Joint Fault Tolerant and Latency-Aware Design of Multilayer Optical Networks

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

Abstract—**When designing a multilayer IP-over-WDM optical network, an important requirement is finding the most cost-effective design that is tolerant of a set of selected failures arbitrarily defined, so the 100% of the IP traffic survives for them. A second concern is the need to guarantee a maximum end-to-end latency for IP traffic (e.g. 50 ms in continental US). Unfortunately, fault tolerance and latency at the IP layer are performance merits extremely challenging to enforce in multilayer networks, since they depend on both the IP routing over the lightpaths and the lightpath routing over the fibers, and how they change during recovery processes. There is little research considering one of these concerns separately, and to the best of the authors' knowledge, just preliminary works on how to design the network jointly considering both requirements. Note that a joint optimization is needed, since e.g. backup routings that do not guarantee end-to-end latency may be useless. In this paper, we address this problem, and present a multilayer planning algorithm to dimension IP/OSPF-over-WDM networks with three main recovery schemes: IP-only restoration, 1+1 optical protection followed by IP restoration, and optical-followed-by-IP restoration. Results are provided to validate the algorithm, and feed a discussion on the joint fault tolerant and latency-aware multilayer network design under those schemes.**

Keywords—Net2Plan; multilayer network design; IP-over-WDM; OSPF; ECMP; latency-awareness; traffic survivability

I. INTRODUCTION

In the context of IP-over-WDM multilayer backbone networks, IP and WDM layers are often operated independently, by separate departments. On the one hand, IP departments optimize the routing over the topology of IP links, consisting of end-to-end lightpaths. Each lightpath is seen by the IP layer as a *pipe*, a direct connection between two routers. On the other hand, the optical layer is operated by different personnel, in charge of providing the requested lightpaths: (i) deciding on its route through the physical topology of fibers, and (ii) managing the transponder banks, optical add/drop multiplexers and the rest of optical equipment.

Nonetheless, the joint design of both layers in IP over WDM networks has shown its benefits in multiple previous studies [1], and is a force towards coordinating the departments operating the network [2]. This work is targeted to contribute in this process.

We focus on two particular situations in which a coordinated view of the network is both challenging and necessary. First, there is a need to create multilayer designs which are tolerant to some selected failure states (e.g. single link cut). In these cases, services at the IP layer (those observed by users) will survive depending on the reactions taken at both layers, and multiple non-trivial trade-offs appear. For instance, IP layer dimensioning depends on whether lightpaths are 1+1 protected at the optical layer or lightpaths are not optically recovered at all. Second situation is the need to guarantee a given maximum end-to-end latency at the IP level, as needed for live video or online-gaming services. Here, end-to-end latency of an IP flow depends on both of the sequence of lightpaths traversed (IP routing), and the total length of each lightpath defined by its traversed fibers (optical routing).

Interestingly, existing approaches do not focus on this kind of fault tolerant and end-to-end latency aware designs in multilayer networks [3].

The main contribution of this work is presenting a complete joint fault tolerant and latency-aware design algorithm for IP/OSPF-over-WDM networks. This contribution is itself threefold:

- OSPF-ECMP (Equal-Cost Multi-Path) routing is considered in the IP layer, as usual in current networks [4], which does not provide the full control over individual paths that MPLS permits, and requires a more challenging and careful design.
- A 100% survivability or fault tolerance is guaranteed for a set of user defined failure scenarios. The algorithm can satisfy this for three different multilayer recovery types, with different agility requirements at the optical layer: (i) IP-only restoration, for static optical layers where lightpaths are not recovered; (ii) 1+1 optical protection followed by IP restoration in the surviving topology, for static optical layers where each IP link is protected through two disjoint lightpaths; and (iii) optical-followed-by-IP restoration, or multilayer restoration.
- A maximum end-to-end latency value at the service (IP) layer is guaranteed by design for any pair of nodes *under any of the considered network failure states*. The end-to-end latency is computed

This work was partially supported by the FPU fellowship program of the Spanish Ministry of Education, Culture and Sports (ref. no. FPU14/04227) and by the Spanish project grant TEC2014-53071-C3-1-P (ONOFRE) and TEC2015-71932-REDT (ELASTIC).

considering just propagation delays, which are a major source of delay in backbone networks, where the high bit rates make queuing delays negligible (unless congestion occurs) respect to propagation times.

The algorithm will be tested and validated for different reference topologies and traffic matrices using the open-source Net2Plan framework [5,6]. This will emphasize one of the use cases of the algorithm: the assessment of the network throughput or cost for different recovery strategies. Also, we will demonstrate the importance of latency constraints to eliminate the appearance of IP flows with a significantly high end-to-end latency across different network states.

The rest of the paper is organized as follows. In Section 2, we review some previous works in multilayer network planning algorithms considering survivability and latency requirements. In Section 3, we describe our novel multilayer planning algorithm considering three different recovery methods. In Section 4, we report and discuss the results of our case study. Finally, Section 5 concludes the paper.

II. Related Work

A. *Survivability-Aware Design in Multilayer Networks*

There is a limited number of works in the literature proposing multilayer design algorithms adapted to provide fault tolerance under different recovery schemes.

The most straightforward design scenario of IP/OSPF-over-WDM networks is the one in which there is no optical recovery mechanism, so that the IP layer is responsible for traffic rerouting and enough capacity must be provided beforehand. A brief survey on related dimensioning techniques is presented in [4]. The main difference between the discussed approaches is the level of integration of the components of the survivable multilayer network design problem: design and dimensioning of the IP topology, traffic routing, lightpath routing and failure analysis. At the end, authors conclude that the right approach is to consider the previous factors altogether to maximize benefits from multilayer networking.

Alternatively, authors in [7] discuss the case in which multilayer restoration is applied. In the studied strategy, recovery starts from the optical layer and escalates towards the IP layer. Therefore, blocked traffic in the lower layer may still be recovered through IP rerouting. Then, multilayer restoration is compared with IP-only restoration, showing that the former is a cost-efficient alternative. Coordination between layers is mandatory in order to ensure the correct process of restoration. In this case, hold-off timers are defined so that restoration mechanisms at the IP layer are not triggered until a certain time has passed since the failure, giving time to the optical layer to react (applying its recovery scheme or just waiting failures to be repaired).

By contrast, [8] provides a different vision by comparing dedicated path protection in the optical layer with multilayer restoration. Again, multilayer restoration leads to significant cost savings.

B. *Latency-Aware Design in Multilayer Networks*

Although end-to-end latency is an important QoS (Quality of Service) parameter into SLA (Service Level Agreement) contracts, it has not been specifically addressed in multilayer networks, or at least in a non-best-effort manner. The most common approach is to route over shortest paths across all layers, and then dimensioning the links accordingly [4]. On the one hand, pure IP routing is based on OSPF with enhanced ECMP splitting rules, which by definition uses shortest path algorithms. On the other hand, lightpath routing in the optical layer often applies classical shortest path algorithms [9] like Dijkstra's algorithm [10] or related variants for node/link-disjoint dedicated path protection [11].

The fact is that, even from the perspective of single-layer networks, enforcing end-to-end latency guarantees at the IP layer is quite challenging, since OSPF-ECMP routing requires a careful tuning of link weights in order to achieve SLA and/or TE (Traffic Engineering) objectives [12]. Unfortunately, this problem itself is not amenable for integer linear programming (ILP) approaches, besides that path-based constraints like end-to-end latency cannot be easily introduced into ILPs for hop-by-hop routing schemes like OSPF-ECMP. In addition, heuristic-based approaches for weight tuning cannot be always able to fulfill QoS/SLA and/or TE requirements unless objective function considers them explicitly [13]. Nonetheless, we would like to remark that operators usually prefer to assign naïve link weight assignment strategies, i.e. using hop-based or distance-based approaches [4], rather than perform weight tuning according to network conditions.

Consequently, most research in latency-aware design for IP networks considers MPLS routing. In this case, paths can be explicitly defined for each IP demand according to SLA and TE objectives [14]. For example, authors in [15] present a single layer network design based on explicit routing to enforce equalization of end-to-end latency between different IP demands, whereas keeping a maximum end-to-end latency value.

C. *Joint Fault Tolerant and Latency-Aware Multilayer Design*

Although fault tolerance and end-to-end latency are important QoS/SLA parameters naturally coupled, there are few relevant works in the literature considering both aspects.

Focusing on single-layer networks, we find that few works have proposed solutions for OSPF-ECMP variants, due to the aforementioned weight tuning problem. In [16], authors present a heuristic algorithm to determine the optimal set of link weights so that the network survives to a set of predefined failures. Latency constraints are not defined by default, but authors give guidelines to include them.

Regarding to multilayer networks, the need for joint consideration of survivability and latency constraints has been discussed in the literature [3], but it has received few attention by the community.

To the best of our knowledge, the only proposal jointly considering both was presented in [6]. Authors present a multilayer algorithm able to guarantee 100% survivability under any bidirectional fiber failure and maximum end-to-end latency, when OSPF-ECMP is applied in the IP layer, and 1+1 protection is applied in the optical layer using two link-disjoint lightpaths for each IP link. A modified version of the classical IGP-WO algorithm [12] is considered to find the set of link weights that minimizes the network congestion. The objective function is tuned as proposed in [15] to penalize solutions violating latency constraints.

In this paper, we significantly extend the previous work to consider additional multilayer recovery options, whereas link weights are left untouched. Consequently, the new algorithm can be applied by operators to plan their networks according to their preferences and, for the first time, taking into account survivability and latency constraints together, in a much broader set of situations.

Finally, it is worth mentioning the absence of proposals for hybrid OSPF-MPLS routing in multilayer networks to deal with this problem. Hybrid routing schemes have received some attention in the past, for either unprotected [17] or protected [18] designs in pure IP networks. Essentially, different IP routing schemes are applied to different traffic profiles: high-priority traffic with stringent QoS/SLA requirements can be appropriately routed using fine-grained MPLS paths, whereas the rest of the traffic is routed using OSPF. A candidate extension for multilayer networks would imply MPLS routing at the IP layer plus 1+1 protection at the optical layer for high-priority traffic, routing the remaining traffic with OSPF and undergoing (multilayer) restoration.

III. Algorithm Description

In this section, we describe the planning algorithm for dimensioning multilayer IP-over-WDM networks taking into account three different recovery mechanisms: 1+1 optical protection followed by IP restoration, optical-followed-by-IP restoration (or multilayer restoration) and IP-only restoration. As aforementioned, IP routing employs OSPF-ECMP according to related standards, with weights equal to one. At the optical layer no wavelength conversion is allowed, and no intermediate regeneration is required.

The algorithm ensures a maximum end-to-end latency for every IP traffic demand according to an input parameter, as well as survivability to a set of user defined failure scenarios modeled by shared-risk groups (SRGs), also introduced as an input. Note that, by defining no faults (SRGs) to be tolerant to, or a very high end-to-end latency requirement, the algorithm can be used to find designs only latency-aware or only survivability-aware, respectively.

A. Implementation

It is well-known that multilayer network design problem is *NP*-complete, thus a heuristic algorithm is implemented. In our case, we decided to develop an algorithm based on the Greedy Randomized Adaptive Search (GRASP) [19] meta-heuristic. GRASP is an iterative meta-heuristic composed of two phases per iteration: construction and local search. In the former, a feasible random solution is constructed using a greedy algorithm. A solution is said to be feasible if satisfies all the problem constraints, which are 100% survivability as well as meeting maximum end-to-end latency in our case. If a feasible solution cannot be found, the algorithm moves to the next iteration. Otherwise, new feasible solutions improving the objective function are explored using a local search algorithm. Finally, the best historical solution after all GRASP iterations is returned.

Below, we provide a high-level description of the algorithm. Due to the space limitation we do not show all the implementation details, but the source code is publicly available in [5].

The core of the algorithm is based on an IP-over-WDM module used to apply RWA (Routing and Wavelength Assignment) and recovery mechanisms, as it should work in a real-world deployment. Specifically, this module reacts to events generated by the main algorithm performing the corresponding actions of the defined behavior. For example, each time a new lightpath must be added, the IP-over-WDM module receives an event 'add-lightpath' and applies the RWA mechanism accordingly. Following this methodology, we are able to generate algorithms for offline dimensioning by reusing existing online algorithms for dynamic lightpath establishment or multilayer recovery, as depicted in Fig. 1.

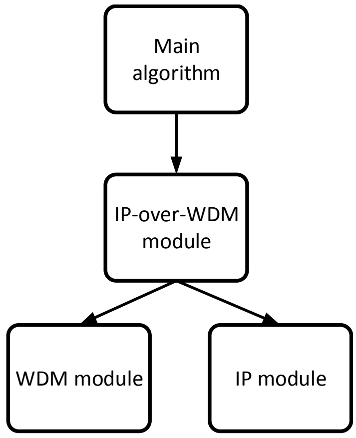

Fig. 1. Algorithm modules hierarchy.

The pseudo-code of the construction phase is presented below:

Step 1: Compute all blocked IP traffic, oversubscribed IP links and maximum end-to-end latency. If no constraint is violated in any predefined failure state, end.

Step 2: Randomly select an ingress-egress pair of nodes among those violating some constraint.

Step 3: Establish a new IP link (lightpath) between the selected nodes. Then go back to Step 1.

Fig. 2. Pseudo-code for construction phase.

We would like to remark that during Step 1, all the predefined failure states are evaluated to check whether survivability and latency constraints are fulfilled. Essentially, we iterate over all failure states to emulate the behavior of the recovery schemes using the IP-over-WDM module. Since in our case each failure state is modeled by an SRG, we send an event 'SRG-failed' to that module to execute the corresponding actions. Then, we check the resulting state and revert back to normal operation.

In Step 3, the RWA policy is used to allocate the lightpath. In our case, we choose a fixed-alternate routing strategy along with first-fit wavelength allocation. In other words, we use a set of precomputed candidate paths and try to find the shortest path selecting the common free wavelength of lower index. During this step, the algorithm might fail to allocate new lightpaths due to RWA constraints, which indicates inability to produce a feasible design for the given network configuration and traffic matrix.

Finally, the local search phase works as follows:

Step 1: Randomly select an ingress-egress pair of nodes with IP links between them.

Step 2: Compute all blocked IP traffic, oversubscribed IP links and maximum end-to-end latency considering all IP links between those nodes as removed. If no constraint is violated, remove permanently those IP links and go to Step 1.

Step 3: Repeat Steps 1 and 2 until all node pairs have been visited and no additional IP link bundle can be removed.

Step 4: Randomly select an ingress-egress node pair of nodes with IP links between them.

Step 5: Try to remove one IP link while guaranteeing no constraint violation. If an IP link is removed, then go to Step 1.

Step 6: Repeat Steps 4 and 5 until all node pairs have been visited and no additional IP link can be removed. Then, end.

Fig. 3. Pseudo-code for local search phase.

It is important to highlight that during the second phase, we are only trying to improve a feasible solution coming from the first stage. If the solution cannot be improved during stage 2, the solution of the corresponding iteration will be the one from stage 1.

B. IP-over-WDM Module and Recovery Schemes

Our IP-over-WDM module is developed to support the three considered recovery mechanisms. We use as reference a bottom-up multilayer approach [9] as applied in [7]. Upon any event, the general module invokes first the WDM module, and then the IP module. Such modules are able to determine whether the received event must be processed, and take the appropriate actions. Specifically, the IP module is common for all recovery mechanisms as implements the OSPF rerouting process according to the surviving IP topology. Differences come from the WDM module.

First, for 1+1 optical protection followed by IP restoration, the WDM module realizes each IP link with two SRG-disjoint lightpaths. Hence, upon failure, it is able to maintain traffic survivability.

Second, for IP-only restoration, the WDM module is only in charge of applying the RWA mechanism to allocate new IP links as lightpaths in the physical topology.

Finally, for multilayer restoration, the WDM module routes new lightpaths as for IP-only restoration but, upon failure events, tries to reroute a lightpath before triggering any IP layer alarm.

C. Latency Computation

Regarding to the latency constraint, it is worth noting that end-to-end latency comprises several metrics such as propagation delay, transmission delay, queuing delay or processing times. To simplify our study, we will consider that latency is dominated by propagation delay, as usual in backbone networks [20].

An efficient computation of end-to-end propagation delay in multilayer networks with OSPF-ECMP in the IP layer is not trivial. We use the path reconstruction method detailed in [21] to convert hop-by-hop routing tables to explicit paths. Then, for each IP demand, we compute end-to-end latency as the worst-case end-to-end propagation delay among all paths carrying traffic for that demand.

IV. Case Study

In this section, we collect a set of testing results obtained for exploring the trade-off between network cost, total throughput and maximum end-to-end latency appearing in the choice among a network operated with the recovery schemes described in Section III.B.

Results have been obtained using the offline network design tool of Net2Plan. This tool assists the user in the process of offline network design and planning, and related data analysis procedures. Inputs to this tool are topology and traffic information, algorithm and algorithm parameters. The algorithm is loaded as a Java class implementing the appropriate interfaces defined in the documentation. We would like to remark that Net2Plan and the algorithm implemented in this paper are publicly available for inspection and validation on the website [5].

A. Testing Scenario

Our testing scenario includes two IP-over-WDM networks composed of a set of fiber links and networks sites following the well-known NSFNet and Internet2 topologies, as well as reference traffic matrices for them. Each node is equipped with IP routing equipment and a ROADM. We assume 40 WDM channels per fiber and each lightpath appears as a single link in the IP layer with a capacity of 40 Gbps. We also assume that no signal regeneration between node pairs is needed and the propagation speed of light is 200,000 km/s. All topology and traffic information is stored in the Net2Plan repository, accessible in [5].

We run the algorithm for different total offered traffic values, from 250 to 4000 Gbps in steps of 250 Gbps, using scaled versions of the reference traffic matrices. Maximum end-to-end latency limit is set to 85 ms for the NSFNet network and 50 ms for Internet2.

Each SRG implies the failure of two fibers (one in each direction), considered as a cut in the ducts carrying those fibers. Those cuts may be caused by human errors (civil engineering activities), natural disasters or programmed maintenance. It is worth noting that no Mean Time To Repair (MTTR) or Mean Time To Fail (MTTF) is assigned to the SRGs since the algorithm is designed to guarantee the survivability of the network against all possible single SRG failures, irrespectively of the unavailability time. This is not a loss of generality respect to multiple failures. For instance, to test multiple simultaneous duct cuts with our algorithm, we can just define new SRGs where the failing resources for them are the combined ones of the multiple cuts.

After all the executions have finished we collect the total number of transponders necessary to establish the planned design for each network and recovery scheme. This is tightly related to the number of IP ports and thus it is a good indicator of the overall network cost [22]. It is computed as the sum of the maximum between incoming and outgoing lightpaths for each node, considering bidirectional transponders with independent tunable transmitter and receiver [23]. We also store the worst-case end-to-end latency for each IP traffic demand under any possible network state, considering single SRG failure.

B. Results

Figs. 4 and 5 show the total number of transponders to be installed across the network to allocate all the offered traffic and provide 100% survivability and maximum end-to-end latency guarantees for both topologies. These figures reveal several interesting findings.

First, given a certain total offered traffic value, transponders' requirements are ranked from high to low as follows: 1+1 optical protection followed by IP restoration (1+1), IP-only restoration (IP-R) and optical-followed-by-IP restoration (Op-IP-R). As expected, optical restoration is able to recover failing lightpaths by rerouting over the surviving topology and, consequently, some failures become unnoticeable to the IP topology at the cost of moderated over-dimensioning. On the other hand, additional lightpaths must be provisioned prior to network operation in the case of IP-only restoration to ensure traffic survivability, since failing lightpaths cannot be recovered until physical resources go back to operational state. Finally, 1+1 optical protection requires the largest number of transponders as each IP link is realized through two physical disjoint lightpaths and requires two pairs of transponders.

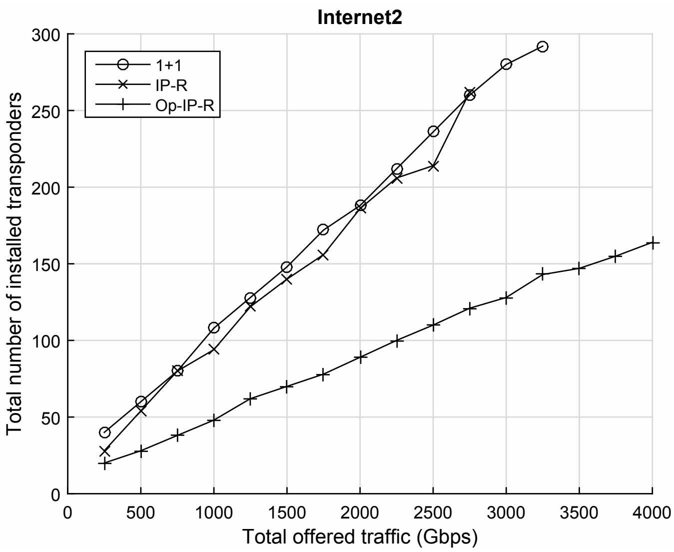

Fig 4. Total number of transponders for Internet2 network.

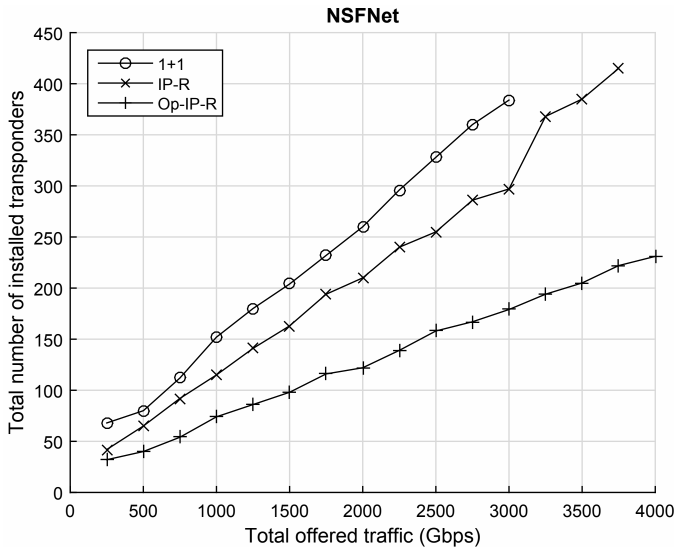

Fig 5. Total number of transponders for NSFNet network.

Second, we observe that for a given number of transponders the most efficient variant is optical-followed-by-IP restoration since it allows to reach higher network throughput or, in other words, it allows to introduce more traffic in the network before requiring a major upgrade (i.e. to install more transponders or to deploy more fibers).

Figs. 6 and 7 illustrate the histogram for worst-case latency distribution for each scenario. In the upper side of each figure, we show the latency distribution in some tests conducted without considering the latency constraints. In the lower side, the equivalent distribution considering such constraint is presented. Vertical dashed lines are used to show the value of the constraint. We observe that for latency-aware scenarios the maximum latency limit is correctly satisfied, and that the end-to-end latency distribution is more concentrated than for the latency-unaware scenarios. In contrast, when the latency limit is not considered, end-to-end latencies at the IP level is anyhow satisfied in many flows, but reach significantly high values in a relevant amount of other flows. This supports and motivates the importance of a joint fault tolerant and latency-aware design.

C. Discussion and Further Work

Results show that optical-followed-by-IP restoration is a better option than 1+1 optical protection or over-dimensioning for IP-only restoration. However, we assume that optical restoration, if possible, takes place in a negligible period of time. Unfortunately, at this moment, lightpath setup or reconfiguration is typically slow by nature due to channel equalization processes, even in the order of minutes [1].

As a further work, it would be worth investigating the performance of our algorithm for other RWA schemes, not only for fixed-alternate routing. Specifically, it would be interesting to include IPFRR (IP Fast Rerouting) considerations for 1+1 protection and IP-only restoration. Maximizing IPFRR coverage, as proposed in [24], will guarantee that a large majority of the traffic is restored at a sub-second timescale.

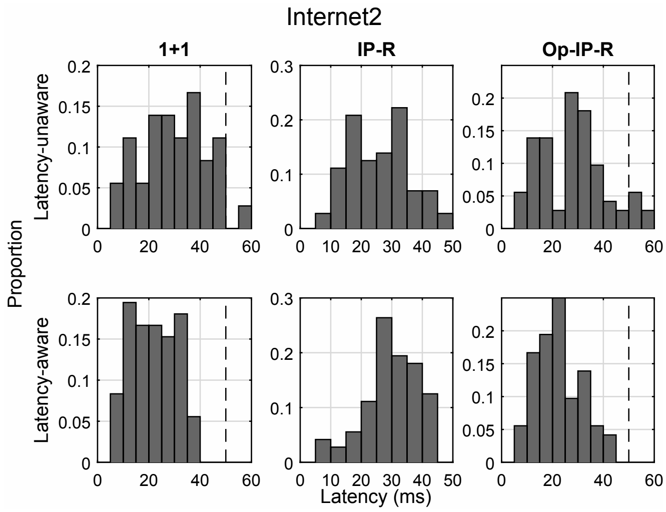

Fig 6. Worst-case end-to-end latency distribution for Internet2 network.

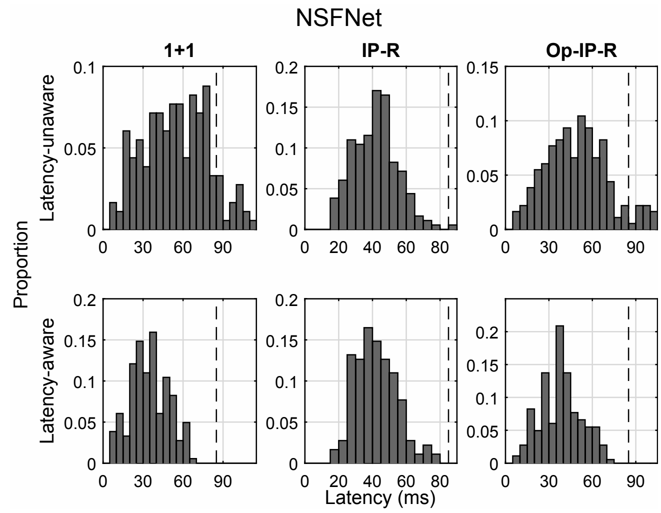

Fig 7. Worst-case end-to-end latency distribution for NSFNet network.

V. CONCLUSIONS

In this paper, we present an algorithm for multilayer IP-over-WDM network planning taking into account survivability and maximum end-to-end latency for IP traffic. The algorithm can consider three different multilayer survivability schemes in the design. We note that the algorithm can cover also the case of multilayer latency-aware or fault tolerant cases separately, by just not defining SRGs, or setting very high latency limits, respectively. Results shown validate our proposal, and support our view that the joint consideration of faults and latencies is needed, e.g. to avoid designs where end-to-end latency for some flows becomes excessive under failures.

REFERENCES

[1] O. Gerstel *et al.*, "Multi-Layer Capacity Planning for IP-Optical Networks," *IEEE Communications Magazine*, vol. 52, no. 1, pp. 44–51, Jan. 2014.

[2] J.-P. Fernandez-Palacios *et al.*, "IP and Optical Convergence: Use Cases and Technical Requirements," White Paper, Jan. 2014.

[3] O. Gerstel, C. Filsfils, and W. Wakim, "IP-Optical Interaction during Traffic Restoration," in *Proceedings of the Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference 2013* (OFC/NFOEC 2013), Anaheim (United States), 2013.

[4] E. Palkopoulou *et al.*, "Impact of IP Layer Routing Policy on Multilayer Design [Invited]," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 7, no. 3, pp. A396–A402, Mar. 2015.

[5] Net2Plan – The open-source network planner. [Online] Available: http://www.net2plan.com [Last accessed: April 2016]

[6] 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*, vol. 29, no. 5, pp. 90–96, Sep. 2015.

[7] J.-L. Izquierdo-Zaragoza and P. Pavon-Marino, "Assessing IP vs optical restoration using the open-source Net2Plan tool," in *Proceedings of the 16th International Telecommunications Network Strategy and Planning Symposium* (NETWORKS 2014), Funchal (Portugal), 2014.

[8] A. Mayoral *et al.*, "Minimizing Resource Protection in IP Over WDM Networks: Multi-layer Shared Backup Router [Invited]," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 7, no. 3, pp. A440–A446, Mar. 2015.

[9] J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery: Protection and Restoration, SONET-SDH, IP, and MPLS*, 1st ed. Morgan Kaufmann, 2004.

[10] E. W. Dijkstra, "A Note on Two Problems in Connexion with Graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, Dec. 1959.

[11] J. W. Suurballe and R. E. Tarjan, "A Quick Method for Finding Shortest Pairs of Disjoint Paths," *Networks*, vol. 14, no. 2, pp. 325–336, 1984.

[12] B. Fortz and M. Thorup, "Internet Traffic Engineering by Optimizing OSPF Weights," in *Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies* (INFOCOM 2000), Tel Aviv (Israel), 2000.

[13] A. Nucci and K. Papagiannaki, *Design, Measurement and Management of Large-Scale IP Networks: Bridging the Gap between Theory and Practice*, 1st ed. Cambridge University Press, 2008.

[14] V. Foteinos, K. Tsagkaris, P. Peloso, L. Ciavaglia, and P. Demestichas, "Operator-Friendly Traffic Engineering in IP/MPLS Core Networks," *IEEE Transactions on Network and Service Management*, vol. 11, no. 3, pp. 333–349, Sep. 2014.

[15] M. Yu, M. Thottan, and L. Li, "Latency Equalization as a New Network Service Primitive," *IEEE/ACM Transactions on Networking*, vol. 20, no. 1, pp. 125–138, Feb. 2012.

[16] L. S. Buriol, M. G. C. Resende, and M. Thorup, "Survivable IP network design with OSPF routing," *Networks*, vol. 49, no. 1, pp. 51–64, Jan. 2007.

[17] M. Zhang, B. Liu, and B. Zhang, "Multi-Commodity Flow Traffic Engineering with Hybrid MPLS/OSPF Routing," in *Proceedings of the IEEE Global Telecommunications Conference 2009* (GLOBECOM 2009), Honolulu (United States), 2009.

[18] A. Mereu, D. Cherubini, A. Fanni, and A. Frangioni, "Primary and Backup Paths Optimal Design for Traffic Engineering in Hybrid IGP/MPLS Networks," in *Proceedings of the 7th International Workshop on Design of Reliable Communication Networks* (DRCN 2009), Washington D.C. (United States), 2009.

[19] T. A. Feo and M. G. C. Resende, "Greedy Randomized Adaptive Search Procedures," *Journal of Global Optimization*, vol. 6, no. 2, pp. 109–133, Mar. 1995.

[20] L. Kleinrock, "The Latency/Dandwidth Tradeoff in Gigabit Networks," *IEEE Communications Magazine*, vol. 30, no. 4, pp. 36–40, Apr. 1992.

[21] P. Pavon-Marino, *Optimization of Computer Networks: Modelling and Algorithms – A Hands-On Approach*, 1st ed. Wiley, 2016.

[22] R. Hülsermann, M. Gunkel, C. Meusburger, and D. A. Schupke, "Cost modeling and evaluation of capital expenditures in optical multilayer networks," *Journal of Optical Networking*, vol. 7, no. 9, p. 814-833, Sep. 2008.

[23] J. M. Simmons, *Optical Network Design and Planning*, 2nd ed. Springer, 2014.

[24] J.-L. Izquierdo-Zaragoza, J.-J. Pedreno-Manresa, and P. Pavon-Marino, "Maximizing IP Fast Rerouting Coverage in Survivable IP-over-WDM Networks," in *Proceedings of the 41st European Conference on Optical Communication* (ECOC 2015), Valencia (Spain), 2015.