performance. Since the choice of restoration method becomes less relevant for the actual recovery percentage, an operator can choose a restoration method based on other criteria, such as notification time or manageability, given that enough WCs are available.

The benefit of the standard or enhanced wavelength assignment scheme highly depends on the number of WCs that is provided. If WCs are limited, all restoration methods can benefit from the converter saving properties of the enhanced scheme. Especially, span and segment restoration can significantly increase their performance with the enhanced scheme. If span or segment restoration is the preferred restoration method and WCs are very limited, it is advisable to use the enhanced scheme to minimize performance penalty compared to end-to-end restoration.

Acknowledgement

This work has been partially supported by the European Commission through the Network of Excellence e-Photon/ONE+.

References

- [1] P. Blair, J.B. (Ovisor Technologies), ROADMs – the traffic cops of agile optical networks, Europhotonics, European Coverage of Optics, Lasers, Imagine, Fibre Optics, Electro-Optics and Photonic Component Manufacturing, Laurin Publications, February/March 2007.
- [2] B. Mukherjee, *Optical WDM Networks*, Springer, New York, 2006.
- [3] W.D. Grover, *Mesh-Based Survivable Networks – Options and Strategies for Optical, MPLS, SONET, and ATM Networking*, Prentice-Hall, New Jersey, 2004.
- [4] J.-P. Vasseur, M. Pickavet, P. Demeester, *Network Recovery, Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*, Morgan-Kaufmann Publishers/Elsevier, Los Altos/Amsterdam, 2004.
- [5] J. Wang, L. Sahasrabudde, B. Mukherjee, Path vs. subpath vs. link restoration for fault management in IP-over-WDM networks: performance comparisons using GMPLS control signaling, *IEEE Commun. Mag.* 40 (11) (2002) 80–87.
- [6] J. Zhang, B. Mukherjee, A review of fault management in WDM mesh networks: basic concepts and research challenges, *IEEE Network* 18 (2) (2004) 41–48.
- [7] E. Mannie, D. Papadimitriou (Eds.), *Recovery (protection and restoration) terminology for generalized multi-protocol label switching (GMPLS)*, RFC 4427.
- [8] S. Ruepp, J. Buron, N. Andriolli, Wavelength-converter saving span restoration in GMPLS controlled WDM optical networks, in: Proc. Conference of Optical Internet, COIN, 2006.
- [9] J. Buron, S. Ruepp, N. Andriolli, Blocking reduction of span restoration requests in GMPLS controlled WDM optical networks, in: Proc. ECOC, 2006.
- [10] S. Ruepp, J. Buron, N. Andriolli, H. Wessing, Span restoration in optical networks with limited wavelength conversion, Proc. ChinaCom.
- [11] C. Gruber, A comparison of bandwidth requirements of path protection mechanisms, ICN 2005, LNCS 3420, 2005, pp. 133–143.
- [12] S. Ruepp, J. Buron, N. Andriolli, Increasing restorability for local-to-egress restoration in GMPLS controlled networks with limited wavelength conversion, in: Proc. ICCS, 2006.
- [13] E. Mannie (Ed.), *Generalized multi-protocol label switching (GMPLS) architecture*, RFC 3945.
- [14] B. Ramamurthy, B. Mukherjee, Wavelength conversion in WDM networking, *IEEE J. Select. Areas Commun.* 16 (7) (1998) 1061–1073.
- [15] K.-C. Lee, V. Li, A wavelength-convertible optical network, *IEEE J. Lightwave Technol.* 11 (5/6) (1993) 962–970.
- [16] J. Elmighani, H. Mouftah, All-optical wavelength conversion: technologies and applications in DWDM networks, *IEEE Commun. Mag.* 38 (3) (2000) 86–92.
- [17] S. Yoo, Wavelength conversion technologies for WDM network applications, *IEEE J. Lightwave Technol.* 14 (6) (1996) 955–966.
- [18] J. Iness, B. Mukherjee, Sparse wavelength conversion in wavelength-routed WDM optical networks, *Photon. Network Commun.* 1 (3) (1999) 183–205.
- [19] G. Shen, S. Bose, T. Cheng, C. Lu, T. Chai, Operation of WDM networks with different wavelength conversion capabilities, *IEEE Commun. Lett.* 4 (7) (2000) 239–241.
- [20] M. Swaminathan, K. Sivarajan, Practical routing and wavelength assignment algorithms for all optical networks with limited wavelength conversion, in: Proc. ICC, 2002, pp. 2750–2755.
- [21] Q.-D. Ho, M.-S. Lee, Converter-aware wavelength assignment in WDM networks with limited-range conversion capability, *IEICE – Trans. Commun.* E89-B (2) (2006) 436–445.
- [22] G. Suixiang, J. Xiaohua, H. Chuanhe, D. Ding-Zhu, An optimization model for placement of wavelength converters to minimize blocking probability in WDM networks, *IEEE J. Lightwave Technol.* 21 (3) (2003) 684–694.
- [23] N. Andriolli, J. Buron, S. Ruepp, F. Cugini, L. Valcarengi, P. Castoldi, Label preference schemes in GMPLS controlled networks, *IEEE Commun. Lett.* 12 (10) (2006) 849–851.
- [24] M. Ali, Shareability in optical networks: beyond bandwidth optimization, *IEEE Commun. Mag.* 42 (2) (2004) S11–S15.
- [25] M. Ali, Generalized sharing in survivable optical networks, *IEEE/ACM Trans. Network.* 14 (6) (2006) 1388–1399.
- [26] A.M.C.A. Koster, A. Zymolka, Tight LP-based lower bounds for wavelength conversion in optical networks, *Stat. Neerland.* 61 (1) (2007) 115–136.
- [27] N. Andriolli, J. Buron, S. Ruepp, F. Cugini, L. Valcarengi, P. Castoldi, Signaling protocol extensions for converter-saving wavelength assignment in GMPLS optical networks, in: Proc. HPSR, 2006.
- [28] A. Fumagalli, L. Valcarengi, IP restoration vs. WDM protection: is there an optimal choice?, *IEEE Network* 14 (6) (2000) 34–41.
- [29] S. Ramamurthy, L. Sahasrabudde, B. Mukherjee, Survivable WDM mesh networks, *IEEE J. Lightwave Technol.* 21 (4) (2003) 870–883.
- [30] A.V. Sichani, H.T. Mouftah, Performance evaluation of dynamic restoration techniques for survivable optical networks, in: Proc. ISCC, 2006, pp. 961–966.
- [31] Y. Luo, N. Ansari, Restoration with wavelength conversion in WDM networks, *Electron. Lett.* 38 (16) (2002) 900–901.
- [32] G. Conte, M. Listanti, M. Settembre, R. Sabella, Strategy for protection and restoration of optical paths in WDM backbone networks for next-generation internet infrastructures, *IEEE J. Lightwave Technol.* 20 (8) (2002) 1264–1276.
- [33] Y. Luo, N. Ansari, Split restoration with wavelength conversion in WDM networks, in: Proc. ICC, 2003, pp. 1423–1427.
- [34] M. Keshtgary, F. Al-Zahrani, A. Fahad, M. Jahangir, A. Jayasumana, Survivability performance evaluation of WDM networks with wavelength converters, *Photon. Network Commun.* 11 (1) (2006) 15–27.
- [35] M. Azim, X. Jiang, P. Ho, S. Horiguchi, Models of restoration probability in WDM networks employing active restoration, *Photon. Network Commun.* 10 (2) (2005) 141–153.
- [36] M. Azim, X. Jiang, P. Ho, S. Horiguchi, Behaviors of active lightpath restoration in all-optical networks with and without wavelength converters, in: Proc. of Asia Pacific Conference, 2005, pp. 941–945.
- [37] Z. Zhang, Y. Yang, On-line optimal wavelength assignment in WDM networks with shared wavelength converter pool, *IEEE/ACM Trans. Network.* 15 (1) (2007) 234–245.
- [38] P.-H. Ho, H. Mouftah, A novel survivable routing algorithm for shared segment protection in mesh WDM networks with partial wavelength conversion, *IEEE J. Lightwave Technol.* 22 (8) (2004) 1548–1560.
- [39] L. Berger (Ed.), *Generalized multi-protocol label switching (GMPLS) signaling functional description*, RFC 3471.

- [40] D. Papadimitriou, E. Mannie (Eds.), Analysis of generalized multi-protocol label switching (GMPLS)-based recovery mechanisms (including protection and restoration), RFC 4428.
- [41] J. Zheng, B. Zhang, H. Mouftah, Dynamic path restoration based on multi-initiation for GMPLS-based WDM networks, in: Proc. ICC, 2004, pp. 1639–1643.
- [42] S. Albarrak, Pre-allocated restoration technique in a distributed GMPLS-based IP/WDM mesh network, in: Proc. of Information and Computer Technologies (ICTTA), 2006, pp. 2236–2241.
- [43] H. Zang, J. Jue, B. Mukherjee, A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks, Opt. Networks Mag. 1 (1) (2000) 47–60.
- [44] OPNET Technologies, Inc. <<http://www.opnet.com>>.
- [45] R. Inkret, A. Kuchar, B. Mikac (Eds.), Extended final report of COST Action 266 – advanced infrastructure for photonic networks, Faculty of Electrical Engineering and Computing, University of Zagreb, 2003.
- [46] A. Farrel (Ed.), Crankback signaling extensions for MPLS and GMPLS RSVP-TE, RFC 4920.



Sarah Ruepp obtained a B.Sc. in electronic and computer engineering from the Engineering College of Copenhagen in 2002. In 2004, she obtained a M.Sc. in telecommunications from the Technical University of Denmark. She is currently working towards a Ph.D. degree on the topic of “Dynamic Protection of Optical Networks” at the Networks Competence Area, DTU Fotonik, Technical University of Denmark. Her research interests include topics related to designing, controlling and modelling of reliable optical networks.



Nicola Andriolli received the Laurea degree in telecommunications engineering from the University of Pisa, Italy, in 2002, and the Diploma and Ph.D. degrees from Scuola Superiore Sant’Anna, Pisa, in 2003 and 2006, respectively. He was a visiting graduate student at BUTE, Budapest, Hungary, and at DTU Fotonik, Kgs. Lyngby, Denmark, and he was a visiting trainee at NICT, Tokyo, Japan. He is currently an assistant professor with Scuola Superiore Sant’Anna. His research interests include network modelling and simulation,


QoS and fault tolerance in optical networks, GMPLS architecture, and optical packet switching.



Jakob Buron received his M.Sc. Eng. in telecommunications engineering from the Technical University of Denmark in 2004, working on power consumption and dynamics of semiconductor optical amplifiers. He holds a Ph.D. degree in optical network design from the same university. His current research interests include network modelling, optical network architectures, survivability and design of optical switching nodes.



Lars Dittmann was born in 1962 and received the M.Sc. EE and Ph.D. from the Technical University of Denmark in 1988 and 1994, and is currently professor at the university within the area of integrated networks. He has since January 1999 been heading the network competence area (covering both optical and electrical networks) within DTU Fotonik at the Technical University of Denmark and was prior to that responsible for electronic switching and ATM networks at the Center for Broadband Telecommunication. Main research area is network architectures and performance.



Lars Ellegaard works as senior Systems Architect at Vitesse focusing on Carrier Ethernet and optical transmission. In prior positions he worked with, and became responsible for, definition and specification of the Tellabs 6300 portfolio of Ethernet/MPLS, SDH and WDM products. He has since 1986 held several positions as development engineer and project manager within R&D, and has also participated in more EU research programs. He has co-supervised several M.Sc. and Ph.D. projects carried out at DTU Fotonik. He received his M.Sc. EE and his Industrial Ph.D. from the Technical University of Denmark in 1985 and 1989, respectively.