

# Idle Protection Capacity Reuse in Multiclass Optical Networks

Nicola Andriolli, Alessio Giorgetti, Luca Valcarenghi, *Member, IEEE*, and Piero Castoldi, *Associate Member, IEEE*

**Abstract**—Optical networks carrying traffic belonging to different survivability classes must ensure not only the proper survivability differentiation but also the efficient network resource utilization. Current approaches improve network resource utilization by carrying low-class preemptable traffic along high-class spare resources, which are idle (i.e., unutilized) during failure-free conditions. In this paper, the proposed idle protection capacity reuse (IR) framework generalizes the practice of reutilizing idle spare resources to both failure-free and failure conditions. The IR framework is based on the idea of exploiting idle high-class lightpath protection resources not only for routing low-class lightpaths during the provisioning phase (i.e., provisioning-phase IR, P-IR) but also for dynamically restoring low-class lightpaths during the recovery phase (i.e., recovery-phase IR, R-IR). Both P-IR and R-IR have the potential to improve network utilization while providing multiclass lightpaths with the required survivability differentiation. The numerical evaluation shows that the utilization of R-IR in dynamic restoration results in an improved low-class lightpath survivability with respect to the utilization of stub release. Moreover, if P-IR and R-IR are alternatively exploited, then the former improves the provisioning performance, while the latter guarantees a higher survivability. In the end, if P-IR and R-IR are concurrently employed in the provisioning and in the dynamic restoration of two different low classes of lightpaths, respectively, an inherent survivability differentiation is achieved.

**Index Terms**—Capacity reuse, distributed control, dynamic restoration, multiclass optical networks, shared path protection (SPP), stub release (SR), survivability differentiation.

## I. INTRODUCTION

**N**EXT-GENERATION optical networks will likely support different services requiring different degrees of fault tolerance [1]–[3]. On the one hand, providing different degrees of fault tolerance requires more complex resilience schemes; on the other hand, it often saves network resources. A common practice for saving network resources is to exploit idle (i.e., unutilized) high-class protection resources to transmit low-class traffic under failure-free network conditions [4]–[8]. Furthermore, in multilayer networks, spare resources allocated for protecting connections at one layer (e.g., the optical layer) can be shared by upper layer (e.g., the IP layer) resilience schemes, as in the common-pool approach [9], [10]. In this way, it is

possible to recover connections from failures affecting different network layers while improving network utilization.

In this paper, the practice of reutilizing idle protection resources within a single network layer is generalized through a novel and more comprehensive idle protection capacity reuse (IR) framework. The IR framework aims to enhance network utilization while providing different connection classes with the required survivability differentiation.<sup>1</sup> The schemes defined within the IR framework exploit idle resources allocated as protection of high-class connections to ease not only the provisioning (i.e., provisioning-phase IR, P-IR) but also the dynamic recovery (i.e., recovery-phase IR, R-IR) of unprotected low-class connections.

R-IR differs from the common-pool approach in that it is applied to a single-layer multiclass scenario instead of a multilayer single-class scenario. Indeed, the common-pool approach is based on an interlayer sharing of idle protection resources to ease the recovery of single-class different granularity connections affected by failures at different layers. In contrast, R-IR is based on an intralayer sharing of idle high-class protection resources to ease the recovery of low-class same granularity connections contemporarily affected by the same failure (e.g., a link disruption).

The scenario to which IR-based schemes are applied includes one high all-optical connection class (i.e., gold-class lightpath) and two low all-optical connection classes (i.e., silver- and bronze-class lightpaths). Gold lightpaths require full recovery upon any single link failure, and thus, they are provided with shared path protection (SPP) [11]–[14]. Silver lightpaths do not use IR during provisioning, i.e., they are allocated on free resources, while they resort to R-IR for their dynamic restoration. Bronze lightpaths utilize P-IR during provisioning, while no IR is used during recovery.

With respect to [15]–[17], this paper systematically defines the IR framework for the first time, and it introduces novel extensions to the SPP implementation proposed in the study in [12] for supporting it. In addition, the comparison between the performance achieved by the use of stub release (SR) [18], [19] and of R-IR during dynamic restoration is presented.

Numerical results show that the advantage of R-IR utilization is twofold. First, it guarantees a higher likelihood of surviving a failure to the lightpaths utilizing it (i.e., silver lightpaths), as compared to the case in which only free resources are utilized. Furthermore, its utilization is more beneficial than the utilization of SR for dynamically restoring silver lightpaths. Instead,

<sup>1</sup>The survivability is defined as the ratio between the number of connections surviving a failure and the number of connections active upon failure.

Manuscript received October 11, 2006; revised February 1, 2007. This work was supported by NoE e-photon/ONE+ (IST-6FP-027497).

The authors are with the Scuola Superiore Sant'Anna, 56124 Pisa, Italy (e-mail: nicola.andriolli@sssup.it; alessio.giorgetti@sssup.it; luca.valcarenghi@sssup.it; piero.castoldi@sssup.it).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JLT.2007.893928

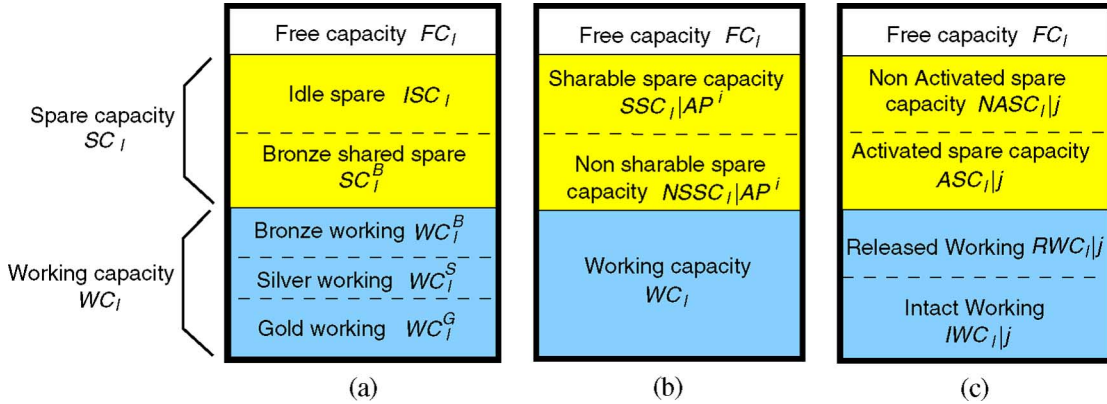


Fig. 1. Capacity categories along link  $l$ . (a) AP computation. (b)  $BP^i$  computation. (c) RP computation upon failure of link  $j$ .

P-IR ensures better provisioning performance to bronze lightpaths at the expense of a lower survivability upon failure. In general, if both silver and bronze lightpaths are contemporarily routed in the network, then an inherent survivability differentiation is obtained: Silver-lightpath survivability is higher than bronze-lightpath survivability because silver lightpaths benefit from R-IR, while bronze lightpaths may be preempted. Moreover, bronze-lightpath survivability strongly depends on the amount of spare capacity allocated for gold lightpaths, while silver-lightpath survivability does not.

## II. IDLE PROTECTION CAPACITY REUSE

The idle protection capacity reuse (IR) framework consists of the definition of a hierarchy of lightpath classes and of schemes that differently exploit unutilized (i.e., idle) spare resources during the provisioning phase (i.e., P-IR) or the recovery phase (i.e., R-IR). Conceptually, several hierarchies of lightpath classes can be devised by defining a high class, based on an arbitrary protection scheme, and several low classes, based on schemes consisting of various combinations of P-IR and R-IR relaxing (or not) the exclusive use of specific resources during the provisioning or the recovery phase.

For proof of concept, in this paper, a three-class scenario (one high class and two low classes) and two IR-based schemes, namely the R-IR scheme and the P-IR scheme, are considered. The R-IR scheme is a dynamic restoration scheme that exploits R-IR to increase the probability of a low-class lightpath to overcome network failures. The P-IR scheme is a provisioning scheme that exploits P-IR to increase the probability of a low-class lightpath to be established.

High-class (i.e., gold) lightpaths are provided with SPP: During lightpath provisioning, each gold lightpath is assigned an active path (AP) and a shared risk link group (SRLG) disjoint backup path (BP). Two low-class lightpaths (i.e., silver and bronze) are considered. Silver lightpaths are assigned an AP provisioned by just utilizing free network resources. Disrupted silver APs are recovered by exploiting the R-IR scheme. In the R-IR scheme, a restoration path (RP) is dynamically computed upon failure occurrence utilizing the spare resources that are not involved in the protection of disrupted gold lightpaths besides the free capacity. In addition, silver RPs may preempt bronze APs routed along the gold protection resources, if

necessary. Bronze lightpaths are provisioned by exploiting the P-IR scheme. In the P-IR scheme, bronze APs are allowed to utilize the idle protection resources of gold lightpaths besides the free capacity. Upon failure, bronze APs may be preempted by the activation of gold BPs and silver RPs. Disrupted and preempted bronze APs just utilize free capacity to dynamically find the RPs.

SR may be optionally utilized for low-class lightpath dynamic restoration. The SR consists in reutilizing the bandwidth that was previously used by disrupted APs for dynamic restoration. Therefore, when the SR is applied, RPs are computed that also consider the network resources released by APs disrupted by the failure, thus potentially improving their likelihood of being set up.

The following sections categorize the utilized link capacities and detail the proposed multiclass lightpath provisioning and restoration schemes.

### A. Link Capacity Categories

Each lightpath is assumed to require the amount of bandwidth provisioned by one wavelength. Hence, the various link capacity categories are expressed in terms of number of wavelengths. Fig. 1 shows the total capacity  $W_l$  along a generic link  $l$ .  $W_l$  is classified into three categories: free capacity ( $FC_l$ ), spare capacity ( $SC_l$ ), and working capacity ( $WC_l$ ).  $FC_l$  is a capacity that can be reserved to establish any AP, BP, and RP.  $SC_l$  and  $WC_l$  are further divided into subcategories as a function of the considered phase (i.e., provisioning or recovery) and of the lightpath to be established (i.e., AP, BP, or RP).

During the provisioning phase and upon establishment of any AP, the classification shown in Fig. 1(a) is valid.  $SC_l$  consists of two subcategories: the idle  $SC_l$  ( $ISC_l$ ) allocated for carrying gold BPs and available for carrying bronze APs and the bronze shared  $SC_l$  ( $SC_l^B$ ), which is shared between allocated gold BPs and bronze APs.  $WC_l$  consists of three subcategories: gold, silver, and bronze  $WC_l$  ( $WC_l^G$ ,  $WC_l^S$ , and  $WC_l^B$ ), which are the capacities utilized by gold, silver, and bronze APs, respectively.

During the provisioning phase for the computation of  $BP^i$  (i.e., the BP of the  $i$ th gold lightpath routed along  $AP^i$ ), the  $WC_l$  classification remains unchanged. Instead, the  $SC_l$  is reclassified into two subcategories: sharable  $SC_l$  ( $SSC_l|AP^i$ ) and non-sharable  $SC_l$  ( $NSSC_l|AP^i$ ), as shown in Fig. 1(b) and in

the study in [14].  $SSC_l|AP^i$  is the  $SC_l$  reserved for BPs of gold lightpaths not belonging to the same SRLG of  $AP^i$  (i.e., having their AP disjoint from  $AP^i$ ).  $NSSC_l|AP^i$  is the  $SC_l$  reserved for BPs of gold lightpaths belonging to the same SRLG of  $AP^i$  (i.e., having their AP partially overlapping with  $AP^i$ ).

During the recovery phase, upon failure of link  $j$ , the  $SC_l$  is divided into activated  $SC_l$  ( $ASC_l|j$ ) and non-activated  $SC_l$  ( $NASC_l|j$ ), as illustrated in Fig. 1(c).  $ASC_l|j$  is the  $SC_l$  activated to protect disrupted gold-class lightpaths.  $NASC_l|j$  is the  $SC_l$ , which is not utilized by gold BPs upon failure of link  $j$ . Moreover, if SR is implemented upon failure of link  $j$ , the  $WC_l$  is divided into released  $WC_l$  ( $RWC_l|j$ ) and intact  $WC_l$  ( $IWC_l|j$ ) on link  $l$ .  $RWC_l|j$  is the capacity released by lightpaths disrupted by the failure of link  $j$ .  $IWC_l|j$  is the capacity utilized by lightpaths not affected by failure of link  $j$ .

### B. Multiclass Lightpath Provisioning and Restoration

Gold-, silver-, and bronze-class lightpath APs, BPs, and RPs are computed by source nodes by applying the Dijkstra's routing algorithm to the network graph  $G(\mathcal{N}, \mathcal{L})$ . In  $G(\mathcal{N}, \mathcal{L})$ ,  $\mathcal{N}$  and  $\mathcal{L}$  represent the set of  $N = |\mathcal{N}|$  nodes and  $L = |\mathcal{L}|$  unidirectional links, respectively. The graph  $G(\mathcal{N}, \mathcal{L})$  is modified, depending on the type of path to be found, as follows.

1) *Provisioning Phase*: To compute both gold and silver APs, links with just  $FC_l$  are considered (i.e., links where  $FC_l > 0$ ). For bronze AP computation, the P-IR scheme is applied: Links with  $FC_l$  or  $ISC_l$  are considered (i.e., links where  $FC_l + ISC_l > 0$ ), assigning a negligible cost to links with  $ISC_l$  (i.e., links where  $ISC_l > 0$ ) to minimize  $FC_l$  utilization. For the computation of a gold BP, after pruning all links used by the relative gold AP, links with  $FC_l$ ,  $SSC_l|AP$ , or  $WC_l^B$  are considered (i.e., links where  $FC_l + SSC_l|AP + WC_l^B > 0$ ), assigning a negligible cost to links with  $SSC_l|AP$  or  $WC_l^B$  (i.e., links where  $SSC_l|AP + WC_l^B > 0$ ) to maximize the capacity sharing.

2) *Recovery Phase*: Upon failure of link  $j$ , each node originating disrupted lightpaths sequentially triggers the recovery procedures for gold, silver, and bronze lightpaths. Gold lightpaths are recovered by activating BPs: The activation of the spare resources reserved for gold lightpaths implies the preemption of bronze APs that share the  $ASC_l|j$ . The dynamic restoration of silver lightpaths utilizes the R-IR scheme: Silver RPs are computed by considering links with  $FC_l$  or  $NASC_l|j$  (i.e., links where  $FC_l + NASC_l|j > 0$ ). Thus, silver RPs may preempt bronze APs that share  $NASC_l|j$ . The dynamic restoration of bronze lightpaths (disrupted by the failure or preempted by gold BPs or silver RPs) is performed by computing bronze RPs just on links with  $FC_l$  (i.e., links where  $FC_l > 0$ ). Bronze RPs cannot be preempted by silver RPs. If the SR is applied, both silver and bronze RPs are computed by also considering links with  $RWC_l|j$  (i.e., links where  $RWC_l|j > 0$ ).

## III. DISTRIBUTED IMPLEMENTATION

Several distributed implementations of SPP are proposed in the literature, which can be extended for supporting IR-based schemes and SR in the considered multiclass optical network

scenario. While the works in [15]–[17] propose an extension of the SPP implementation described in the study in [13], this paper extends the one introduced in the study in [12]. This choice is driven by the fact that the SPP implementation proposed in the study in [12] is not only fully distributed but also minimizes the amount of network state information to be flooded throughout the network and maintained in each network node.

The following sections detail the network state information stored in each network node, the network state information update procedure, and how the information used for lightpath computations is retrieved. In the end, the SR implementation is described. The description is general in the sense that all three lightpath classes are supposed to be concurrently present in the network: If one of the low-class lightpaths is missing, the relative information is simply disregarded.

### A. Stored Network State Information

In the rest of this paper, the downstream node of each unidirectional link is referred to as the link master node, and it is assumed to know the up-to-date state of all incoming links [5].

To perform the network graph  $G(\mathcal{N}, \mathcal{L})$  modifications described in Section II-B, each node stores and updates the following set of vectors of size  $1 \times L$  that contain the relevant network state information.

- 1)  $\mathbf{f}$  is the  $FC_l$  vector, whose  $l$ th element represents the amount of  $FC_l$  on link  $l$  (i.e.,  $f_l = FC_l$ ).
- 2)  $\mathbf{s}$  is the  $SC_l$  vector, whose  $l$ th element represents the amount of  $SC_l$  on link  $l$  (i.e.,  $s_l = SC_l$ ).
- 3) **FailSelf**( $\mathbf{m}$ ) is the failself vector per each incoming link  $m$ , whose  $l$ th element (i.e.,  $\text{FailSelf}(m)[l]$ ) represents the bandwidth that is reserved on network link  $l$  to protect the failure of link  $m$ .
- 4) **FailOther**( $\mathbf{m}$ ) is the failother vector per each incoming link  $m$ , whose  $l$ th element (i.e.,  $\text{FailOther}(m)[l]$ ) represents the bandwidth that is reserved on link  $m$  to protect against failures of any other link  $l$ .
- 5)  $\mathbf{i}$  is the  $ISC_l$  vector, whose  $l$ th element represents the amount of  $ISC_l$  reserved on link  $l$  (i.e.,  $i_l = ISC_l$ ).
- 6)  $\mathbf{b}$  is the  $WC_l^B$  vector, whose  $l$ th element represents the amount of  $WC_l^B$  reserved on link  $l$  (i.e.,  $b_l = WC_l^B$ ).

Master nodes may optionally store the values of  $WC_m^S$  and  $WC_m^G$  for each incoming link  $m$ .

Network nodes directly obtain, from vector  $\mathbf{f}$ , the value of  $FC_l$ , which is needed for the computation of all types of path. The vectors  $\mathbf{s}$ , **FailSelf**( $\mathbf{m}$ ), and **FailOther**( $\mathbf{m}$ ), which were introduced in [12] for SPP implementation, are also utilized in this paper for the computations needed in the R-IR scheme, as detailed in the following sections. The two vectors  $\mathbf{i}$  and  $\mathbf{b}$  are introduced in this paper for supporting the P-IR scheme.

The amount of information to be stored in each network node is of the order of  $O(L)$  and does not depend on the network load.

### B. Network State Information Update

An extended version of the Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) [20], [21], as the one proposed

TABLE I  
UPDATE OF LINK  $m$  CAPACITY CATEGORIES:  $X^Y$  INDICATES A PATH OF TYPE  $X \in \{AP, BP, RP\}$  OF A LIGHTPATH REQUEST OF CLASS  $Y \in \{G, S, B\}$

	$WC_m^G$	$WC_m^S$	$WC_m^B = b_m$	$ISC_m = i_m$	$SC_m^B$	$FC_m = f_m$
Provisioning phase						
$AP^G$	++					--
$AP^S$		++				--
$AP^B   (ISC_m > 0)$				--	++	
$AP^B   (ISC_m = 0)$			++			--
$BP^G   (SSC_m   AP > 0)$						
$BP^G   (SSC_m   AP = 0, WC_m^B > 0)$			--		++	
$BP^G   (SSC_m   AP = 0, WC_m^B = 0)$				++		--
Recovery phase						
$BP^G   (ISC_m > 0)$	++			--		
$BP^G   (ISC_m = 0)$	++				--	
$RP^S   (SC_m > ASC_m   j, ISC_m > ASC_m   j)$		++		--		
$RP^S   (SC_m > ASC_m   j, ISC_m \leq ASC_m   j)$		++			--	
$RP^S   (SC_m = ASC_m   j)$		++				--
$RP^B$			++			--

in the study in [12], is assumed to be utilized as a signaling protocol to dynamically activate and tear down lightpaths. In addition, since network nodes have to differently manage lightpath requests of different classes, all RSVP-TE signaling messages are further extended to inform traversed nodes about the lightpath request class (i.e., gold, silver, or bronze), the path type (i.e., AP, BP, or RP), and some specific link capacity category values during either the provisioning or the recovery phase.

During the provisioning phase, each master node uses the received extended RSVP-TE messages to dynamically update vectors  $\mathbf{FailSelf}(m)$  and  $\mathbf{FailOther}(m)$  and the entries  $s_m$ ,  $f_m$ ,  $i_m$ , and  $b_m$  of vectors  $\mathbf{s}$ ,  $\mathbf{f}$ ,  $\mathbf{i}$ , and  $\mathbf{b}$  for each incoming link  $m$ . The work in [12] details the update procedure of  $s_m$  and vectors  $\mathbf{FailSelf}(m)$  and  $\mathbf{FailOther}(m)$ . Vectors  $\mathbf{f}$ ,  $\mathbf{i}$ , and  $\mathbf{b}$  are updated according to Table I, where each row describes, with a C-like notation, the link-capacity categories that are updated when a specific type of lightpath is activated. To reserve resources for BPs, network nodes must know the value  $SSC_m | AP$ , as shown in Table I. This information is obtained by including the values of  $SSC_l | AP$  for each link  $l$  in the RSVP-TE path message sent along the BP. Each master node floods the network with the updated values of  $f_m$ ,  $i_m$ ,  $b_m$ , and  $s_m$  utilizing the Open Shortest Path First-Traffic Engineering (OSPF-TE) routing protocol extensions [22]–[24]. Remote network nodes are, therefore, able to keep the vectors  $\mathbf{f}$ ,  $\mathbf{i}$ ,  $\mathbf{b}$ , and  $\mathbf{s}$  up-to-date.<sup>2</sup> The amount of information flooded within the network for routing purpose is of the order of  $O(L)$ .

During the recovery phase, BPs and RPs are activated by utilizing RSVP-TE messages based on the ones used during provisioning. To apply the updates described in Table I, the additional information about the failed-link ID (e.g.,  $j$ ) is carried in the RSVP-TE messages. Upon the activation of each BP, the  $\mathbf{FailOther}(m)[j]$  vector is locally updated by each network node as follows:  $\mathbf{FailOther}(m)[j] \leftarrow \mathbf{FailOther}(m)[j] - 1$ . Thus, net-

work nodes can locally compute  $ASC_m | j = \mathbf{FailOther}(m)[j]$  to reserve silver RP resources (see Table I).<sup>3</sup>

### C. Path Computation Information Retrieval

Each source node directly obtains, from locally stored vectors  $\mathbf{f}$ ,  $\mathbf{i}$ , and  $\mathbf{b}$ , the information needed for computing APs and bronze RPs, in accordance with the multiclass lightpath provisioning and restoration strategies described in Section II-B.

The computation of gold BPs and silver RPs requires additional information: the  $SSC_l | AP$  and the  $NASC_l | j$ , respectively. The work in [12] describes the procedures utilized for collecting the  $NSSC_l | AP$  value during the signaling phase along the gold AP. Source nodes are then able to obtain  $SSC_l | AP$  by subtracting the collected  $NSSC_l | AP$  from the  $SC_l$  value locally stored in the  $\mathbf{s}$  vector:

$$SSC_l | AP = SC_l - NSSC_l | AP = s_l - NSSC_l | AP. \quad (1)$$

To allow source nodes to compute the  $NASC_l | j$  value, this paper proposes an extension of the fault-notification message (e.g., the notify message defined in the study in [21]) to inform disrupted lightpath source nodes about the faulted link ID (e.g.,  $j$ ). In addition, the master node of the failed link inserts into the fault-notification message the locally stored  $\mathbf{FailSelf}(j)$  vector, which informs remote nodes of the  $(ASC_l | j)$  value in each network link  $l$  (i.e.,  $\mathbf{FailSelf}(j)[l] = ASC_l | j$ ). Source nodes obtain the  $NASC_l | j$  value by subtracting the received  $ASC_l | j$  from the  $SC_l$  value locally stored in the  $\mathbf{s}$  vector:

$$NASC_l | j = SC_l - ASC_l | j = s_l - \mathbf{FailSelf}(j)[l]. \quad (2)$$

### D. Stub Release

To exploit SR, each source node must compute silver and bronze RPs by considering the value of  $RWC_l | j$  on

<sup>2</sup>To apply the multiclass lightpath provisioning and restoration strategy, described in Section II-B, it is sufficient to know whether  $FC_l$ ,  $ISC_l$ , and  $WC_l^B$  are equal to zero or not. To avoid too frequent link state advertisements, the aforementioned quantities are, therefore, flooded only when they become null or non-null.

<sup>3</sup>Note that in Table I, the capacities allocated to silver and bronze RPs are classified as  $WC_m^S$  and  $WC_m^B$ , respectively.

each link  $l$ . To implement SR, an additional release vector  $\mathbf{Release}(\mathbf{m})$  of size  $1 \times L$  is stored in each master node per each incoming link  $m$ . The  $l$ th element of  $\mathbf{Release}(\mathbf{m})$  (i.e.,  $\text{Release}(m)[l]$ ) represents the bandwidth that will be released on link  $l$  upon failure of link  $m$ .

Under failure-free network conditions, these vectors are dynamically updated in each network node by signaling messages sent along APs. The update procedure of  $\mathbf{Release}(\mathbf{m})$  vectors is similar to the one utilized for  $\mathbf{FailSelf}(\mathbf{m})$  and  $\mathbf{FailOther}(\mathbf{m})$  vectors.

Upon failure of link  $j$ , the master node inserts in the fault-notification message the  $\mathbf{Release}(j)$  vector which informs the other network nodes about the  $\text{RWC}_{i|j}$  value in each network link  $l$  (i.e.,  $\text{Release}(j)[l] = \text{RWC}_{i|j}$ ). This information is utilized by other network nodes to compute RPs.

#### IV. SIMULATION SCENARIO

The IR-based schemes performance analysis has been conducted by means of a custom network simulator developed in C++. The considered test network is the 15-node 24-bidirectional link Pan-European topology utilized in the study in [15]. Additionally, a 15-node ring topology and a  $5 \times 5$  Manhattan street topology have been used for comparison: The former has the same node count as the Pan-European topology and is sparser, while the latter has a higher node count than the Pan-European one with the same average nodal degree. Each network link carries up to  $W_l = 16$  wavelengths, and each node is equipped with wavelength conversion capability. The traffic model is incremental, i.e., bidirectional lightpath requests are served sequentially, and accepted lightpaths are never torn down. Bidirectional lightpaths requests are randomly generated by following a uniform distribution among all node pairs. A single link failure scenario is considered: Upon reaching a given load, all possible single link failures are generated.

##### A. Performance Metrics

Various performance metrics are considered to assess the performance of the proposed schemes during both the provisioning and the recovery phase. They are detailed as follows.

- 1) The provisioning blocking probability  $P_b^P$  is the ratio between the number of accepted and generated lightpath requests. Gold-class  $P_b^P$  is further divided into two additive components: the blocking probability  $P_b^{P,AP}$ , due to the inability of finding the AP, and the blocking probability  $P_b^{P,BP}$ , due to the inability of finding the BP after a successful AP computation.
- 2) The lightpath failure probability  $P_f$  is the ratio between the number of faulted lightpaths, either disrupted by the failure or preempted, and the number of active lightpaths upon failure (i.e., the lightpaths accepted until failure).
- 3) The restoration blocking probability  $P_b^R$  is the ratio between the number of unrecovered and faulted lightpaths.
- 4) The survivability  $S$ , which is defined in [25] and [26], is the ratio between the number of lightpaths surviving a failure, being either not failed or successfully recovered, and the number of active lightpaths upon failure.

Given the definitions above, the following holds:

$$S = 1 - P_f \cdot P_b^R. \quad (3)$$

Gold-class survivability is unitary under the assumption of single link failure because gold lightpaths are protected. The lower bound for silver- and bronze-class survivability is obtained from (3) when no lightpath can be recovered (i.e.,  $P_b^R = 1$ ), and it is equal to  $1 - P_f$ .

The  $P_b^P$  is evaluated as a function of the per-class generated requests, which are the total number of per-class lightpath requests received by the network during each simulation trial. The other performance metrics (i.e.,  $P_f$ ,  $P_b^R$ , and  $S$ ) are evaluated as a function of the per-class accepted requests, which are the average number of per-class active lightpaths in the network upon link failure. The average is calculated on all simulation trials performed with a specific number of generated requests.

##### B. Network State Synchronization in the Recovery Phase

During the provisioning phase, network state information in all nodes is assumed to be up-to-date and synchronized. However, upon failure occurrence, many restoration attempts for disrupted silver- and bronze-class lightpaths are almost simultaneously triggered. Thus, some nodes might not have an information consistent with the current network state. In this paper, two different network state synchronization scenarios are evaluated during the recovery phase: 1) In the global-updating procedure (GU for short), the network state, which is stored in each node, is assumed to be instantaneously updated after each restoration attempt. Therefore, each source node computes RPs considering the resources utilized by previously restored silver and bronze lightpaths, and 2) in the local-updating procedure (LU for short), each node has an up-to-date knowledge of the resources utilized just by the recovered lightpaths originating from itself. Therefore, each source node computes RPs considering the pre-failure network state information with only local updates.

#### V. SIMULATION RESULTS

Three network scenarios are considered to determine the impact of the utilization of IR-based schemes on network performance: Each of them is reported in a separate section. Each simulation set consists of at least 500 independent trials for each traffic value. The simulation is stopped when either the desired confidence (i.e., a confidence interval at 95% confidence level within 10% of the estimated average) or the maximum number of independent trials (i.e., 5000) is reached. All results are plotted with the confidence interval at 95% confidence level.

In Section V-A, the performance obtained by the use of R-IR versus the one obtained by the use of SR during dynamic restoration is assessed in a scenario with gold- and silver-class lightpaths. In Section V-B, two alternative two-class scenarios are compared, where gold lightpaths share network capacity with either silver lightpaths (exploiting R-IR) or bronze lightpaths (exploiting P-IR). Finally, in Section V-C, a three-class scenario is assessed, with the contemporary presence of




Fig. 2. Silver-class restoration blocking (a) probability and (b) survivability with GU: Pan-European topology.

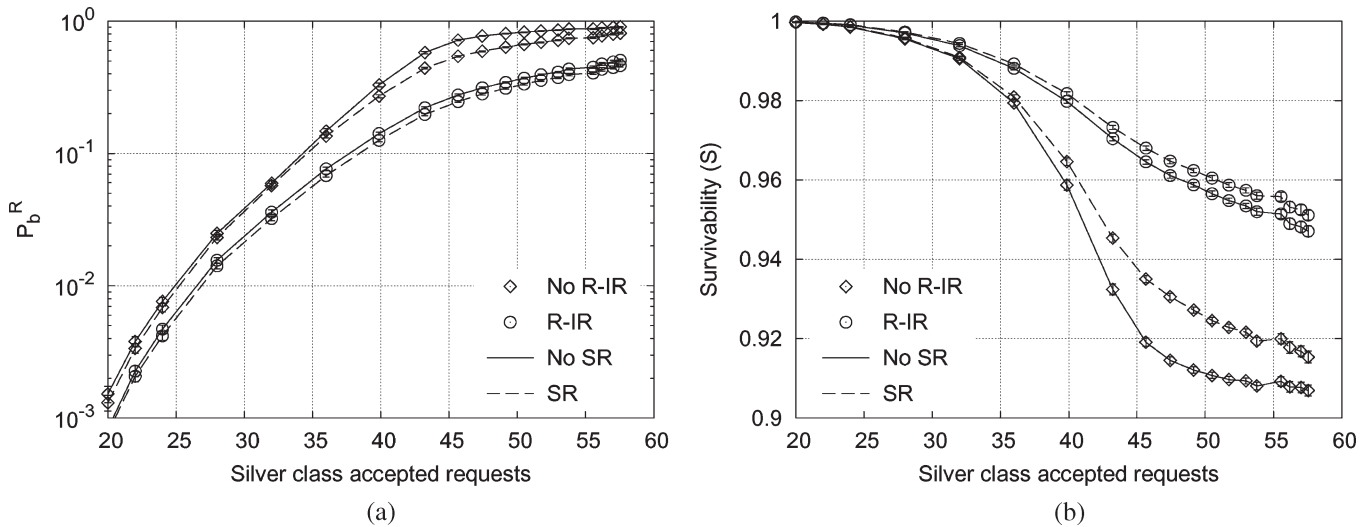


Fig. 3. Silver-class restoration blocking (a) probability and (b) survivability with LU: Pan-European topology.

gold, silver (exploiting R-IR), and bronze (exploiting P-IR) lightpaths.

A. R-IR versus SR

To compare the performance of the use of R-IR and SR during dynamic restoration, a two-class scenario is considered, where an equal amount of gold and silver lightpaths is generated. R-IR and SR can be utilized for silver lightpath dynamic restoration either alternatively or jointly since they exploit different resources. Provisioning blocking probability  $P_b^P$  and failure probability  $P_f$  are omitted because they are not affected by the considered schemes.

Figs. 2 and 3 show that the use of R-IR provides a more significant reduction of  $P_b^R$  and increase of  $S$  than the use of SR, with respect to the case in which only  $FC_l$  is utilized (i.e., no R-IR–no SR). For instance, in Fig. 2(a), at a load of 40 accepted requests, the R-IR scheme achieves a decrease of 86% in  $P_b^R$ , while the utilization of SR obtains just a 30% decrease.

Moreover, the utilization of SR, in conjunction with R-IR, is not beneficial: The  $P_b^R$  decrease is 89%, which is comparable with the decrease obtained by the R-IR scheme alone.

Fig. 2(b) shows that if neither R-IR nor SR is adopted, the survivability value is less than one (i.e., some silver lightpaths begin not to survive the failure), even for low load, and it drops to values around 91% eventually. The introduction of R-IR significantly improves silver lightpath survivability:  $S$  starts dropping from the unitary value at a larger load, and for high network load, it reaches values around 96%. On the other hand, the introduction of SR, either alone or in combination with R-IR, brings only minor improvements. If only SR is utilized,  $S$  starts decreasing for the same traffic as in the scheme without R-IR and SR but with a slightly smoother slope; final values approach 92%. Moreover, if SR is utilized in combination with R-IR, then the survivability improves just marginally than if R-IR is utilized alone.

The comparison between the plots depicted in Figs. 2 and 3 shows that the inaccurate network state information stored




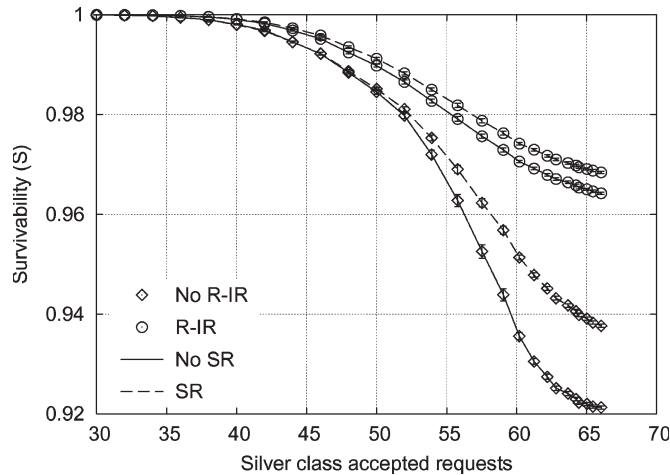


Fig. 4. Silver-class survivability with LU: 15-node ring topology.

Fig. 5. Silver-class survivability with LU:  $5 \times 5$  Manhattan street topology.

with LU compared to GU negatively impacts on the restoration blocking probability and on the survivability. However, all schemes maintain the same performance ranking achieved with GU. In general, with the LU procedure, the improvement achieved by R-IR is still evident, and the marginal contribution of SR, which also requires additional link state information as detailed in Section III-D, is confirmed.

To investigate the impact of the network topology, Figs. 4 and 5 present the survivability results with LU for the 15-node ring and the  $5 \times 5$  Manhattan street topology, respectively. In the ring, the survivability for all the considered schemes decreases because of the lower average nodal degree, but the impact of R-IR is still very relevant, while SR is less useful. In the  $5 \times 5$  Manhattan street topology, the survivability is higher, but the same ranking of the schemes is maintained.

Based on the results shown in this section, the following sections focus only on schemes where silver-class lightpaths exploit the R-IR scheme (i.e., dynamic restoration with R-IR but not SR). Moreover, the LU procedure is considered, which provides a lower bound on the R-IR scheme performance: With the accurate information provided by GU, the improvement obtained by the R-IR scheme is expected to be larger.

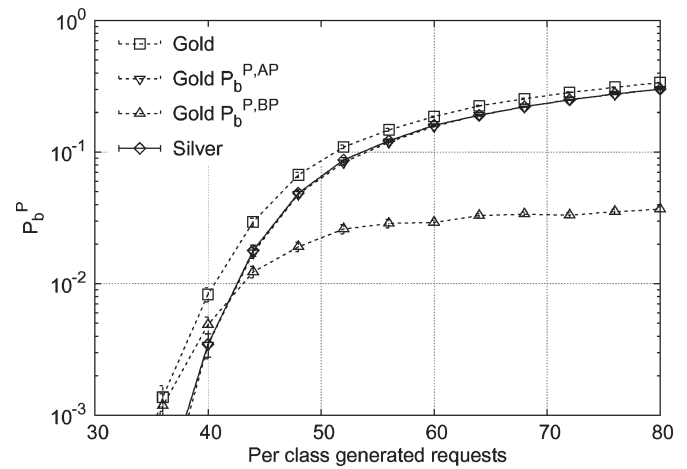


Fig. 6. Provisioning blocking probability in the gold-silver scenario: Pan-European topology.




Fig. 7. Provisioning blocking probability in the gold-bronze scenario: Pan-European topology.

### B. Two-Class Scenario: Gold-Silver versus Gold-Bronze

In this section, two different two-class scenarios are considered, where an equal amount of gold and, respectively, silver and bronze lightpath requests are generated. Thus, in the gold-silver scenario, idle protection resources are used only after a failure (i.e., R-IR), while in the gold-bronze scenario, they are used only before the failure (i.e., P-IR).

Fig. 6 refers to the gold-silver scenario and shows that gold-class  $P_b^P$  is higher than silver-class  $P_b^P$ , given the same amount of generated requests. The additional blocking of gold-class lightpaths is due to the unsuccessful search for a BP.

Fig. 7 refers to the gold-bronze scenario and shows that the  $P_b^P$  of both classes is lower than the one experienced by either class in the gold-silver scenario. This is explained by the fact that the sharing of resources between gold and bronze lightpaths during the provisioning phase facilitates the path search for both classes. Indeed, gold-class  $P_b^{P,AP}$  and bronze  $P_b^P$  are identical, but in this scenario, gold-class  $P_b^{P,BP}$  is null.

The better behavior of the bronze class in the provisioning phase with respect to the silver class is counterbalanced by a worse behavior after a failure occurrence. This fact is

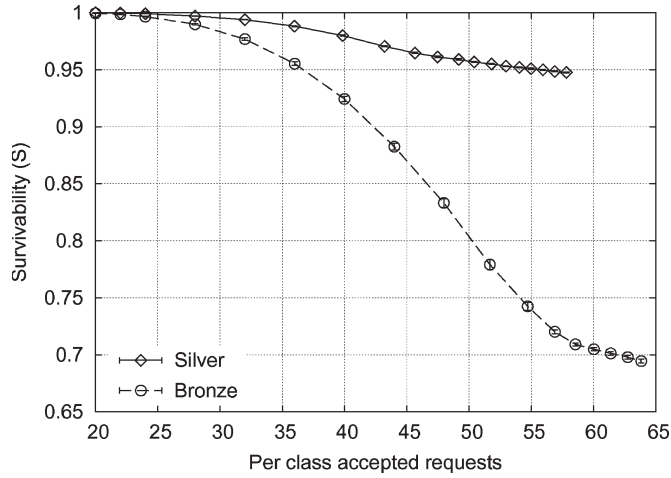


Fig. 8. Survivability in gold-silver and gold-bronze scenarios: Pan-European topology.




Fig. 9. Failure probability and restoration blocking probability in gold-silver and gold-bronze scenarios: Pan-European topology.

highlighted in Fig. 8, showing silver- and bronze-class survivabilities as a function of the per-class accepted requests. While silver  $S$  remains high (around 95%) even for heavy network load, bronze  $S$  quickly decreases, reaching very low values (below 70%).

The worse survivability performance of bronze lightpaths can be explained by looking into its components  $P_f$  and  $P_b^R$ , shown in Fig. 9, as a function of the per-class accepted requests.  $P_b^R$  is similar for the two classes: Bronze  $P_b^R$  is just slightly higher, because bronze lightpaths just exploit free resources during the recovery phase. In contrast, bronze  $P_f$  is about four times the silver  $P_f$ , because in case of failure, many bronze lightpaths undergo preemption.

In conclusion, in the gold-silver scenario, the R-IR scheme guarantees a good survivability at the expense of a larger provisioning blocking probability: At the load ensuring a typical provisioning blocking probability (i.e.,  $P_b^P = 0.01$ ), the survivability is still 97.5%. In contrast, in the gold-bronze scenario, the P-IR scheme improves the provisioning performance but heavily impacts the survivability: At the load ensuring  $P_b^P = 0.01$ , the survivability drops to 77%.

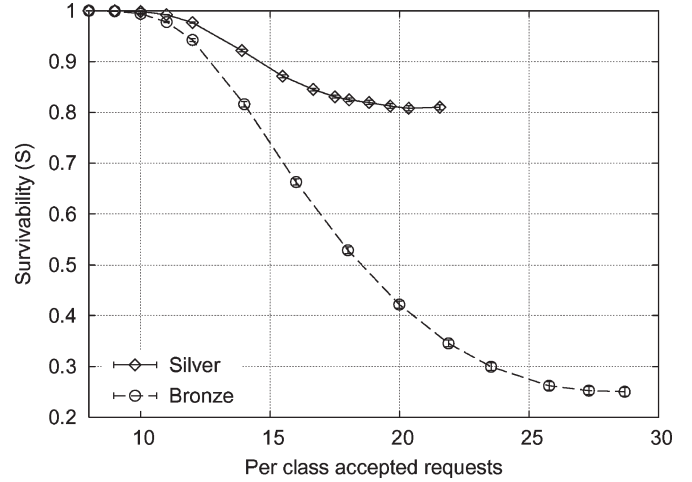


Fig. 10. Survivability in gold-silver and gold-bronze scenarios: 15-node ring topology.

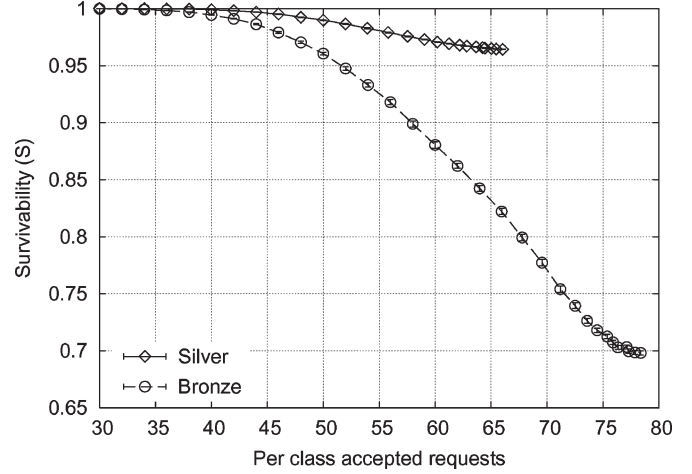


Fig. 11. Survivability in gold-silver and gold-bronze scenarios:  $5 \times 5$  Manhattan street topology.

Figs. 10 and 11 show the survivability for the 15-node ring and the  $5 \times 5$  Manhattan street topology. The results confirm the performance obtained in the Pan-European topology. The particularly low values of bronze  $S$  obtained in the ring for heavy load are due to the low average nodal degree of this topology.

### C. Three-Class Scenario

This section shows the behavior of the proposed IR-based schemes in a three-class scenario, where gold, silver (exploiting the R-IR scheme), and bronze (exploiting the P-IR scheme) lightpaths are concurrently established. Two traffic patterns are reported in two distinct sections: In the first one, a variable equal amount of gold, silver, and bronze lightpath requests is generated; in the second one, the amount of generated gold requests is fixed, while only the amount of silver and bronze requests equally varies.

1) *Variable Gold-Silver-Bronze Traffic Pattern*: Figs. 12–14 refer to the traffic pattern with an equal amount of gold, silver, and bronze generated requests. Fig. 12 shows the provisioning blocking probability  $P_b^P$  as a function each class of



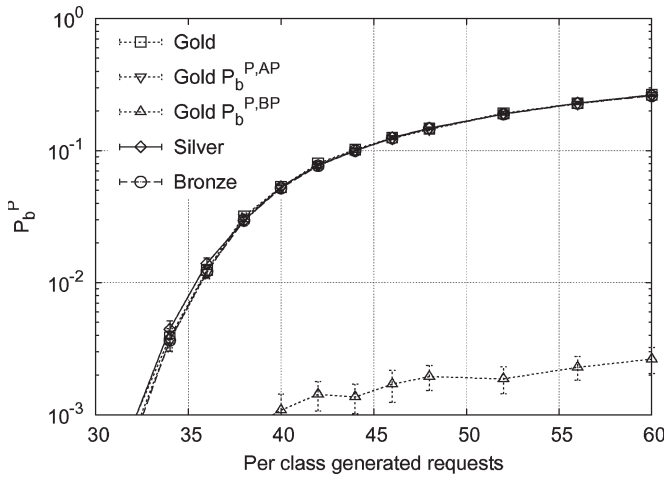


Fig. 12. Provisioning blocking probability in the three-class scenario with variable gold-silver-bronze requests: Pan-European topology.




Fig. 14. Failure probability and restoration blocking probability in the three-class scenario with variable gold-silver-bronze requests: Pan-European topology.




Fig. 13. Survivability in the three-class scenario with variable gold-silver-bronze requests: Pan-European topology.

generated requests. The three classes experience a similar  $P_b^P$ , because they all exploit a similar amount of capacity for routing their APs. Specifically, gold-class lightpath requests are not penalized by the necessity to find both an AP and a BP ( $P_b^{P,BP}$  is negligible as compared to  $P_b^{P,AP}$ ), because gold BPs share resources with bronze APs.

As shown in Fig. 13, the class differentiation is achieved during the recovery phase. Gold  $S$  is unitary under the assumption of single link failure, while silver and bronze  $S$  decrease as the network load increases. However, while silver  $S$  remains high (above 95%), also for heavy load, bronze  $S$  quickly decreases, reaching very low values (below 65%). Even in the scenario, which is shown for comparison, where silver RPs cannot preempt bronze APs routed along idle gold BP resources (i.e., R-IR without bronze preemption), the difference is still significant (silver  $S$  above 90% and bronze  $S$  below 70% for heavy load), demonstrating that the differentiation effectiveness is inherent in the IR. The worse survivability performance of the bronze lightpaths can be explained again by considering its components  $P_f$  and  $P_b^R$ , which are shown in Fig. 14, as a function of the accepted requests. As in Fig. 9, also in this scenario, bronze  $P_f$  is about four times higher than

silver  $P_f$ , because many bronze lightpaths are preempted to recover gold- and silver-class lightpaths. Moreover, bronze  $P_b^R$  is much higher than silver  $P_b^R$  for all network loads. This is due to the fact that bronze RPs are routed just on  $FC_l$ , while silver RPs, by utilizing the R-IR scheme, can exploit also the  $NASC_{l|j}$ , preempting the bronze APs using it. In particular, for high network load, bronze  $P_b^R$ , which is different from silver  $P_b^R$ , approaches one. Thus, bronze  $S$  tends to the lower bound  $1 - P_f$ , corresponding to the probability of not being disrupted.

2) *Fixed Gold Traffic Pattern:* This section analyzes the dependence of the performance experienced by different lightpath classes on the amount of gold-class lightpaths. Three sets of simulations are run with a fixed amount of generated gold lightpath requests (respectively 20, 40, and 60 requests) and a variable amount of silver and bronze lightpath requests (from 0 to 60 for both classes in all simulations).

Fig. 15 shows that for each amount of generated gold requests, silver  $S$  is higher than bronze  $S$ . The values of  $S$  for both classes decrease as the amount of generated gold requests increases because the overall network load increases. However, on the one hand, silver  $S$  is monotonically decreasing as a function of the accepted silver requests and bounded to values higher than 95%. On the other hand, bronze  $S$  shows a non-monotonic behavior as a function of the accepted bronze requests (reaching values lower than 55%), and it heavily depends on the amount of generated gold-class requests.

The survivability behavior can be explained by the plots in Figs. 16 and 17. Fig. 16 shows that for the same amount of gold requests, bronze  $P_b^R$  is higher than silver  $P_b^R$  for any value of the silver and bronze accepted requests. Moreover, the relative difference between silver and bronze  $P_b^R$  is higher for increasing amount of gold-class lightpaths: If more gold lightpaths are active in the network, the  $NASC_{l|j}$  upon failure occurrence is greater, and the R-IR scheme, which is utilized by silver lightpaths, is more efficient. Again, for heavy load, bronze  $P_b^R$  approaches one, and bronze  $S$  reduces to  $1 - P_f$ . In Fig. 17, it is possible to observe that silver  $P_f$  is fixed at 10% and that it does not vary in function of both silver and bronze accepted requests and gold requests because silver lightpaths

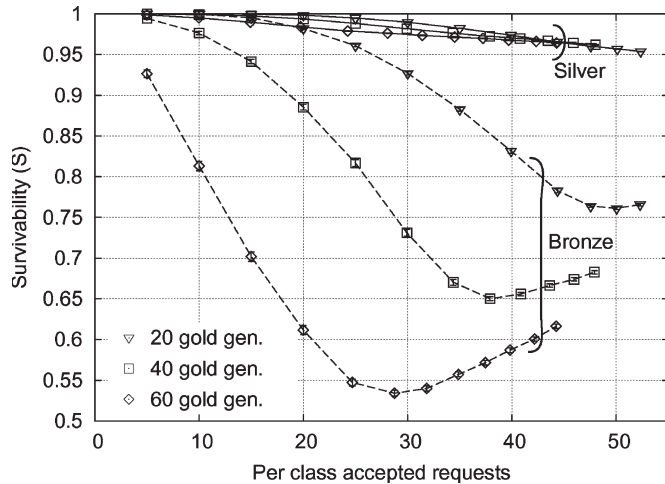


Fig. 15. Survivability in the three-class scenario with fixed gold requests: Pan-European topology.

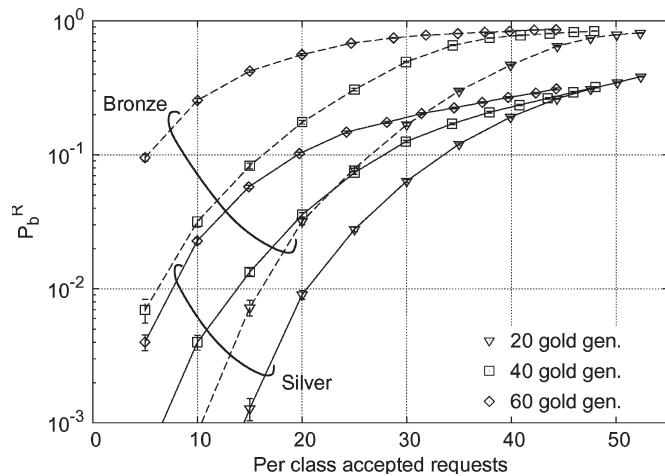


Fig. 16. Restoration blocking probability in the three-class scenario with fixed gold requests: Pan-European topology.




Fig. 17. Failure probability in the three-class scenario with fixed gold requests: Pan-European topology.

cannot be preempted. Thus, silver  $P_f$  is just the probability that a lightpath link fails, and it depends only on the topology, the traffic pattern, and the routing algorithm. On the other hand,

bronze  $P_f$  also accounts for preempted lightpaths, and thus, it greatly depends on the amount of gold, silver (i.e., preempting), and bronze (i.e., preemptable) requests accepted in the network. In particular, for a fixed amount of silver and bronze accepted lightpaths, bronze  $P_f$  increases with the number of gold generated requests, since in case of failure more bronze lightpaths are preempted by gold BPs. For a fixed amount of gold generated requests, bronze  $P_f$  decreases for increasing bronze accepted requests, since the fraction of bronze requests exploiting gold SC<sub>l</sub> decreases. Thus, in the three-class scenario, bronze  $S$ , which is different from silver  $S$ , may significantly vary as a function of all classes requests.

## VI. CONCLUSION

This paper focused on a multiclass optical network scenario, where high-class lightpaths exploit SPP and low-class lightpaths resort to best-effort dynamic restoration. Two idle protection capacity reuse (IR)-based schemes for increasing the network resource utilization while guaranteeing lightpath class survivability differentiation have been assessed: the existing P-IR and the novel R-IR. P-IR and R-IR schemes utilize idle high-class lightpath protection resources for the provisioning and the restoration of low-class lightpaths, respectively. R-IR has also been compared with a commonly utilized postfailure reuse scheme, i.e., SR. A distributed implementation of the proposed schemes has been detailed, where the R-IR scheme does not require additional network state information with respect to the information required by SPP. On the other hand, the P-IR scheme and SR require that each network node stores and periodically updates additional network state vectors.

Numerical results have shown that, if R-IR is utilized, low-class lightpath survivability is increased. Moreover, R-IR has proven to be significantly more beneficial than SR. If low-class lightpaths alternatively exploit P-IR and R-IR, a tradeoff between provisioning and restoration performance is achieved. The provisioning performance of high- and low-class lightpaths is improved with P-IR compared to R-IR, but upon failure occurrence, the low-class lightpath survivability is higher with R-IR than with P-IR. Furthermore, if two low-class lightpaths are concurrently utilized (i.e., silver exploiting R-IR and bronze exploiting P-IR), a performance differentiation in the recovery phase is inherently achieved: Silver survivability is bounded and is always higher than bronze survivability. In addition, bronze survivability heavily depends on the amount of spare capacity allocated for high-class lightpaths.

## REFERENCES

- [1] M. Tacca *et al.*, "Differentiated reliability in optical networks: Theoretical and practical results," *J. Lightw. Technol.*, vol. 21, no. 11, pp. 2576–2586, Nov. 2003.
- [2] M. Durvy *et al.*, "Network availability based service differentiation," in *Proc. IWQoS*, Jun. 2003, pp. 305–324.
- [3] S. Dong *et al.*, "Differentiated-resilience provisioning for the wavelength-routed optical network," *J. Lightw. Technol.*, vol. 24, no. 2, pp. 667–673, Feb. 2006.
- [4] N. Andriolli *et al.*, "Separate wavelength pools for multiple-class optical channel provisioning," in *Proc. NETWORKS*, Jun. 2004, pp. 379–384.
- [5] L. Berger, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*, Jan. 2003. IETF RFC 3471.

- [6] J. Lang *et al.*, *Generalized Multi-Protocol Label Switching (GMPLS) Recovery Functional Specification*, Mar. 2006. IETF RFC 4426.
- [7] E. Mannie and D. Papadimitriou, *Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)*, Mar. 2006. IETF RFC 4427.
- [8] D. Papadimitriou and E. Mannie, *Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-Based Recovery Mechanisms (Including Protection and Restoration)*, Mar. 2006. IETF RFC 4428.
- [9] P. Demeester *et al.*, "Resilience in multilayer networks," *IEEE Commun. Mag.*, vol. 37, no. 8, pp. 70–76, Aug. 1998.
- [10] M. Pickavet *et al.*, "Recovery in multilayer optical networks," *J. Lightw. Technol.*, vol. 24, no. 1, pp. 122–134, Jan. 2006.
- [11] S. Ramamurthy *et al.*, "Survivable WDM mesh networks," *J. Lightw. Technol.*, vol. 21, no. 4, pp. 870–883, Apr. 2003.
- [12] G. Li *et al.*, "Efficient distributed restoration path selection for shared mesh restoration," *IEEE/ACM Trans. Netw.*, vol. 11, no. 5, pp. 761–771, Oct. 2003.
- [13] Y. Liu *et al.*, "Approximating optimal spare capacity allocation by successive survivable routing," *IEEE/ACM Trans. Netw.*, vol. 13, no. 1, pp. 198–211, Feb. 2005.
- [14] P.-H. Ho and H. Moutfah, "Shared protection in mesh WDM networks," *IEEE Commun. Mag.*, vol. 53, no. 1, pp. 70–76, Jan. 2004.
- [15] A. Giorgetti *et al.*, "Failure-aware idle protection capacity reuse," in *Proc. GLOBECOM*, St. Louis, MO, Dec. 2005, vol. 4, pp. 1895–1899.
- [16] A. Giorgetti *et al.*, "Impact of idle protection capacity reuse on multi-class optical networks," in *Proc. DRCN*, Ischia, Italy, Oct. 2005, pp. 149–154.
- [17] N. Andriolli *et al.*, "Impact of spare capacity reuse on multi-class light-path survivability differentiation," in *Proc. ICTON*, Barcelona, Spain, Jul. 2005, vol. 2, pp. 9–12.
- [18] R. R. Iraschko *et al.*, "Optimal capacity placement for path restoration in STM or ATM mesh-survivable networks," *IEEE/ACM Trans. Netw.*, vol. 6, no. 3, pp. 325–336, Jun. 1998.
- [19] D. Schupke and R. Prinz, "Capacity efficiency and restorability of path protection and rerouting in WDM networks subject to dual failures," *Photonic Netw. Commun.*, vol. 8, no. 2, pp. 191–207, Sep. 2004.
- [20] A. Banerjee *et al.*, "Generalized multi-protocol label switching: An overview of signaling enhancements and recovery techniques," *IEEE Commun. Mag.*, vol. 39, no. 7, pp. 144–151, Jul. 2001.
- [21] L. Berger, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE) Extensions*, Jan. 2003. IETF RFC 3473.
- [22] A. Banerjee *et al.*, "Generalized multi-protocol label switching: An overview of routing and management enhancements," *IEEE Commun. Mag.*, vol. 39, no. 1, pp. 144–150, Jan. 2001.
- [23] K. Kompella and Y. Rekhter, *Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*, Oct. 2005. IETF RFC 4202.
- [24] K. Kompella and Y. Rekhter, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*, Oct. 2005. IETF RFC 4203.
- [25] S. C. Liew and K. W. Lu, "A framework for characterizing disaster-based network survivability," *J. Lightw. Technol.*, vol. 12, no. 1, pp. 52–58, Jan. 1994.
- [26] T. Z. Jiang, "A new definition of survivability of communication networks," in *Proc. MILCOM*, Nov. 1991, pp. 901–904.



**Nicola Andriolli** received the Laurea degree in telecommunications engineering from the University of Pisa, 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 Student at Budapest University of Technology and Economics (BUTE), Budapest, Hungary, and at Department of Communications, Optics and Materials of the Technical University of Denmark (COM-DTU), Copenhagen, Denmark, and he was a Visiting Trainee at National Institute of

Information and Communications Technology (NICT), Tokyo, Japan. He is currently a Postdoctoral Fellow with Scuola Superiore Sant'Anna. His research interests include network modeling and simulation, QoS and fault tolerance in optical networks, Generalized Multi-Protocol Label Switching (GMPLS) architecture, and optical packet switching.



**Alessio Giorgetti** received the Laurea degree in telecommunications engineering from the University of Pisa, Pisa, Italy, in 2002 and the Ph.D. degree from Scuola Superiore Sant'Anna, Pisa, in 2006.

He is currently a Postdoctoral Fellow with Scuola Superiore Sant'Anna. His research interests include wavelength-division-multiplexing (WDM) network design, fault tolerance in optical networks, GMPLS control plane for wavelength-routed optical networks, next-generation Synchronous Digital Hierarchy (SDH) technologies, and WDM-passive

optical network (PON) access networks design.



**Luca Valcarengi** (S'00–M'02) received the Laurea degree in electronics engineering from the Politecnico di Torino, Torino, Italy, in 1997, and the M.S. degree in electrical engineering and the Ph.D. degree in electrical engineering-telecommunications from the University of Texas at Dallas, Richardson, TX, in 1999 and 2001, respectively.

Between January 2002 and August 2002, he was a Research Associate with the Optical Networking Advanced Research Laboratory, Erik Jonsson School of Engineering and Computer Science, University of Texas at Dallas. Since September 2002, he has been an Assistant Professor with the Scuola Superiore Sant'Anna, Pisa, Italy. He coauthored more than 50 papers published in international journals and presented in leading international conferences. His main research interests are optical network design, analysis, and optimization, artificial intelligence optimization techniques, communication network reliability, IP over wavelength division multiplexing (WDM) networking, QoS in network infrastructures for Grid computing, and broadband wireless access (IEEE802.16).

Dr. Valcarengi has been part of the Organizing and Technical Program Committees of International conferences.



**Piero Castoldi** (S'93–A'96) received the Laurea degree (*cum laude*) in electrical engineering from the University of Bologna, Bologna, Italy, in 1991 and the Ph.D. degree from the University of Parma, Parma, Italy, in 1996.

In 1996 and 1997 and in the summers of 1999 and 2000, he was a Postdoc at the Department of Electrical Engineering, Princeton University, Princeton, NJ. From December 1997 to February 2001, he was an Assistant Professor at the Faculty of Engineering, University of Parma, where he became tenured.

Since March 2001, he has been an Associate Professor in telecommunications with the Scuola Superiore Sant'Anna, Pisa, Italy, where he is responsible for the networks group. In January and February 2002, he was a Visiting Professor at the University of Texas at Dallas, Richardson, TX. He also collaborates in several projects of the Inter-University National Consortium for Telecommunications (CNIT), and since March 2003, he has been the Director of the CNIT National Photonic Networks Laboratory in Pisa. He has held, as an Instructor, several courses at the undergraduate and graduate level at the University of Parma and at Scuola Superiore Sant'Anna. His scientific activity has covered the area of digital transmission, telecommunication-network performance, and, more recently, provisioning, protection, and switching techniques for the next-generation optical networks. He is the author of more than 60 technical papers published in international conferences and journals and of an international book on code division multiple access.