Experimental Demonstration of a GMPLS-enabled Impairment-Aware Lightpath Restoration Scheme

Jordi Perelló, Salvatore Spadaro, Fernando Agraz, Marianna Angelou, Siamak Azodolmolky, Yixuan Qin, Reza Nejabati, Dimitra Simeonidou, Pannagiotis Kokkinos, Emmanouel (Manos) Varvarigos, and Ioannis Tomkos

***Abstract*—Fast lightpath restoration becomes particularly challenging in all-optical networks. First, the optical transparency complicates failure localization and isolation procedures, as loss of light alarms stemming from a failure propagate downstream from the failure point. Besides, such a transparency implies that optical signals must traverse relatively long distances without electrical regeneration. In view of this, backup path computations must also account for all degradations introduced in the physical end-to-end path, thus ensuring their feasibility. Looking toward the optical core networks of the future, the EU DICONET Project has worked on cross-layer solutions to enhance network control and management with the impairment-awareness needed to govern the underlying optical layer. This includes a network planning and operation tool (NPOT) that implements impairment-aware routing and wavelength assignment algorithms, along with a failure localization mechanism. This paper reports the experimental demonstration of a dynamic impairment-aware restoration scheme that benefits from enhanced NPOT features for fast lightpath restoration. To this end, a prioritized scheduler to provide differentiated resilience support, the implementation of the quality of transmission estimator module on field programmable gate array hardware, and a fast resource pre-reservation protocol are presented in this work. The performance of the proposed impairment-aware lightpath restoration scheme has been evaluated experimentally on a 14-node all-optical network test-bed, showing average restoration times of 1.16 and 1.64 s for high and low priority traffic classes, respectively.**

***Index Terms*—Failure restoration; Impairment-awareness; Optical networks; Transparency.**

Manuscript received June 29, 2011; revised December 21, 2011; accepted March 19, 2012; published April 10, 2012 (Doc. ID 150020).

Jordi Perelló (e-mail: perello@ac.upc.edu), Salvatore Spadaro, and Fernando Agraz are with the Universitat Politècnica de Catalunya (UPC), Jordi Girona 1–3, 08034 Barcelona, Spain.

Marianna Angelou is with the Universitat Politècnica de Catalunya (UPC), Jordi Girona 1–3, 08034 Barcelona, Spain, and is also with Athens Information Technology (AIT), 19.5 km Markopoulo Ave., Peania 19002, Athens, Greece.

Siamak Azodolmolky is with the Universitat Politècnica de Catalunya (UPC), Jordi Girona 1–3, 08034 Barcelona, Spain, was with Athens Information Technology (AIT), 19.5 km Markopoulo Ave., Peania 19002, Athens, Greece, and is now with the University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK.

Yixuan Qin, Reza Nejabati, and Dimitra Simeonidou are with the University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK.

Pannagiotis Kokkinos and Emmanouel (Manos) Varvarigos are with the Research Academic Computer Technology Institute (RACTI), University of Patras, Patras, Greece.

Ioannis Tomkos is with Athens Information Technology (AIT), 19.5 km Markopoulo Ave., Peania 19002, Athens, Greece.

Digital Object Identifier 10.1364/JOCN.4.000344

I. Introduction

Future all-optical core networks are envisioned to run away from the opaqueness of today's transport networks. Currently, networks are built on point-to-point wavelength division multiplexing (WDM) channels, which are electrically terminated and regenerated at every intermediate node of the end-to-end path. In contrast, all-optical networks will leverage data transmission and switching operations entirely in the optical domain. Therefore, operators will be relieved from deploying expensive and not scalable service-dependent intermediate electronic stages, enabling ultrafast, highly scalable, and cost-effective backbone networks [1,2].

Transparent optical networks will also be enhanced with automated connection (i.e., lightpath) provisioning and restoration functionalities, easing the adoption of reliable on-demand transport services at low cost [3]. Cost savings will come here from the avoidance of manual resource provisioning operations plus the high network capacity efficiency provided by the restoration schemes [4]. In traditional protection mechanisms, extra network capacity is allocated during the lightpath setup, so as to ensure fast traffic switchover under failure conditions. Conversely, restoration makes use of additional network capacity reactively, only once a lightpath failure is detected. This enables extensive sharing of the spare capacity amongst non-simultaneously failed lightpaths. To achieve these goals, the generalized multi-protocol label switching (GMPLS) protocols [5], resulting from the Internet Engineering Task Force (IETF) efforts toward a standard distributed control plane, become the prime candidates to be adopted.

However, the transition from opaque to transparent optical networks poses some challenges that should be properly addressed. First of all, transparency implies the transmission of signals over very long physical distances without electrical regeneration. This may lead to unfeasible lightpaths due to the physical layer degradations that accumulate on the signal along the path. To address this issue, regenerators (i.e., optical–electronic–optical, OEO) can be strategically placed at a selected number of sites. This increases the total distance that can be spanned by the lightpaths, thus making them feasible. Networks following this configuration are typically referred to as translucent networks (e.g., see [6]).

In either fully transparent or translucent optical networks, physical layer impairment (PLI) information should be taken into account in the routing and wavelength assignment (RWA) algorithms, turning them into impairment-aware RWA (IA-RWA) algorithms. In this way, the physical feasibility of the lightpaths under establishment can be predicted and regeneration, if available, can be utilized. The Q-factor of a lightpath (i.e., a measure of its quality of transmission, QoT) is a parameter typically used for these purposes [7]. In [8], the IA-RWA process is performed in two phases. First a lightpath computation step takes place where the shortest path is selected and, then, a lightpath verification step based on a bit error rate (BER) threshold is performed. In [9], the authors propose an IA-RWA algorithm in which the wavelength is initially selected by means of a first-fit algorithm (unaware of PLI constraints) and, then, a shortest path for that wavelength is computed, considering the noise variance of the PLI as the link cost. Apart from separated RWA solutions, some heuristics intend to solve the IA-RWA problem jointly [10,11]. Besides, a number of IA-RWA algorithms for translucent networks, considering both linear and non-linear physical effects, are evaluated in [12].

From the network control plane perspective, the IETF is also working on the definition of the proper extensions to the GMPLS protocols to account for PLIs in optical networks [13]. Following the same trend, [14] introduces extensions to the GMPLS signaling protocol, which include the QoT parameters that characterize the optical layer. Using this extended protocol, distributed IA-RWA schemes are proposed. Moreover, [15] discusses GMPLS routing and signaling protocol extensions for translucent optical networks. However, such extensions do not consider the inter-channel effects. For the interested reader, a survey on physical layer IA-RWA algorithms is presented in [16].

Another challenge inherent in transparent networks is that, without electrical termination at network nodes, loss of light (LoL) alarms stemming from a single failure propagate downstream from the failure point. This issue complicates failure localization and isolation procedures, which are of paramount importance when targeting fast lightpath restoration.

A straightforward solution to address failure localization and isolation problems in all-optical networks is the deployment of dedicated out-of-band optical supervisory channels (OSCs) per fiber link. This solution, however, requires a WDM channel and a transmitter–receiver pair per link just for monitoring purposes. In order to improve the performance of this basic scheme, more sophisticated out-of-band techniques have also been proposed, aiming to minimize network monitoring costs while detecting failures without ambiguity (e.g., see [17–19]). In addition, in-band monitoring schemes have also been presented [20,21], which allow finer soft-failure detection and localization per WDM channel.

In particular, the concept of monitoring trails (m-trails), has been proposed as the most general monitoring structure for these out-of-band monitoring techniques, achieving detection and localization of hard failures that affect any single fiber link. In [18] the authors formulate an integer linear program (ILP) for optimal m-trail design, with the objective of minimizing the overall monitoring cost in the network. Instead, in [19] an algorithm based on random code assignment (RCA) and random code swapping (RCS) is developed for solving the m-trail design problem.

In this context, the DICONET project [22] has addressed both challenges, looking toward the future Internet backbone. The main outcome of the project has been the development of a network planning and operation tool (NPOT) that collects real-time optical layer performance measurements and incorporates them into IA-RWA algorithms. As a result, lightpath routes are computed, so that potential blockings due to insufficient QoT are avoided. In this process, the effects of new lightpath establishments on currently active lightpaths in the network are also evaluated, avoiding degradations that could even compromise their operation.

Moreover, the designed NPOT also implements scalable failure localization and isolation procedures, addressing the aforementioned all-optical core networks' needs. The proposed solution is based on m-trails, which minimize the number of out-of-band OSCs that need to be deployed in the network in order to achieve unambiguous failure localization. Thus, once the failure is localized and isolated, restoration actions can be triggered for all the affected lightpaths.

However, the complexity of the IA-RWA is generally much higher than in the case of traditional RWA, where only topological and wavelength availability information is considered. This complexity impacts drastically on the final setup times experienced by incoming lightpath requests, which may be even unsuitable for highly dynamic network services such as lightpath restoration. In fact, the most time-consuming steps in the IA-RWA process are the QoT estimations of the lightpath under establishment and the potentially affected ones. With this in mind, the QoT estimator module in the NPOT has been implemented on field programmable gate array (FPGA) hardware. As will be shown in the experimental results, this enables impairment-aware lightpath restoration times of around a second, totally applicable in real network environments.

The remainder of this article continues as follows. Section II describes the functionality of the proposed NPOT, paying special attention to those modules involved in the dynamic impairment-aware restoration. Once the basic NPOT features are explained, Section III presents the enhancements to the NPOT required to match the fast and differentiated restoration needs in future optical networks. Section IV details the 14-node all-optical network test-bed where the experimental evaluation of our proposals has been conducted. Section V presents and discusses the obtained experimental results. Finally, Section VI concludes the paper.

II. Network Planning and Operation Tool

The proposed NPOT will become the enabler of the aforementioned impairment-awareness in both network planning and operation phases. In particular, the network planning phase is devoted to allocating a large, already known set of demands (i.e., the traffic matrix) in the most efficient way. In the context of transparent optical networks, this entails optimal IA-RWA to maximize network resource usage, optimal regenerator placement to guarantee acceptable QoT for all lightpaths, while minimizing the number of regenerators needed, as well as optimal monitor placement to localize failures unambiguously but with the minimum number of monitors.

In contrast, in the network operation phase, lightpath requests are served on a one-by-one basis upon arrival. The goal here is to allocate incoming requests as efficiently as possible considering the current network state. In this regard, two alternative approaches have been proposed for integrating the NPOT in the GMPLS-enabled control plane, namely, distributed and centralized [23].

In the distributed approach, each network node runs an instance of the NPOT that interacts with the impairment-aware GMPLS control plane to assess the QoT of the lightpaths under establishment. For this, end-to-end routes are computed only accounting for topological and wavelength availability information. In fact, real-time PLI values are collected during the Resource Reservation Protocol with Traffic Engineering Extensions (RSVP-TE) signaling process [24] to allow the NPOT at the destination node to estimate the QoT of the new lightpath. A new lightpath may degrade the QoT of those active lightpaths sharing any of the spans along the signaled route. In light of this, the NPOT also contacts with the NPOTs located at the end nodes of the potentially disrupted lightpaths, asking them to recompute their QoT considering the new network scenario. This verification step is performed to ensure that QoT requirements are still met despite the new lightpath establishment. Thus, if no QoT violation occurs, the requested lightpath establishment is finally completed.

Conversely, in the centralized approach, a single NPOT instance is responsible for providing the impairment-awareness required by the GMPLS control plane. To this goal, such a centralized NPOT is fed with real-time PLI values, as well as the wavelength availability, through an extended version of the GMPLS Open Shortest Path First with Traffic Engineering Extensions (OSPF-TE) routing protocol [25]. This information is incorporated into an online IA-RWA engine that computes QoT compliant routes for the new lightpath requests.

A particular situation that has to be managed in the network operation phase is the restoration of those lightpaths whose connectivity has been interrupted due to network failures. The remainder of this section concentrates on the NPOT modules and functionalities involved in the lightpath restoration process. For a complete description of the NPOT functionalities for both network planning and operation phases, the interested reader may refer to [23].

A. Online IA-RWA Module

The online IA-RWA module in the NPOT (Fig. 1) is used in the network operation phase to compute the most appropriate routes and wavelengths for the incoming lightpath requests. In the centralized approach, such a functionality is implemented by the multi-parametric IA-RWA algorithm presented in [26]. This algorithm deals with the effects of the different PLIs indirectly, instead of estimating their absolute values to select the optimal route and wavelength for a new lightpath request, which is time-consuming. This is achieved by considering impairment generating sources, such as the path length, number of hops, or number of crosstalk sources (among others). The algorithm computes a set of candidate lightpaths, whose QoT is approximated by means of a function that combines the impairment generating parameters. These candidate lightpaths are accepted or rejected depending on whether their QoTs remain above or below a predefined threshold, respectively. The feasible lightpaths resulting from

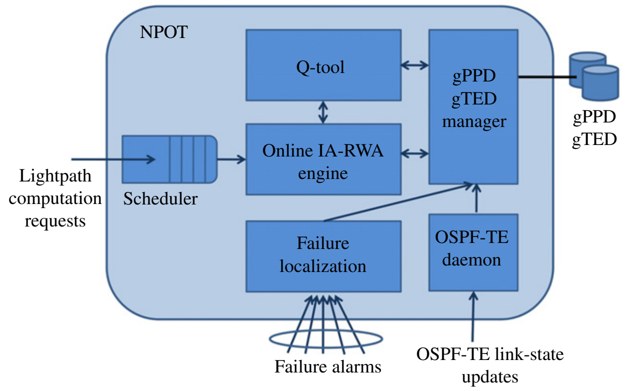

Fig. 1. (Color online) Network planning and operation tool (NPOT): main blocks involved in impairment-aware dynamic lightpath restoration.

the algorithm, as well as the active ones potentially disrupted by the new lightpath establishment, are finally checked using the NPOT Q-tool module (detailed in the following subsection), ensuring in this way their physical feasibility.

B. QoT Estimator Module: Q-tool

Under request from the online IA-RWA module, the Q-tool inside the NPOT (Fig. 1) is responsible for estimating the QoT of a lightpath given the network topology, physical layer characteristics, and currently active lightpaths.

The NPOT maintains two different databases, namely, the global physical parameters database (gPPD) and the global traffic engineering database (gTED), which store all the information requested by the Q-tool module to predict whether or not a specific lightpath will be feasible. Specifically, the gPPD contains the physical characteristics of the nodes, links, and components in the network, together with information about the currently active lightpaths. In turn, the gTED stores topology and wavelength availability information. Such databases are initially populated through configuration files and updated dynamically through the GMPLS control plane by means of the OSPF-TE protocol. In Fig. 1, the gPPD/gTED manager is the module responsible for managing both databases, as well as interfacing them with the other NPOT modules.

As detailed in [23], the Q-tool takes both linear and non-linear PLIs into account in order to estimate the QoT of a certain lightpath. Regarding the linear impairments, the amplifier spontaneous emission (ASE) noise, chromatic dispersion (CD), filter concatenation (FC), and polarization mode dispersion (PMD) are considered. For those non-linear impairments that depend on the characteristics of the active lightpaths at a certain moment, the self-phase modulation (SPM), cross phase modulation (XPM), and four wave mixing (FWM) are also evaluated.

C. Failure Localization and Isolation

As mentioned before, failure localization and isolation functionalities are challenging in all-optical networks. The failure localization module in the NPOT (Fig. 1) addresses

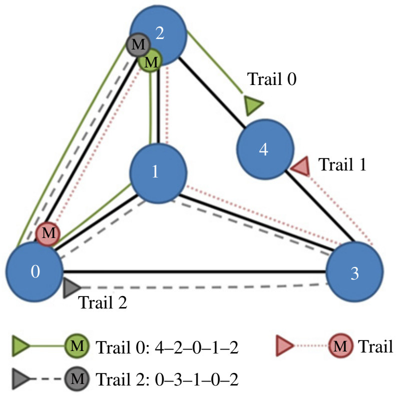

Link	T2	T1	T0
(0, 1)	1	0	1
(0, 2)	1	1	1
(0, 3)	1	0	0
(1, 2)	0	1	1
(1, 3)	1	1	0
(2, 4)	0	0	1
(3, 4)	0	1	0

Fig. 2. (Color online) Example of an m-trail design in a five-node network consisting of three trails. The resulting alarm code table is also shown.

both objectives at low network capital expenditure (CAPEX). Basically, this module localizes a network failure by mapping the failure alarms coming from the optical node controllers to the m-trail monitoring solution designed in the network planning phase. For the sake of clarity, an m-trail could be seen as an OSC spanning several network links, with a signal transmitter located at the m-trail source node, and a monitor at the m-trail destination node with the receiver. By definition, an m-trail can traverse a certain network node many times, but a network link only once. Therefore, when a link failure occurs, all those m-trails supported on the failed link detect LoL in the receiver, generating a failure alarm with the identifier of the disrupted m-trail. The optimal m-trail design must ensure unambiguous failure localization, but using the minimum number of m-trails and, thus, expensive monitors.

Figure 2 depicts an example of an m-trail design in a five-node network. In this case, three m-trails are enough to detect and localize any link failure in the network without ambiguity. This is reflected in the table shown in the same figure, where each network link presents a unique alarm code. It is worth mentioning here that, if an OSC had been deployed per link (i.e., straightforward solution), seven monitors instead of three would have been required for failure localization purposes, which increases the final network CAPEX.

The optimal m-trail design in transparent optical networks is typically obtained by means of ILP formulations, which try to minimize the number of m-trails and their length, subject to a unique alarm code per link (e.g., see [18]). Nonetheless, this problem may take lots of time to solve, especially in large network scenarios. With this in mind, the monitor placement module in the centralized NPOT runs a meta-heuristic for the m-trail design based on tabu search (TS), which reaches an almost optimal design in very low execution times. The details of this meta-heuristic, called meta-heuristic for monitoring trail assignment (MeMoTA), can be found in [27].

The MeMoTA meta-heuristic is executed in the network planning phase to obtain the optimal configuration of m-trails enabling successful failure detection and localization. Then, based on its output, an alarm code table such as the one in Fig. 2 is constructed. This alarm code table is later used by the NPOT failure localization module to localize failures based on the failure alarms received from the GMPLS control plane nodes. Note that the NPOT requires information about all disrupted m-trails to successfully localize a given failure. To this end, a hold-off timer is set once a failure alarm is detected. Once the timer expires, localization is performed on the alarm code table, using the failure alarms received so far. Having the exact location of the failure, all the affected resources are put into failure state, isolating the failure from future lightpath requests. Finally, restoration actions are triggered for all the affected lightpaths.

III. NPOT Enhancements for Fast Lightpath Restoration

As experimentally demonstrated in [23], the centralized approach performs significantly better in terms of lightpath blocking probability. However, it shows increased lightpath setup times due to the high computational complexity of the IA-RWA process at the centralized NPOT. While such setup times, of the order of a few seconds, may be acceptable for a new lightpath establishment, they would lead to large data losses in the case that failure restoration was implemented in the network. In view of this, the present work aims at enhancing the centralized NPOT, so as to successfully fulfill the failure restoration needs in future core optical networks, while still benefitting from its accurate impairment-aware path computation feature. The improvements introduced in this work comprise a prioritized NPOT scheduler to efficiently support traffic with differentiated resilience requirements, a hardware-accelerated Q-tool implementation, which drastically reduces the QoT estimation time, that is, the bottleneck in the IA-RWA process, as well as a resource pre-reservation protocol allowing consecutive backup path computations.

A. Prioritized NPOT Scheduler

In the network operation phase, the NPOT scheduler (Fig. 1) queues the incoming requests until the online IA-RWA module becomes available, since the NPOT is only able to process one path computation at a time. This operation becomes of critical importance during the failure restoration procedures, where a large number of restoration requests must be served almost simultaneously. Besides, as pointed out before, the IA-RWA process is a computationally intensive task. Hence, the online IA-RWA engine requires a significantly large amount of time to process each restoration request (even of the order of a few hundreds of milliseconds with the hardware acceleration features presented in this paper). This means that the latest restoration requests to arrive at the centralized NPOT experience high queuing delays, which could even compromise the success of their restoration.

Taking into account the large variety of data applications that will run over next-generation optical networks, each one requiring differentiated treatment in case of failure (e.g., a maximum recovery time), the NPOT has been equipped with a prioritized scheduler aiming to serve the restoration requests with higher resilience requirements (i.e., higher priority) first, incurring lower recovery times for them. The prioritized scheduler, governed by the NPOT scheduler manager, implements n software-based first-in-first-out (FIFO) queues, one for each lightpath priority supported on the network. This solution becomes interesting provided that a reasonable number of priorities are defined in the network, since no sorting algorithm is needed in the prioritized scheduler. In particular, two different priorities are considered in this work.

The operation of the scheduler is the following. At the moment that the NPOT scheduler manager detects that any queue of the prioritized scheduler has path requests to be served, it checks the availability of the online IA-RWA engine. If this is free, the manager starts checking the states of the priority queues, from the highest to the lowest priority (e.g., from 1 to n). If it finds any restoration request in any of the queues, the manager pops the request that arrived at the queue first, sends the request to the online IA-RWA engine, and stops the search. Upon receiving the request, the online IA-RWA engine changes its state to busy, notifying the NPOT scheduler manager accordingly, and processes the request. Once the request is processed, the online IA-RWA engine changes its state back again to free and notifies the NPOT scheduler manager. Upon this notification, if any request is waiting in the scheduler, the same operation described above is performed again and again, thus emptying the queues in the scheduler according to their priority.

B. Hardware-Accelerated Q-tool

The online IA-RWA module in the NPOT invokes the Q-tool in order to guarantee acceptable QoT of the newly established lightpaths and the potentially disrupted active ones. The Q-tool takes both linear (ASE, CD, FC, and PMD) and non-linear (SPM, XPM, and FWM) impairments into consideration. Computation of all these impairments by analytical expressions and numerical simulations to come up with a lightpath QoT prediction is a computationally intensive task, which may incur significant running times on a standard computer. In light of this, the Q-tool module in the centralized NPOT has been implemented on FPGA hardware, coupled to the NPOT following a client–server model.

The main goals targeted in the design of the hardware-accelerated Q-tool were 1) drastically reduced QoT estimation time compared with a purely software-based Q-tool implementation and 2) customizability and programmability. Specifically, the former aimed at making the Q-tool suitable for incorporation in dynamic network environments with delay-sensitive services, such as lightpath restoration. The latter, in turn, attempted to provide a highly customizable platform easing the implementation of new IA-RWA algorithms with fast turnaround time.

To the former goal, two contributions to the total Q-tool response time were identified, namely, the actual QoT estimation time and the communication overhead. Between them, the QoT estimation time clearly dominates over the second contribution, especially in highly loaded network scenarios, where the QoT of a large number of lightpaths has to be computed. Hence, most research and engineering effort was devoted to the hardware implementation of the QoT estimation over the FPGA through paralleled pipelining. In particular, four separated parallel pipelines are implemented, executing

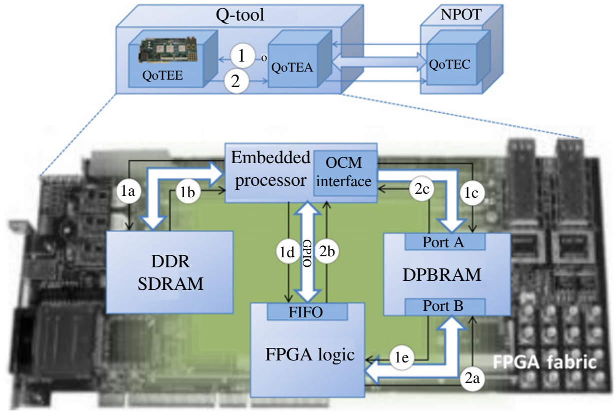

Fig. 3. (Color online) Hardware implementation of the Q-tool. The numbers and letters show the sequence of operations to estimate the QoT of an incoming lightpath request.

five top-level functions, namely, EyeClosure (involving SPM, CD, and FC), XPM, FWM, PMD, and ASE. The last function, ASE, has inputs that depend on the outcome of the function EyeClosure; thus EyeClosure and ASE are executed sequentially in pipeline 1, whereas the remaining functions are executed one in each pipeline. In this way, each function provides the impact of individual PLIs, which are eventually combined to come up with a single Q-factor figure.

Regarding customizability and programmability, implementations from scratch were identified as a major cause of high turnaround times that should be avoided. To address this, common computation modules shared by most custom Q-factor algorithms (e.g., FFT, IFFT, and trigonometric functions) were identified, implemented, and provided as extensions in the Q-tool, so that the algorithm programmers can select any subset of such modules that they need. It should be mentioned that, in the implementation of the modules, an ideal balance between hardware and software-based functions was sought, so as to achieve flexibility and performance, while avoiding software bottlenecks.

Figure 3 shows a top-level picture of the hardware-accelerated Q-tool with its main architectural modules. Specifically, the Q-tool is composed of two different modules: the QoT estimation agent (QoTEA) and estimation engine (QoTEE). The QoTEE module is deployed in a Xilinx Virtex IV FPGA and is responsible for the actual QoT estimation. In turn, the QoTEA runs on a 300 MHz IBM PowerPC 405 hard core embedded inside the FPGA fabric with 1 GB DDR2 memory, and is responsible for receiving the lightpath QoT estimation requests coming from the Q-tool client (QoTEC) in the online IA-RWA module and sending them to the QoTEE. The DDR SDRAM is used to store input data coming from the QoTEC in the request. For the computation, the data are shifted to the dual port block RAM (DPBRAM), so as to be rapidly accessible from the FPGA logic. The general purpose IO (GPIO) is used by the software program to initialize the estimation of the QoT in the FPGA and make sure that all computation requests are served. When the computation ends, the FPGA sends the Q-factor figures back to the DPBRAM and triggers an interruption signal through the GPIO. Upon this signal, the software module retrieves the results from the

DPBRAM and sends them back to the QoTEC. More details about the hardware-accelerated Q-tool implementation and testing can be found in [28].

C. Resource Pre-reservation Protocol

Multiple requests must be processed one after another during the restoration procedures. Therefore, the state of the resources involved in a backup path computation must be rapidly updated for subsequent requests. In the initial NPOT implementation, a guard time was left between two consecutive route computations. This allowed the gPPD and gTED databases to be fed with the new PLI and wavelength availability information through the OSPF-TE protocol, once the signaling phase of the related connection was completed. As highlighted in [23], however, this operation forced the NPOT to remain idle for several seconds between successive computations, that is, the time it takes OSPF-TE to update, network wide, the new state of the involved resources. In order to avoid the aforementioned idle periods in the centralized NPOT, which would likely lead to undesirably high restoration times, a resource pre-reservation protocol is proposed in this work, preventing the use of resources assigned to lightpaths under establishment when processing new requests.

It is worth mentioning that the centralized NPOT follows a stateful approach, keeping track of both available network resources and active lightpaths in the network (stored in the gPPD database, as mentioned in Subsection II.B). Compared to the standard path computation element (PCE) [29], where a stateless implementation may be enough to compute end-to-end routes, based only on the network resource availability, statefulness is mandatory in the centralized NPOT, so as to account for the non-linear PLIs (SPM, XPM, and FWM) when assessing the feasibility of the incoming lightpath requests. The proposed pre-reservation protocol exploits the statefulness in the centralized NPOT to achieve its purposes.

Aiming to avoid the gPPD and gTED inaccuracies between successive path computations, two alternative states are introduced for the lightpaths stored in the gPPD database, namely, *under establishment* and *active*. When a route is successfully computed for an incoming lightpath request, and before replying to the source node that performed the request, the lightpath is stored in the gPPD and its state is set to under establishment. Moreover, those resources supporting the lightpath are blocked, thus preventing them from being used in successive lightpath computations, while waiting to be eventually committed. Specifically, a timeout counter is set for every lightpath under establishment. In this way, if no commit request is received before the timeout counter expires, the lightpath is released and its respective resources are unblocked. Finally, the computed route and wavelength for the lightpath, together with the identifier of the lightpath, are sent back to the source node of the request.

Upon receiving the route and wavelength from the centralized NPOT, the source node of the connection starts the signaling procedures through the standard GMPLS RSVP-TE protocol. In the case that the lightpath is successfully established, a commit message is sent from the source node to the centralized NPOT, containing the identifier of the lightpath whose state has to be changed to active, and its related network resources committed. Conversely, if the signaling of the lightpath fails for any reason, the source node sends a rollback message, containing the identifier of the lightpath under establishment that has to be released and its resources unblocked.

(a)
```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Message type          |        Message length         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             LSP Id                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Src node Id                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Dst node Id                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

(b)
```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Message type          |        Message length         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             LSP Id                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             Result                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Wavelength                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            Node Id 1                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Local if Id 1                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              . . .                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Node Id N-1                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Local if Id N-1                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            Node Id N                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

(c)
```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Message type          |        Message length         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             LSP Id                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             action                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Fig. 4. Pre-reservation protocol message formats: (a) request message, (b) reply message, (c) action message.

For the sake of illustration, Fig. 4 depicts the format of the different pre-reservation protocol messages. As seen, in the request message, the identifiers of the source and destination nodes, together with the identifier of the new lightpath request (LSP Id) generated by the source node, are included. In turn, the reply message returns the result of the request (whether or not a route and a wavelength have been found), the assigned wavelength, and, then, the succession of node and output interface identifiers throughout the path. Finally, the action message serves the purpose of informing the centralized NPOT about the success or failure of the signaling process. In such cases, the action to be performed is set to commit (action = 0) or rollback (action = 1), respectively. Note that, even though the proposed protocol is proprietary, it has been designed to tightly match the information needed by standard control plane protocols, such as RSVP-TE.

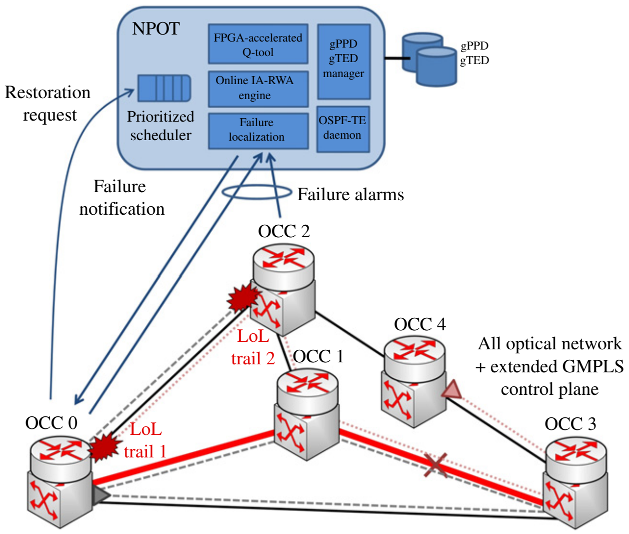

Fig. 5. (Color online) Impairment-aware restoration example: a bidirectional lightpath is established between nodes 0 and 3, and a failure occurs on the link connecting nodes 1 and 3; m-trails 1 and 2 become disrupted.

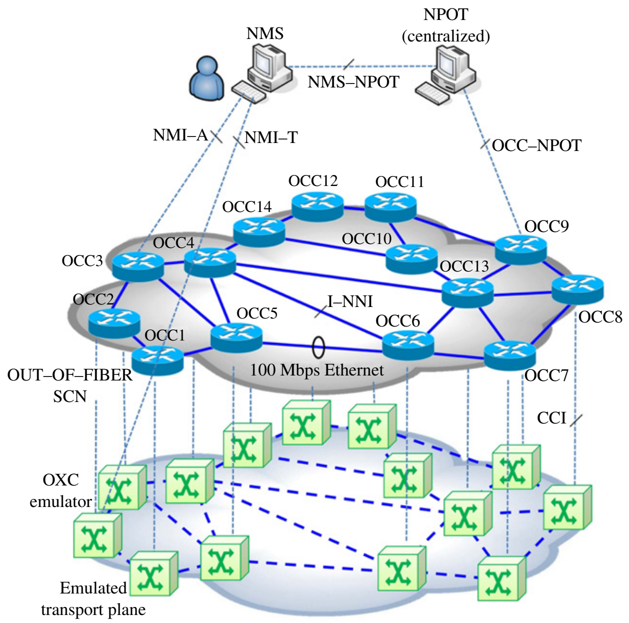

Fig. 6. (Color online) Test-bed used for the experiments: a 14-node all-optical meshed transport plane with 23 links is emulated.

D. Impairment-Aware Lightpath Restoration Example

Now that the enhancements to the NPOT for fast restoration have been presented, this subsection shows an example with all the interactions among the NPOT modules to achieve a successful restoration. For this purpose, the scenario depicted in Fig. 5 will be used, which represents the same network and m-trail design as in Fig. 2.

In the figure, assume that a lightpath is established between nodes 0 and 3, and a failure occurs on the link that connects nodes 1 and 3. This failure affects both the active lightpath, whose connectivity has to be restored as soon as possible, and the m-trails supported on the failed link, namely, m-trails 1 and 2. Specifically, m-trail disruptions will be detected on the trail termination nodes 0 and 2, respectively, where monitors are located. These LoL events are reported to the respective optical connection controllers (OCCs).

On detecting the failure, the OCCs also inform the centralized NPOT about the failure state and the m-trail that has detected it. Therefore, the failure localization module in the NPOT can localize the failure by a simple alarm code table lookup. Indeed, looking at the alarm code table in Fig. 2, one can easily see that code 6 (i.e., T2 = 1, T1 = 1, T0 = 0) identifies the failed link (1–3). After localizing the failure, the module isolates it by updating the gPPD and gTED databases accordingly. Next, it notifies the source nodes of all affected lightpaths, so that they can start the restoration procedures.

In the implemented centralized restoration scheme, failure restoration is delegated to the GMPLS control plane, in particular, to the standard RSVP-TE protocol. Therefore, as soon as the source node is notified about the failure, it requests a backup lightpath computation from the centralized NPOT for restoration purposes. The request is initially queued in the NPOT scheduler according to its priority, until served by the online IA-RWA module. If a feasible backup route is found, the lightpath is stored as under establishment in the gPPD database, and the involved resources are pre-reserved.

Finally, the source OCC triggers the RSVP-TE signaling to establish the backup path. Upon successful backup lightpath establishment, the source OCC sends a commit action message back to the centralized NPOT in order to commit the pre-reserved resources and change the state of the lightpath to active. Conversely, if the backup lightpath establishment is unsuccessful, the source OCC notifies the centralized NPOT with a rollback action message accordingly, which releases the lightpath under establishment and unblocks the previously pre-reserved resources.

IV. Evaluation Scenario

The performance of the proposed centralized lightpath restoration approach has been validated on the 14-node all-optical network test-bed depicted in Fig. 6. The test-bed describes the same topology as the 14-node Deutsche Telekom (DT) network [23], where 10 bidirectional wavelengths per fiber link at 10 Gbps have been assumed.

Each network node is composed of an OCC and a wavelength selective switch (WSS) based optical cross connect (OXC) emulator, both interconnected through the connection controller interface (CCI). This interface runs an extensible markup language (XML) based proprietary TCP protocol for the configuration of the OXC. Moreover, a simple network management protocol (SNMP) trap listener is implemented in the OCC side to gather possible LoL alarms coming from the managed OXC. The network management system (NMS) implements the network management interfaces, namely, NMI–A and NMI–T, which enable the communication to the OCCs and the OXCs, respectively. Besides, all OCCs

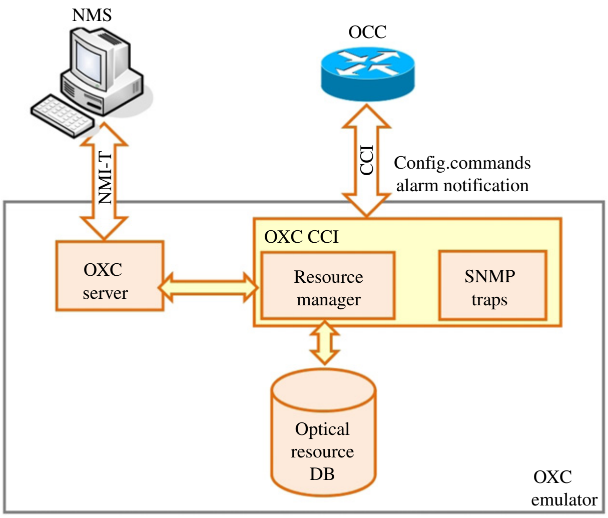

Fig. 7. (Color online) OXC emulator: comprising blocks and interfaces.

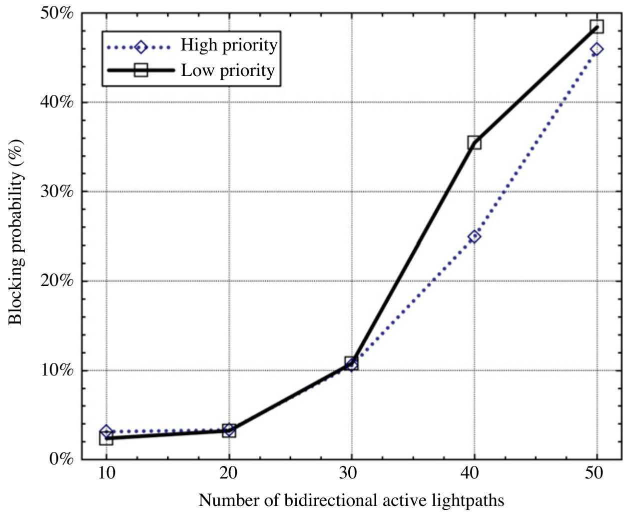

Fig. 8. (Color online) Lightpath restoration blocking probability as a function of the number of active lightpaths in the network for high and low priority traffic classes.

are interconnected to the NPOT through the OCC–NPOT interface. These interfaces implement TCP-based proprietary protocols.

The connectivity between OCCs is supported over 100 Mbps point-to-point Ethernet links, describing an out-of-fiber control plane configuration with the same topology as the emulated all-optical data plane. OCCs are deployed by means of Pentium IV Linux-based routers at 2 GHz (512 MB RAM memory) running the whole GMPLS protocol set, that is, the standard GMPLS RSVP-TE protocol for signaling, an extended version of the OSPF-TE protocol able to disseminate PLI information for routing purposes, and the GMPLS LMP protocol [30] for the initial resource discovery, as well as the maintenance of the control channels between OCCs.

Moreover, the centralized NPOT in the test-bed has been deployed on an Intel Core i3 530 Linux server at 2.93 GHz. This computer is directly connected to the FPGA board implementing the NPOT Q-tool over a Gigabit Ethernet full-duplex link. The specific details of the FPGA used can be found in Section III.

The software-based emulated OXCs also run over Linux-based machines. As depicted in Fig. 7, each OXC emulator implements a database with the physical resources that would be equipped in the real node. The status of these resources is managed during the network operation. In this way, when a lightpath is signaled and an OCC requests the allocation of a particular resource through the CCI, the managed OXC emulator updates the optical resource database, changing the status of that resource from available to reserved. A slightly more complex operation is needed for emulating failures in the network. Specifically, link failures can be forced at any moment from the NMS, which maintains information about the network resources and the m-trail configuration in the network. When a failure is forced for a particular link, the NMS obtains the m-trail endpoint that should have LoL due to the failure. Next, a message informing about the failure is sent through the NMI–T interface to that node. The m-trail endpoint receiving the message changes the state of the affected m-trail monitor to LoL, and triggers the failure restoration procedures as detailed in the previous section.

The following performance evaluation approach is adopted. Five different network scenarios with 10, 20, 30, 40, and 50 bidirectional lightpaths between randomly selected node pairs with acceptable QoT are considered. A high or a low priority is assigned to these lightpaths according to a 25%–75% high–low priority ratio. In each of these scenarios, we load the network until we have the desired number of active lightpaths. Then, a failure is forced on a randomly selected link among those links supporting at least one lightpath. Such operations are performed 100 times for each scenario, ensuring in this way a statistically relevant number of experiments.

V. Experimental Results

Figure 8 depicts the observed restoration blocking probability in every evaluated scenario, that is, the probability that a restoration request is blocked due to either unacceptable QoT or resource unavailability. As can be observed, this probability increases with the number of active lightpaths in the network. This can be justified by the fact that a lower number of spare resources exist in the network. Hence, while some lightpath restoration requests cannot be completed due to resource unavailability, others also fail due to very long physical routes that do not match the QoT requirements. It is interesting to highlight that, in the scenario with 40 active lightpaths, high priority traffic experiences notably lower restoration blocking probability, since high priority restoration requests are served first using the spare network resources. With 50 active lightpaths, however, both curves get closer again because of the general resource scarcity, as only 10 wavelengths are available per fiber link. In fact, this resource scarcity is the main cause of the increased restoration blocking probability observed in the highest loaded network scenarios.

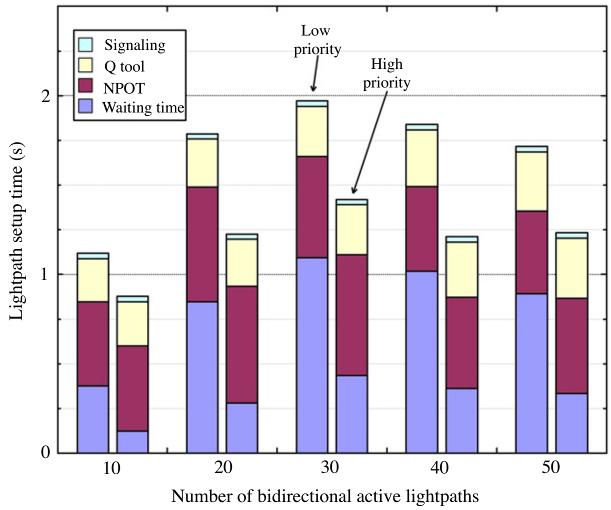

Fig. 9. (Color online) Average setup time breakdown as a function of the number of active lightpaths in the network for high priority (right bars) and low priority (left bars) traffic classes. The waiting time in the prioritized scheduler, NPOT processing time, Q-tool computation time, and backup path signaling time contributions have been considered.

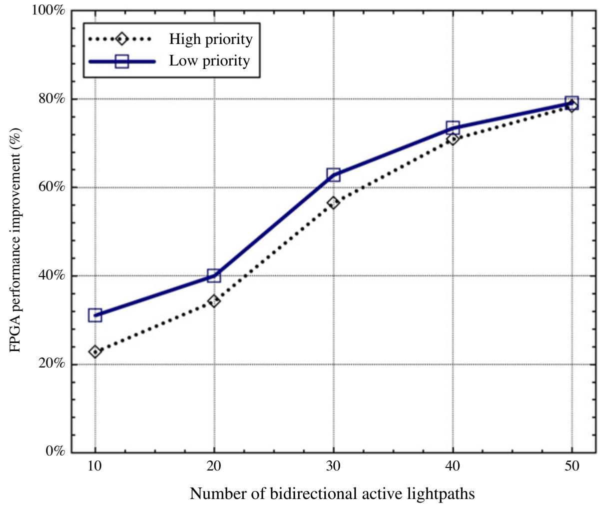

Fig. 10. (Color online) Relative FPGA QoT estimation time improvement as a function of the number of active lightpaths in the network compared to the case where a software-based Q-tool is used.

Figure 9 plots the average restoration times in the different load scenarios, which have been additionally split into their major contributions. Specifically, the average of the waiting time in the prioritized scheduler, the processing time in the NPOT, the Q-tool computation time, and the RSVP-TE backup path signaling time have been considered. The processing time in the NPOT comprises the total amount of time needed to process the incoming route request once it comes out of the prioritized scheduler, execute the IA-RWA algorithm, collect the resulting backup path, and send it back to the source node OCC. It can be observed that the restoration times start to increase with the number of lightpaths in the network. This is mainly caused by the increasing number of connections affected per failure that have to be queued in the prioritized scheduler until they can be processed by the NPOT online IA-RWA module. Indeed, while the other contributions remain quite steady, the waiting time incurred in the prioritized scheduler highly determines the restoration time behavior. This demonstrates the necessity of a resource pre-reservation strategy to ensure fast resource update between successive computations, together with priorities in the NPOT scheduler, guaranteeing low restoration times for the high priority critical traffic.

Focusing on the restoration time differences between high and low priority traffic, these also increase with the number of active lightpaths. For instance, while waiting time differences of approximately 250 ms are observed in the scenario with 10 active lightpaths, these rise up to 650 ms in the scenario with 30 active ones. Toward higher loads, however, the restoration times start to decrease. This can be explained by the fact that backup routes are more difficult to find. Thus, if no available route exists for a request, there is no need to invoke the NPOT Q-tool, which reduces the queuing time of the subsequent requests, still waiting to be served. This restoration time reduction is especially noticeable for low priority lightpaths, which are always the last ones to be served.

The results in Fig. 9 also prove the scalability of the FPGA-accelerated Q-tool that has been developed and integrated in the NPOT. In fact, Q-tool time differences of merely tens of ms are experienced when changing from the scenario with 10 active lightpaths to the one with 50. To assess the benefits of the hardware-accelerated Q-tool, Fig. 10 compares the performance of the accelerated Q-tool with a software-based Q-tool running along with the NPOT over the same hardware platform. In particular, the Q-tool time improvement achieved with the FPGA-based hardware acceleration is depicted for the five different network scenarios considered so far.

As shown, under light traffic loads, the improvements remain quite low (around 20%–30%), since the number of overlapping active lightpaths in the backup route, whose QoT feasibility has to be checked, is generally low. However, as more and more active lightpaths exist in the network, which can be supported on any of the backup route spans, QoT estimations become more computationally intensive. This makes software-based Q-tool implementations inefficient, and thus inappropriate for supporting critical network functionalities like dynamic impairment-aware lightpath restoration. For instance, in the scenario with 50 active lightpaths, FPGA-accelerated QoT estimations take around 330 ms, whereas they increase up to 1.5 s in the software-based Q-tool, leading to an improvement of around 78%. The slight Q-tool time improvement differences between low and high priority restoration requests arise from the fact that, as they are the last ones to be served, low priority backup path computations face a sensibly increased lightpath overlap, as the backup routes are typically longer than the primary ones. This increases the computational complexity in the Q-tool, making the software-based Q-tool perform worse in such situations.

In Fig. 11, the cumulative distribution function (CDF) of the lightpath restoration time for both high and low priority traffic classes has been illustrated. For these results, the entire set of successful restorations in all network scenarios has been used. The CDF function represents the probability that a lightpath

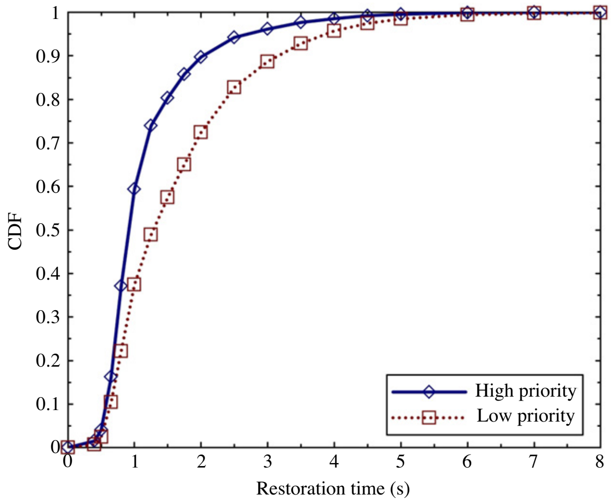

Fig. 11. (Color online) Cumulative distribution function (CDF) of the lightpath restoration time for both high and low priority traffic classes.

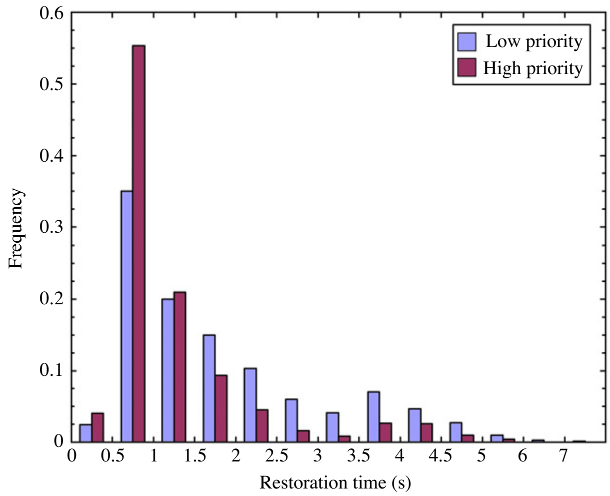

Fig. 12. (Color online) Restoration time frequencies for high and low priority traffic classes.

restoration takes a time lower than or equal to a certain value. As seen, the CDF of the high priority traffic grows significantly faster than the one of the low priority traffic, meaning faster restoration. For instance, it can be observed that more than 60% of the high priority restorations incur sub-second restoration times, and 90% of them are restored in less than 2 s. Moreover, focusing on the low priority traffic, even though it experiences longer restoration times, almost 90% of the low priority restorations are performed in less than 3 s, totally acceptable if we assign this class of service to the best-effort traffic on the network.

For better illustration, Fig. 12 plots the lightpath restoration time frequencies for both high and low priority traffic classes. From the figure, most of the high priority restorations (55%) are concentrated in the interval between 0.5 and 1 s, and none of them experience a restoration time higher than 5 s. In the case of the low priority traffic, it also experiences its peak (35%) in the same interval between 0.5 and 1 s, but with a longer and more pronounced tail that finally results in slightly increased restoration times.

Taking all these results into account, average restoration times of 1.16 and 1.64 s for the high and low priority traffic have been finally obtained, respectively. Such an outcome highlights the success of the resilience differentiation policies applied in the NPOT, so as to meet the diversity of resilience requirements posed by new emerging data applications. Moreover, the obtained values represent important restoration time improvements compared with previously reported NPOT implementations which lack FPGA-accelerated Q-tool and fast resource pre-reservation features for successive backup path computations.

VI. Concluding Remarks

This paper has reported the experimental demonstration of a dynamic impairment-aware restoration scheme implementing differentiated resilience support. The proposed scheme builds upon a centralized NPOT entity, which provides the network with online IA-RWA, as well as failure localization and isolation functionalities upon fiber link failures.

Through the paper, the functionality of the NPOT building modules involved in the impairment-aware lightpath restoration has been first described, followed by the presentation of key enhancements to the NPOT functionality for fast lightpath restoration. These include the implementation of a prioritized scheduler in the NPOT so as to efficiently serve lightpaths with differentiated resilience requirements, the deployment on FPGA hardware of the Q-tool module inside the NPOT, thus breaking the QoT estimation bottleneck in the IA-RWA process, as well as the implementation of a fast resource pre-reservation protocol allowing consecutive path computations without having to remain idle until the respective resource state updates. The dynamic impairment-aware restoration scheme has been experimentally evaluated on a 14-node GMPLS-enabled network test-bed deploying a centralized NPOT accordingly extended with the proposed enhancements for fast restoration. From the experimental evaluation, successful impairment-aware lightpath restoration times of around 1.16 and 1.64 s have been achieved for high and low priority traffic, respectively. These results, totally permissible in real network environments, leverage dynamic restoration as a key element to provide the desired reliability levels to the future optical Internet.

Acknowledgments

This work has been supported by the European Commission through the FP7 DICONET Project and the Spanish Science Ministry through the project ELASTIC (TEC2011-27310).

References

[1] J. Berthold, A. A. M. Saleh, L. Blair, and J. M. Simmons, "Optical networking: Past, present, and future," *J. Lightwave Technol.*, vol. 26, no. 9, pp. 1104–1118, May 2008.

[2] A. A. M. Saleh and J. M. Simmons, "Technology and architecture to enable the explosive growth of the Internet," *IEEE Commun. Mag.*, vol. 49, no. 1, pp. 126–132, Jan. 2011.

[3] A. Jajszczyk, "Automatically switched optical networks: benefits and requirements," *IEEE Commun. Mag.*, vol. 43, no. 2, pp. S8–S13, Feb. 2005.

[4] W. D. Grover, *Mesh-Based Survivable Networks: Options and Strategies for Optical, MPLS, SONET, and ATM Networking*. Prentice-Hall PTR, 2004.

[5] A. Farrel and I. Bryskin, *GMPLS: Architecture and Applications*. Academic Press, 2006.

[6] X. Yang and B. Ramamurthy, "Dynamic routing in translucent WDM optical networks: The intradomain case," *J. Lightwave Technol.*, vol. 23, no. 3, pp. 955–971, Mar. 2005.

[7] V. Anagnostopoulos, C. Politi, C. Matrakidis, and A. Stavdas, "Physical layer impairment aware wavelength routing algorithms based on analytically calculated constraints," *Opt. Commun.*, vol. 270, no. 2, pp. 247–254, Feb. 2007.

[8] H. Yurong, J. P. Heritage, and B. Mukherjee, "Connection provisioning with transmission impairment consideration in optical WDM networks with high-speed channels," *J. Lightwave Technol.*, vol. 23, no. 3, pp. 982–993, Mar. 2005.

[9] J. He, M. Brandt-Pearce, Y. Pointurier, and S. Subramaniam, "QoT-aware routing in impairment-constrained optical networks," in *Proc. of IEEE GLOBECOM 2007*, Nov. 2007.

[10] A. Morea, N. Brogard, F. Leplingard, J. Antona, T. Zami, B. Lavigne, and D. Bayart, "QoT function and A* routing: an optimized combination for connection search in translucent networks," *J. Opt. Netw.*, vol. 7, no. 1, pp. 42–61, Jan. 2008.

[11] S. Rai, C. Su, and B. Mukherjee, "On provisioning in all-optical networks: an impairment-aware approach," *IEEE/ACM Trans. Netw.*, vol. 17, no. 6, pp. 1989–2001, Dec. 2009.

[12] S. Pachnicke, T. Paschenda, and P. Krummrich, "Assessment of a constraint-based routing algorithm for translucent 10 Gbits/s DWDM networks considering fiber nonlinearities," *J. Opt. Netw.*, vol. 7, no. 4, pp. 365–377, Apr. 2008.

[13] Y. Lee, G. Bernstein, D. Li, and G. Martinelli, "A framework for the control of wavelength switched optical networks (WSON) with impairments," *IETF draft*, Jan. 2012 [Online]. Available: http://tools.ietf.org/html/draft-ietf-ccamp-wson-impairments-10.

[14] F. Cugini, N. Sambo, N. Andriolli, A. Giorgetti, L. Valcarenghi, P. Castoldi, E. Le Rouzic, and J. Poirrier, "Enhancing GMPLS signaling protocol for encompassing quality of transmission (QoT) in all-optical networks," *J. Lightwave Technol.*, vol. 26, no. 19, pp. 3318–3328, Oct. 2008.

[15] R. Martinez, R. Casellas, R. Muñoz, and T. Tsuritani, "Experimental translucent-oriented routing for dynamic lightpath provisioning in GMPLS-enabled wavelength switched optical networks," *J. Lightwave Technol.*, vol. 28, no. 8, pp. 1241–1255, Apr. 2010.

[16] S. Azodolmolky, M. Klinkowski, E. Marín Tordera, D. Careglio, J. Solé-Pareta, and I. Tomkos, "A survey on physical layer impairments aware routing and wavelength assignment algorithms in optical networks," *Comput. Netw.*, vol. 53, no. 7, pp. 926–944, May 2009.

[17] H. Zeng, C. Huang, A. Vukovic, and M. Savoie, "Fault detection and path performance monitoring in meshed all-optical networks," in *Proc. of IEEE GLOBECOM 2004*, Nov. 2004.

[18] B. Wu, P. Ho, and K. Yeung, "Monitoring trail: On fast link failure localization in WDM mesh networks," *J. Lightwave Technol.*, vol. 27, no. 23, pp. 4175–4185, Dec. 2009.

[19] J. Tapolcai, B. Wu, and P.-H. Ho, "On monitoring and failure localization in mesh all-optical networks," in *Proc. of INFOCOM 2009*, June 2009.

[20] C. Mas, I. Tomkos, and O. Tonguz, "A failure location algorithm for transparent optical networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 8, pp. 1508–1519, Aug. 2005.

[21] S. Stanic, G. Sahin, H. Choi, S. Subramaniam, and H.-A. Choi, "Monitoring and alarm management in transparent optical networks," in *Proc. of IEEE BROADNETS 2007*, Sept. 2007.

[22] S. Azodolmolky, D. Klonidis, I. Tomkos, Y. Ye, C. V. Saradhi, E. Salvadori, M. Gunkel, K. Manousakis, K. Vlachos, E. Varvarigos, R. Nejabati, D. Simeonidou, M. Eiselt, J. Comellas, J. Solé-Pareta, C. Simonneau, D. Bayart, D. Staessens, D. Colle, and M. Pickavet, "A dynamic impairment aware networking solution for transparent mesh optical networks," *IEEE Commun. Mag.*, vol. 47, no. 5, pp. 38–47, May 2009.

[23] S. Azodolmolky, J. Perelló, M. Angelou, F. Agraz, L. Velasco, S. Spadaro, Y. Pointurier, A. Francescon, C. V. Saradhi, P. Kokkinos, E. Varvarigos, S. Al Zahr, M. Gagnaire, M. Gunkel, D. Klonidis, and I. Tomkos, "Experimental demonstration of an impairment aware network planning and operation tool for transparent/translucent optical networks," *J. Lightwave Technol.*, vol. 29, no. 4, pp. 439–448, Feb. 2011.

[24] L. Berger, "Generalized multi-protocol label switching (GMPLS) signalling resource reservation protocol-traffic engineering (RSVP-TE) extensions," *IETF RFC 3473*, Jan. 2003.

[25] D. Katz, K. Kompella, and D. Yeung, "Traffic engineering (TE) extensions to OSPF version 2," *IETF RFC 3630*, Sept. 2003.

[26] K. Christodoulopoulos, P. Kokkinos, and E. M. Varvarigos, "Indirect and direct multicost algorithms for online impairment-aware RWA," *IEEE/ACM Trans. Netw.*, vol. 19, no. 6, pp. 1759–1772, Dec. 2011.

[27] A. Haddad, E. A. Doumith, and M. Gagnaire, "A meta-heuristic approach for monitoring trail assignment in WDM optical networks," in *Proc. of IFIP/IEEE RNDM 2010*, Oct. 2010.

[28] Y. Qin, S. Azodolmolky, M. Gunkel, R. Nejabati, and D. Simeonidou, "Hardware accelerated impairment-aware control plane for future optical networks," *IEEE Commun. Lett.*, vol. 15, no. 9, pp. 1004–1006, Sept. 2011.

[29] A. Farrel, J. P. Vasseur, and J. Ash, "A path computation element (PCE)-based architecture," *IETF RFC 4655*, Aug. 2006.

[30] J. Lang, "Link management protocol (LMP)," *IETF RFC 4204*, Oct. 2005.

Jordi Perelló received his M.Sc. and Ph.D. degrees in Telecommunications Engineering in 2005 and in 2009, respectively, both from the Universitat Politècnica de Catalunya (UPC), Spain. Currently, he is an Assistant Professor in the Computer Architecture Department (DAC) at UPC. He has participated in various IST FP6 and FP7 European research projects such as EU DICONET, BONE, IST NOBEL 2, e-Photon/ONe+, and COST Action 291. Dr. Perelló has published more than 40 articles in international journals, conference proceedings, and book chapters. His research interests concern resource management, quality of service issues, and survivability of next-generation optical transport networks.

Salvatore Spadaro (M'10) received his M.Sc. (2000) and Ph.D. (2005) degrees in Telecommunications Engineering from the Universitat Politècnica de Catalunya (UPC). He also received a Dr.Ing. degree in Electrical Engineering from Politecnico di Torino (2000). He is currently an Associate Professor in the Optical Communications Group of the Signal Theory and Communications Department of UPC. Since 2000 he has been a staff member of the Advanced Broadband Communications Center (CCABA) of UPC. His research interests are in the fields of all-optical networks with emphasis on network control and management, resilience, and network virtualization.

Fernando Agraz received his M.Sc. degree in Computer Engineering in 2005 from the Universitat Politècnica de Catalunya (UPC). Since 2005 he has been working as a Research Engineer in the Optical Communications Group (GCO) at UPC, also preparing his Ph.D. He has also participated in various European research projects such as DICONET, IST Nobel Phase 2, and e-Photon/ONe+. His current research focuses on network management and routing in GMPLS-based networks.

Marianna Angelou (S'09) received her degree in Computer Science from the Aristotle University of Thessaloniki, Greece, in 2005. She received a M.Sc. in Information and Telecommunication Technologies from Athens Information Technology in 2008 and is currently pursuing a Ph.D. with the Universitat Politècnica de Catalunya, Barcelona. In January 2008 she joined the High-speed Networks and Optical Communications Group of Athens Information Technology, as a Research Scientist, working within the framework of EC funded research projects. Her research activities focus on cross-layer optimization techniques for core optical networks and cover a broad range of topics in that area, including physical layer awareness, energy efficiency, and networking with flexible/adaptive transmission characteristics.

Siamak Azodolmolky received his Computer Hardware Engineering (B.Eng.) degree from Tehran University in Iran in 1994 and his first M.Sc. in Computer Architecture from Azad University in 1998. He has worked with Data Processing Iran (ex-IBM), as a Systems Engineer and Senior R&D Engineer during 1992–2001. He received his second M.Sc. with distinction from the Information Networking Institute (INI) of the Carnegie Mellon University in 2006. In August 2007 he joined Athens Information Technology (AIT) as a Research Scientist. He received his Ph.D. degree from Universitat Politècnica de Catalunya (UPC) (BarcelonaTech) in 2011. He was the main technical investigator of the FP7 ICT EU DICONET project (STREP) and also the FP7 ICT EU BONE project (NoE). In August 2010 he joined the High Performance Networks Research Group of the School of Computer Science and Electronic Engineering of the University of Essex as a Senior Research Officer.

Yixuan Qin is a Senior Research Officer in the High Performance Networks Group (HPNG) at the University of Essex. He has a B.Sc. from Southwest Jiaotong University, China, and an M.Sc. and a Ph.D. from the University of Essex, UK, all in Electronic Engineering. His research interests include high performance computing, flexible networks, passive optical networks, optical burst switching, impairment-aware-based routing and GMPLS networks, high-speed digital system design, embedded system design, and software/hardware co-design. He has published more than 30 articles in high impact conferences, journals, and book chapters.

Reza Nejabati is currently a Lecturer and prior to that he was an RCUK Fellow at the University of Essex. His current research focuses on application of high-speed network technologies, service-oriented and programmable networks, and cross-layer network design. He is author and co-author of over 90 papers and 3 standardization documents.

Dimitra Simeonidou is the Head of the High Performance Networks Group. She joined Essex as an academic in 1998 after four years with Alcatel Submarine Networks. At Essex, she is leading a group of 30 researchers and Ph.D. students and currently involved in 19 national and EU funded projects. Her research focuses on the fields of optical networks, grid and cloud computing, and the future Internet. She is the author and co-author of over 300 papers, 11 patents, and several standardization documents.

Pannagiotis Kokkinos received his Ph.D. in 2010 from the Computer Engineering and Informatics Department of the University of Patras, Greece, in the field of Optical Grid Networks. He also holds an M.Sc. degree (2006) and a Diploma (2003) from the same department. Dr. Kokkinos has a strong technological and research background working in the private sector as a Software Engineer and in several national and European programs as a Researcher. He has a number of publications in highly respected conferences and journals. Dr. Kokkinos' current research interests include routing and scheduling algorithms for distributed systems.

Emmanouel (Manos) Varvarigos received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens in 1988 and M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 1990 and 1992, respectively. He has held faculty positions at the University of California, Santa Barbara (1992–1998, as an Assistant and later an Associate Professor) and Delft University of Technology, the Netherlands (1998–2000, as an Associate Professor). Since 2000 he has been a Professor in the Department of Computer Engineering and Informatics at the University of Patras, Greece. He is also the Scientific Director of the Greek School Network and Network Technologies Sector at the Computer Technology Institute (CTI). Professor Varvarigos has served on the organizing and program committees of several international conferences, primarily in the networking area, and in national committees. He has participated in more than 20 EU funded research projects in the areas of optical networking and grid computing and in many national research projects. He has also worked as a Researcher at Bell Communications Research and has consulted with several companies in the US and in Europe. Dr. Varvarigos has over 200 publications in international journals and conferences. His research activities are in the areas of optical networking, network protocols, grid and cloud computing, wireless ad hoc networks, and network services.

Ioannis Tomkos (B.Sc., M.Sc., Ph.D.) has been with the Athens Information Technology Center (AIT) since September 2002. In the past, he was Senior Scientist (1999–2002) at Corning Inc. USA and a Research Fellow (1995–1999) at the University of Athens, Greece. At AIT he founded and serves as the Head of the "High-Speed Networks and Optical Communication (NOC)" Research Group that was/is involved in many EU funded research projects as well as in national projects (including five running projects), within which Dr. Tomkos is representing AIT as Principal Investigator and has a consortium-wide leading role (e.g., Project Leader of the EU ICT STREP project DICONET, Project Leader of the EU ICT STREP project ACCORDANCE, Technical Manager of the EU ICT STREP project SOFI, Technical Manager of the EU IST STREP project TRIUMPH, Chairman of the EU COST 291 project, WP leader in many other projects).

Dr. Tomkos received in 2007 the prestigious title of "Distinguished Lecturer" of the IEEE Communications Society for the topic of transparent optical networking. Together with his colleagues and students he has authored about 400 peer-reviewed archival articles (an updated list of all published items may be found using the "Publish or Perish" tool, over 240 IEEE sponsored items may be found through IEEE Xplore), including over 100 journal/magazine/book publications and about 300 conference/workshop proceedings papers. Dr. Tomkos has served as the Chair of the International Optical Networking Technical Committee of the IEEE Communications Society (2007–2008) and the Chairman of the IFIP working group on "Photonic Networking" (2008–2009). He is currently the Chairman of the OSA Technical Group on Optical Communications (2009–2010) and the Chairman of the IEEE Photonics Society Greek Chapter (2010). He is involved in several high-level committees (including the OSA/IEEE Tyndall Award Evaluation/Nomination Committee, the FTTH Innovation Award Evaluation Committee, etc.). He is also the Chairman of the working group "Next generation networks" of the "Digital Greece 2020" Forum. He has been General Chair, Technical Program Chair, Subcommittee Chair, Symposium Chair, or/and member of the steering/organizing committees for major conferences (e.g., OFC, ECOC, IEEE GlobeCom, IEEE ICC, ICTON, ONDM, BroadNets, etc.) in the area of telecommunications/networking (more than 100 conferences/workshops). In addition, he is a member of the Editorial Boards of the IEEE/OSA *Journal of Lightwave Technology*, the IEEE/OSA *Journal of Optical Communications and Networking*, the *IET Journal on Optoelectronics*, and the *International Journal on Telecommunications Management*. Among many other guest editorials for special issues, he is the Chief Editor for a 2012 special issue on "The Evolution of Optical Networking" for the prestigious *Proceedings of IEEE*. He is a Senior Member of IEEE and a Fellow of the IET.