

# Assessing IP vs optical restoration using the open-source Net2Plan tool

Jose-Luis Izquierdo Zaragoza, Pablo Pavon-Marino  
Department of Information Technologies and Communications  
Universidad Politécnica de Cartagena  
Plaza del Hospital, 1, 30202 Cartagena, Spain  
Email: {josel.izquierdo, pablo.pavon}@upct.es

**Abstract**—Economic sustainability of the ever-increasing traffic requirements is built on the cost reduction of IP-over-WDM backbone networks. These networks are composed of a set of IP routers connected via all-optical circuits (or lightpaths), bypassing intermediate nodes, thanks to reconfigurable optical equipment. Today, traffic survivability is ensured by over-dimensioning the IP layer and using some IP/MPLS fast convergence/reroute mechanism, without making use of the reconfigurability of the optical layer. In this paper, we provide algorithms to dimension an IP-over-WDM network under two different resilience schemes: IP restoration and optical-followed-by-IP restoration. In both cases, the IP routing is controlled by an OSPF-like algorithm. In the optical-followed-by-IP restoration, OSPF reactions to failures are delayed by appropriately setting the OSPF hold-off timer, giving time to lightpath restoration attempts to be completed. We evaluate the performance, in terms of bandwidth requirements and traffic survivability, of these two resilience schemes. Our tests show that optical-followed-by-IP restoration is a cost-efficient alternative to IP restoration, still improving traffic survivability in most cases, without any degree of over-dimensioning. Results were obtained using the open-source Net2Plan tool, and source code is available on the website.

**Keywords**—IP restoration; multilayer restoration; bandwidth requirements; traffic survivability; IP; OSPF; WDM; Net2Plan

## I. INTRODUCTION

Internet traffic is exponentially growing with a compound annual growth rate (CAGR) larger than 20% [1]. In its turn, revenue growth is only linear, which implies that the revenue-traffic product is decreasing [2]. Hence, in order to maintain the telecom market sustainable, CAPital EXpenditures (CAPEX) and OPERational EXpenditures (OPEX) must be dropped significantly [3].

Backbone networks, and progressively metro networks, are based on the multilayer IP-over-WDM paradigm. In these networks, IP traffic is routed on top of a virtual topology of all-optical circuits (or lightpaths), which are routed over the fiber topology. As shown in Fig. 1, IP traffic is accommodated into lightpaths, which become IP links, going from the ingress node to the egress node. Intermediate IP routers are bypassed via optical cross-connects (OXC) using reconfigurable add/drop multiplexers (ROADMs), which perform optical switching at the WDM layer. Note that in these networks, a single fiber failure implies that multiple lightpaths will fail simultaneously.

This work was partially supported by the FPU fellowship program of the Spanish Ministry of Education, Culture and Sport (ref. no. FPU12/04571) and research project 3704/13TIC.

Considering that lightpaths are expected to carry multi-Gbps traffic, survivability becomes a major concern.

The large majority of backbone networks today adopt only-IP restoration schemes to protect the network under both optical and IP failures, thanks to over-dimensioning of the IP layer and IP/MPLS fast convergence/reroute mechanisms [3].

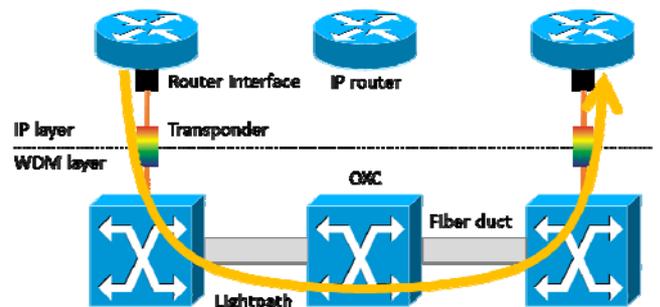


Fig. 1. IP-over-WDM network architecture.

From the point of view of costs, reconfigurability of the WDM layer to provide optical restoration may play here an important role. If the optical layer is able to recover lightpaths between IP routers, savings in CAPEX (i.e. number of IP ports or transponders) and OPEX (i.e. power consumption) may be achieved with minimal over-dimensioning.

In this paper, we investigate this question by comparing two different resilience schemes: IP restoration, and optical-followed-by-IP restoration. On the one hand, the classical IP restoration policy is based on fast IGP (e.g. OSPF or IS-IS) convergence, driven by link-state update messages, to reroute the traffic over the surviving virtual topology. On the other hand, the optical-followed-by-IP restoration policy tries to first reallocate failing lightpaths over the surviving physical topology, and then the IP layer reroutes the traffic on top of the surviving virtual topology. To configure this behavior, the OSPF or IS-IS hold-off times must be dimensioned appropriately, so that IP layer restoration starts when optical restoration has been completed (successfully or not).

In our tests, we use our open-source Net2Plan tool [4] to compare both approaches. First, the offline network design tool is used to implement and execute our proposed multilayer network dimensioning algorithms for each scenario, where the number of established lightpaths is used as a cost indicator.

Then, the resilience simulator is used to evaluate the traffic survivability on each of the two schemes. Interested readers can find all the source code on the Net2Plan repository [5].

The rest of the paper is organized as follows. In Section 2, we review some previous works in IP and/or optical restoration. In Sections 3 and 4, we describe a planning algorithm as well as a provisioning algorithm implementing both resilience schemes. In Section 5, we report the results obtained within Net2Plan. Finally, Section 6 concludes the paper.

## II. RELATED WORK

Survivability in IP-over-WDM networks is a well-discussed topic. Note that IP and WDM layers provide similar recovery functionalities, and some routing instabilities may appear [6]. Therefore, different approaches to coordinate or isolate those functionalities have been proposed in the literature.

In [6], authors describe a framework for multilayer survivability, identifying several recovery options. On the one hand, single layer strategies in the optical layer are described as effective options for recovery of burdensome failures such as fiber cuts. The benefit comes from the fact that only a single failure is to be treated, and the number of required actions will be minimal [7]. On the other hand, IP restoration is able to recover from failures at any layer. The point is that a lot of actions may be required, since multiple entities are involved due to the finer granularity. Hence, coordinated multilayer strategies are presented as the preferred option, due to the complimentary role of restoration at each layer [3].

It is worth noting that in multilayer schemes some coordination rules are mandatory in order to ensure efficient interoperation between the different layers [16]. These rules are known as escalation strategies. Here, hold-off timers are defined so that the restoration mechanism at a given layer is not triggered until a certain time after the failure, to give some time for the other layer(s) to respond. Typical escalation proposals often fall into one of the following categories: bottom-up and top-down [6]. In the former, recovery starts from the optical layer and escalates upwards to more complex higher layer mechanisms. Hence, unrecovered traffic at the lower layer might be recovered at the upper layer. In contrast, the latter triggers IP mechanisms first, and then escalates downwards. The reason behind this rationale is that higher layers are able to differentiate service types and may try to recover traffic following some priority criteria. Recently, all these escalation strategies have been questioned, and novel cross-layer negotiation mechanisms are being introduced [3].

Nonetheless, the design of survivable multilayer networks is not trivial, since the mapping of an IP topology to a physical topology is a challenging task. We remark that in order to guarantee a fully restorable logical topology, enough excess capacity must be provisioned to support rerouting of disrupted traffic. The rest of this section is devoted to examine some previous works on the planning of survivable multilayer networks.

As in [6], authors in [7] present a conceptual framework on IP/WDM layer protection/restoration, and review several single-layer resilience schemes. Then, authors discuss advantages and disadvantages of cross-layer approaches for IP-over-WDM networks. To conclude, they present a heuristic that is able to balance between traffic protected by the WDM layer and traffic being restored by the IP layer, depending on topology constraints and minimization of network cost.

In [3], authors argue that multilayer schemes are mandatory in IP-over-WDM, not only to reduce costs, but also to minimize traffic impact and to fulfill new flexibility requirements of elastic optical networks. Throughout the paper, they emphasize that escalation strategies are insufficient, and pose the need of cross-layer mechanisms involving a negotiation process involving both layers. Moreover, they recognize the non-trivial nature of the revert-back-to-normal process. Then, they describe for the first time a step-by-step mechanism that could be implemented in practice.

Different planning models solving the virtual topology design (VTD) problem have been studied in the literature. The VTD problem involves routing and wavelength assignment of lightpaths (RWA, or lightpath routing) as well as traffic routing on top of the lightpaths. Essentially, all of them focus on determining the spare/excess/backup capacity requirements to support some degree of protection.

In [8], two metrics to assess the quality of IP restoration schemes are presented (load factor and spare factor), investigating the impact of lightpath routing on spare/excess capacity requirements. Then, a two-step algorithm solving an optimization model in each step. In contrast, authors in [9] are interested in determining the bandwidth requirements to ensure 100% IP restoration. Again, they propose a two-stage heuristic scheme. On the other hand, authors in [10] suggest to maximize the bandwidth usage by sharing the backup capacity of wavelength services to support IP traffic, even guaranteeing some capacity for IP services [11].

However, to the best of our knowledge, only few papers are focused on comparing the dimensioning of the IP layer (i.e. number of established lightpaths, or bandwidth requirements using the terminology in [12]) for IP restoration vs optical protection/restoration schemes. For example, optical protection/restoration may be costly in terms of CAPEX that IP restoration. On the other hand, optical equipment is less power-hungry than IP equipment, thus reducing energy consumption and balancing the CAPEX vs OPEX trade-off. Still, depending of the type of traffic (IP services vs wavelength services), a combination of both mechanisms should be the best option.

In [13], authors compare optical protection vs IP restoration in the context of single link failures. For IP restoration, they propose to use a two link-disjoint path pair so that one of them will be always operation. For optical protection, a shared-path approach is presented. Their results show that optical protection should be preferable in terms of bandwidth requirements.

In [12], authors also present a comparison between different recovery schemes. In this case, they also consider multiple-link failures. As expected, results show that optical

protection is preferred over IP restoration in terms of bandwidth requirements.

### III. IP-R AND OP-IP-R PLANNING ALGORITHM

In this section, we describe the planning algorithm for dimensioning multilayer IP-over-WDM networks taking into account the two different resilience policies under study. We assume that IP routing is governed by OSPF, with all IP links (or lightpaths) having a static IGP weight equal to one. Equal-Cost Multi-Path (ECMP) rule is applied according to the standard. The algorithm does not take into account any delay constraint either in the IP layer or in the WDM layer. In addition, we assume that ROADMs do not allow wavelength conversion, and signal regeneration is not required.

Our algorithm is divided into two stages. First, we use a heuristic, inspired in the well-known HDLA algorithm in [14], to solve VTD problem allocating enough lightpaths to carry all the offered traffic to the network without grooming. Then, the second stage, only use for IP-R, is able to set up new lightpaths in such a way that the traffic can be still routed across the network under different possible failure states.

The pseudo-code of the algorithm is given below:

Stage 1: Route traffic over (multiple) single-hop lightpaths

Step 1. Given the input IP traffic matrix  $M$ , make a copy  $R = M$ .

Step 2. Select the ingress-egress node pair with the largest traffic value in  $R$ .

If we have tried all ingress-egress node pairs with non-zero traffic, stage 1 ends.

Step 3. Try to allocate a new single-hop lightpath for the selected demand in the previous step, over a set of pre-computed physical paths using first-fit wavelength assignment.

If allocated, decrease  $R$  accordingly, and go to Step 2. Otherwise, it will fail since the network is unable to carry all traffic in an unprotected manner.

(IP-R only) Stage 2: Enforce traffic survivability under failure

Step 4. Iterate over the set of pre-defined failure states, and compute the maximum carried traffic by each lightpath in any state.

Step 5. Select the highest loaded lightpath.

If no lightpath is over-subscribed, finish the algorithm.

Step 6. Try to allocate a new lightpath connecting the node pair corresponding to the lightpath selected in Step 5.

If allocated, go to Step 4. Otherwise, try with the second most loaded lightpath, and so on. If no lightpath could be added, finish the algorithm

Note that this algorithm (algorithm 1 from now on) establishes multiple single-hop lightpaths between each node pair during

the stage 1, while the excess capacity on lightpaths is shared during failure states thanks to the grooming enforced by OSPF.

The inputs of the algorithm are the following:

- A fiber plant connecting a set of nodes
- The number of wavelengths per fiber
- An IP end-to-end traffic matrix, measured in Gbps
- A binary rate per lightpath, measured in Gbps
- A parameter 'K' indicating the number of k-loopless shortest paths that are pre-computed to allocate lightpaths.
- A parameter 'onlyIPRestoration' that determines if the second stage is applied or not (which happens to overdimension the optical layer in the IP-restoration case).

At the end of this algorithm, we collect the number of established lightpaths, which is used as the cost metric.

### IV. IP-R AND OP-IP-R PROVISIONING ALGORITHM

In this section, we describe the provisioning algorithm (named as algorithm 2), which is able to react to failure/reparations applying either IP-R or Op-IP-R policy. We assume that the control plane is able to notify failures and reparations, and the algorithm returns a set of actions to be applied by the control plane.

Upon failures, the algorithm receives the set of affected fibers. Then, if the Op-IP-R policy is active, the algorithm tries to re-route the failing lightpaths using an adaptive shortest-path algorithm, which has information about wavelength usage, and search for the shortest route having a wavelength available along the path. In case of IP-R no action is made, since we rely on IP rerouting. By the way, new traffic routes are computed over the surviving lightpath topology using the same OSPF weights as in algorithm 1.

Note that algorithm 2 defines a bottom-up escalation strategy for the Op-IP-R case, avoiding triggering recovery mechanisms at both layers simultaneously, which may affect to IGP convergence [15]. Regarding to a practical implementation of this strategy, we are neglecting the effect of the hold-off timer [16], which consists in delaying the IP restoration mechanism is not triggered until a certain time after the failure, to give some time for the optical layer to respond. The study of the influence of the hold-off timer is out of the scope of the paper.

On the other hand, upon reparation the algorithm receives the set of repaired fibers. Here, the algorithm tries to revert as many lightpaths as possible to their original primary routes (determined in algorithm 1). Again, IP rerouting is used to re-route traffic over the new virtual topology. Intuitively, reparation is independent of the restoration policy.

## V. RESULTS

This section collects a set of testing results obtained for exploring the cost and availability trade-off appearing in the

choice between a network operated with (i) IP restoration (IP-R) and (ii) optical-followed-by-IP restoration (Op-IP-R).

#### A. Testing scenario

Our testing scenario is an IP-over-WDM network composed of a set of fiber links between sites following the NSFNET topology [14]. Each network site is equipped with IP routing equipment and an ROADMs, and is connected to other sites via a virtual topology of 40 Gbps lightpaths. Therefore, each lightpath appears as a single link to the IP layer, with a capacity of 40 Gbps. Lightpaths occupy 50 GHz slots in the ITU G.694.1 grid [17], and we assume that the bandwidth limitations of the EDFAs limit the number of available WDM channels to 40 wavelengths per fiber. Topology diameter is around 4500 km, and we assume that lightpaths between any node pair can be allocated without any intermediate signal regeneration process. This is compatible with current capabilities of commercial coherent optical transponders.

The network shared risk groups (SRGs) considered in the network plan are the possible cuts in the ducts carrying the optical fibers, caused by civil engineering activities, natural disasters, etc. A duct cut would mean the simultaneous failure of the two fibers (one in each direction) contained in the conduct, and thus the failure of all the lightpaths traversing them. We assume that each duct cut has a random time to repair, with negative exponential distribution and an average of 12 hours. Time between failures in a duct has also a random negative exponential distribution. Each duct is supposed to have, in average, 1.9 failures per 1000 km of length per year. This is consistent with empirical data collected in [18].

We apply the algorithm 1 to determine the set of lightpaths to be established. All lightpaths are assigned an OSPF weight of one. Then, IP traffic is routed through the minimum possible amount of lightpaths (in general one), and the Equal-Cost Multi-Path (ECMP) rule is applied to balance traffic. The IP-R and Op-IP-R scenarios considered (i) have a different initial design, and (ii) react differently to link failures:

- In IP restoration case, algorithm 1 dimensions the number of lightpaths between each node pair, using the first and second phase of the algorithm. Thus, the algorithm first attempts to allocate enough lightpaths to carry the traffic in the no-failure situation, without oversubscribing any lightpath. Then, the second algorithm stage adds enough lightpaths to guarantee that under any *single SRG failure* (duct cut), all the traffic could be routed without link oversubscription. During network operation, failing lightpaths are not restored, and just appear as inactive links to the IP layer. The network reacts in a similar manner to failure and reparation events: the OSPF algorithm recomputes the routing using only the surviving IP links (lightpaths). As a result, depending on the surviving virtual topology, some lightpaths may be oversubscribed, that is, receive more traffic than what they can carry. Oversubscribed links produce higher delays, packet drops, etc. To reflect these effects, in this paper we account all the traffic traversing an oversubscribed lightpath as lost traffic. Also, all the

traffic for which OSPF does not find a route (which happens if the surviving topology of lightpaths is not connected), is considered lost.

- In the optical-followed-by-IP restoration case, the initial design is determined by the first stage of algorithm 1. This means a lower number of lightpaths than in the previous case, since no extra lightpaths are allocated for guaranteeing the IP traffic is carried under any single SRG failure. This is because the network reacts to lightpath failures attempting to reroute each one of the affected lightpaths over the surviving fibers. This rerouting is supposed to be handled by either a distributed (e.g. GMPLS-like) or a centralized (e.g. PCE-like or supported by a centralized management plane) mechanism. Those lightpaths for which a valid RWA is not found, are just not restored. Once the lightpath restoration stage is completed, the IP/OSPF layer observes the surviving lightpath topology, and starts its IP traffic rerouting in the same manner as the one described in the IP restoration case. Similarly, in our blocking statistics we consider that all the traffic traversing oversubscribed links or that OSPF cannot route, is lost.

The input IP traffic to the network is a scaled version of the reference traffic matrix for this network [14]. We repeat our tests for 35 growing traffic matrices, ranging from a low load situation (1 Tbps of total offered traffic) to a high-load situation (18 Tbps of total offered traffic). We observed that for 18.5 Tbps of offered traffic, stage one of algorithm 1 did not succeed to carry all the traffic without lightpath oversubscription in the no-failure situation. Note that stage two of the algorithm 1, only applied when IP-R, might fail to allocate all the extra lightpaths required to guarantee carrying all the traffic for every single SRG failure, even though in situations in which a feasible solution could be found in the stage one. In our experiments, algorithm 1 was able to complete stage two for all traffic loads.

#### B. Network cost assessment

Fig. 2 provides the comparison results in terms of network cost between the IP-R and Op-IP-R strategies. We use the number of lightpaths to establish as our metric of interest, since this is proportional to the number of transponders needed in the network. Naturally, this number will be always higher in the IP-R case since the second stage of algorithm 1 is executed only in this mode, adding extra lightpaths respect to the Op-IP-R dimensioning, to guarantee no traffic loss under any single SRG failure. We are interested in observing how large this transponder over-dimensioning in the IP-R case is. This information is relevant for those network operators considering lightpath restoration as an option: they would have to balance this quantitative transponder cost, to the alternate more subjective extra cost associated to Op-IP-R case with respect to IP-R. In particular, lightpath restoration requires equipping the network with agile versions of ROADMs, and a more complex WDM line engineering that is robust under lightpath rerouting.

Fig. 2 shows that the IP-R scheme requires in general more than 30% of extra lightpaths, even exceeding 40% in some

cases (from 9 to 13.5 Tbps), to carry all the traffic under any single SRG failure, in contrast to the “unprotected” Op-IP-R case.

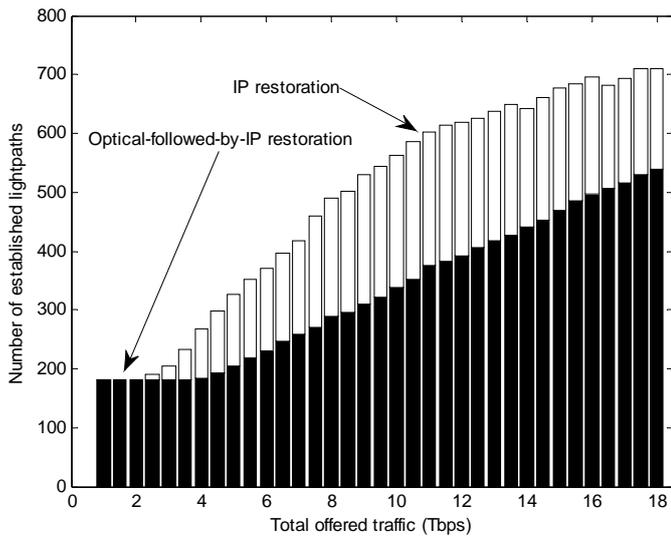


Fig. 2. Cost comparison for the planning phase in NSFNET network.

### C. Traffic survivability assessment

Although resilience schemes would enforce some degree of network connectivity, the focus of these schemes must be on ensuring that as much traffic as possible will survive under different failure scenarios. Fig. 3 provides the comparison results in terms of traffic survivability. Results have been obtained by simulation using the Net2Plan 0.2.3 Resilience Simulation tool. In this tool, failure and reparation events can be generated using a user-defined or built-in scheme, and reactions to these events can be also defined by a user-defined or built-in algorithm. In both cases, algorithms are instances of Java classes implementing the appropriate interface (see [5] for details). For failure and reparation event generation a built-in procedure is used, which randomly creates failures and reparations according to the SRG information provided. The algorithm reacting to failures and reparations matching the behavior described for the IP-R and Op-IP-R case has been implemented as a single class, which can operate in any of the two modes, depending on an input parameter. Recall that this and all the other algorithms in this paper are available at Net2Plan repository [5].

In Fig. 3, we measure traffic survivability as the average fraction of IP traffic that is being blocked by the network (the less the better), when it is randomly observed. In other words, it is equal to one minus the average traffic survivability. IP traffic losses are accounted for all the IP traffic for which there is no valid route, or which traverses an oversubscribed lightpath. Results show that the Op-IP-R scheme outperforms IP-R in terms of blocked traffic when offered traffic is less or equal than of 16 Tbps, even by more than one order of magnitude for traffic below 13 Tbps. Interestingly, when traffic grows over 16 Tbps, IP-R seems to be the best option. This is due to the fact that the optical restoration module is not able to reallocate all the failing lightpaths, which makes surviving

lightpaths become oversubscribed. In contrast, the IP restoration case is able to accommodate more traffic, thanks to the over-dimensioning process at stage 2 of algorithm 1.

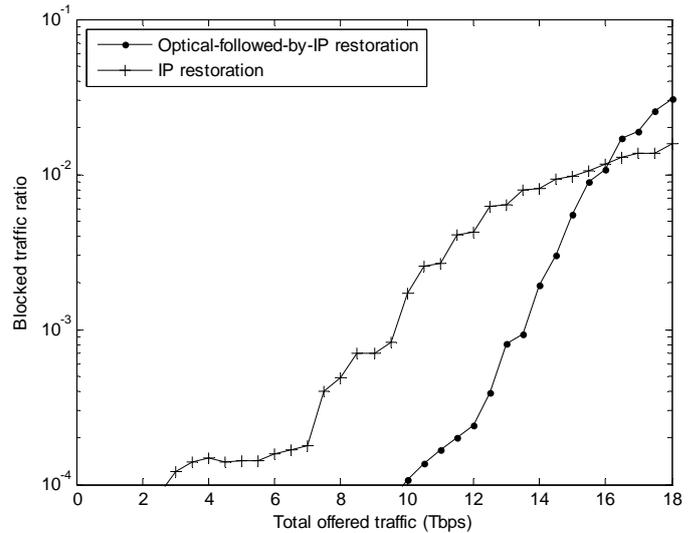


Fig. 3. Traffic blocking ratio in NSFNET network.

### D. Discussion and further work

Results in this paper are preliminary in different senses. On the one hand, we focus our cost model on the number of established lightpaths, but we neglect the effect of the optical architecture. For instance, the IP-R scheme can be implemented with simpler OXCs, since no degree of reconfigurability of the optical layer is required, that is, economic fixed-OADM can be installed into the nodes. In contrast, Op-IP-R requires the deployment of ROADMs, which are much more costly [19]. Even, the cost of state-of-the-art devices providing colorless, directionless and contentionless operation [20] might be prohibitive. A complete cost analysis, involving CAPEX and OPEX, is left as a further work.

In addition, survivability results show that a pure Op-IP-R with no over-dimensioning may fail to provide better performance than IP-R at high traffic loads. The reason behind this behavior is that IP-R pre-reserves many additional lightpaths, during the planning phase, to accommodate traffic from failing lightpaths. In its turn, Op-IP-R only establishes the minimum number of lightpaths required to carry all the traffic, without any protection degree. Even though algorithm 2 makes its best to re-route existing lightpaths, the surviving virtual topology may not be able to handle all the traffic, and it may be desirable to establish additional lightpaths for these situations. In a further work, we could introduce the concept of automatic switched optical networks (ASONS) [21], which have the ability of setting up/tearing down lightpaths on demand, as a cost-effective option to establish new lightpaths between highly congested node pair, tearing them down after reparation of the failed ones.

To conclude, another question to investigate is about IP routing. As stated in [22], typically IGP weights are traffic-engineered for optimal link utilization. When a link failure

occurs, affected traffic is re-routed on the surviving topology using the old weights, resulting in congestion episodes and traffic losses. Then, the approach should be to determine a new optimal set of weights to minimize such congestion during the failure, returning back to the original weight setting after reparation. While this idea works in theory, it is not implemented in practice for several reasons (i.e. computing time, IGP re-convergence time, and so on). Instead, we could try to extend our algorithm to determine ahead-of-time a robust weight system that works for both normal state as well as under different failure states, eliminating the need to determine new weights after each failure dynamically.

## VI. CONCLUSIONS

In this paper, we present a comparison between two recovery schemes or IP-over-WDM networks. Our results suggest that optical-followed-by-IP restoration (Op-IP-R) is a more economical option than IP restoration in terms of the number of required lighpaths, while still improves traffic survivability in most cases. Results suggest that a trade-off solution, applying Op-IP-R with minimal over-dimensioning would be the best future-proof alternative.

## REFERENCES

- [1] Cisco Systems Inc., "Cisco Visual Networking Index: Forecast and Methodology, 2013-2018," White Paper, June 2014.
- [2] O. Gerstel, "The Age of Multi-Layer Networking," in *Proceedings of the Asia Communications and Photonics Conference (ACP2013)*, Beijing (China), November 2013.
- [3] O. Gerstel, C. Filsfil, and W. Wakim, "IP-Optical Interaction during Traffic Restoration," in *Proceedings of the Optical Fiber Communication Conference and the National Fiber Optic Engineers Conference (OFC/NFOEC 2013)*, Anaheim (USA), March 2013.
- [4] P. Pavon-Marino and J.L. Izquierdo-Zaragoza, "Net2Plan: An open-source network planning tool for bridging the gap between academia and industry," *IEEE Network Magazine*, in press.
- [5] Net2Plan – The open-source network planner [Online]. Available: <http://www.net2plan.com/> [Last accessed: July 2014]
- [6] P. Demeester *et al.*, "Resilience in multilayer networks," *IEEE Communications Magazine*, vol. 37, no. 8, pp. 70-76, August 1999.
- [7] A. Fumagalli and L. Valcarenghi, "IP Restoration vs. WDM Protection: Is There an Optimal Choice?," *IEEE Network*, vol. 14, no. 6, pp. 34-41, November-December 2000.
- [8] D.D.J. Kan, A. Narula-Tam, E. Modiano, "Lightpath Routing and Capacity Assignment for Survivable IP-over-WDM Networks," in *Proceedings of the 2009 Workshop on Design of Reliable Communication Networks (DRCN2009)*, Washington (USA), October 2009.
- [9] T. Lin, Z. Zhou, K. Thulasiraman, "Logical topology survivability in IP-over-WDM networks: survivable lightpath routing for maximum logical topology capacity and minimum spare capacity requirements," in *Proceedings of the 2011 Workshop on Design of Reliable Communication Networks (DRCN2011)*, Krakow (Poland), October 2011.
- [10] C.S.K. Vadrevu and B. Mukherjee, "Survivable IP topology design with re-use of backup wavelength capacity," in *Proceedings of the 2009 International Conference on Advanced Networks and Telecommunication Systems (ANTS 2009)*, New Delhi (India), December 2009.
- [11] C.S.K. Vadrevu, M. Tornatore, R. Wang, and B. Mukherjee, "Integrated design for backup capacity sharing between IP and wavelength services in IP-over-WDM networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 4, no. 1, pp. 53-65, January 2012.
- [12] B. Jaumard, M. Bui, B. Mukherjee, and C.S.K. Vadrevu, "IP restoration vs. optical protection: Which one has the least bandwidth requirements?," *Optical Switching and Networking*, vol. 10, no. 3, pp. 261-273, July 2013.
- [13] L. Sahasrabudhe, S. Ramamurthy, and B. Mukherjee, "Fault management in IP-over-WDM networks: WDM protection versus IP restoration," *IEEE Journal of Selected Areas in Communications*, vol. 20, no. 1, pp. 21-33, January 2002.
- [14] R. Ramaswami, and K.N. Sivarajan, "Design of Logical Topologies for Wavelength-routed Optical Networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, pp. 840-851, June 1996.
- [15] M. Pickavet *et al.*, "Recovery in multilayer optical networks," *Journal of Lightwave Technology*, vol. 24, no. 1, pp. 122-134, January 2006.
- [16] J.P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery: Protection and Restoration, SONET-SDH, IP, and MPLS*, Morgan Kaufmann, 2004.
- [17] *Spectral grids for WDM applications: DWDM frequency grid*, ITU-T G.694.1 standard, Rev. 02/12, 2012.
- [18] Y. T'Joens, G. Ester, and M. Vandenhoute, "Resilient Optical and SONET/SDH-based IP networks," in *Proceedings of the 2nd International Workshop on the Design of Reliable Communication Networks (DRCN2000)*, Munich (Germany), April 2000.
- [19] F. Rambach *et al.*, "A Multilayer Cost Model for Metro/Core Networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 5, no. 3, pp. 210-225, March 2013.
- [20] S. Gringeri, B. Basch, V. Shukla, R. Egorov, and T.J. Xia, "Flexible architectures for optical transport nodes and networks," *IEEE Communications Magazine*, vol. 48, no. 7, pp. 40-50, August 2010.
- [21] A. Groebbens *et al.*, "Logical topology design for IP rerouting: ASONs versus static OTNs," *Photonic Network Communications*, vol. 21, no. 2, pp. 170-191, February 2011.
- [22] D. Medhi, "A Perspective on Network Restoration," in *Handbook of Optimization in Telecommunications*, Springer, 2006.