Impact of IP Layer Routing Policy on Multi-Layer Design

E. Palkopoulou[1], O. Gerstel[1], I. Stiakogiannakis[2], T. Telkamp[1], V. López[3], I. Tomkos[4]

[1]*Cisco, Email: epalkopo@cisco.com*
[2]*Institute of Computer Science, Foundation for Research and Technology, Greece*
[3]*Telefonica I+D, Global CTO Unit, Madrid Spain*
[4] *Athens Information Technology Center, Peania, Greece*

Abstract: We evaluate the impact of the IP layer routing policy (Hop-Based and Distance-Based) on the cost and latency of a multi-layer network design. We find that the optical network's regeneration requirements affect the optimal policy.

OCIS codes: (060.0060) Fiber optics and optical communications; (060.4256) Networks, network optimization

1. Introduction

The behavior of the IP layer is often absent from papers on IP-Optical network planning. This leads to an inaccurate assessment of the impact of a certain IP layer design over an optical topology – affecting, for example, studies of router bypass. This paper aims at bringing out the important characteristics of the IP layer that impact such multi-layer planning. The increasing pressure to improve the cost efficiency of networks has led to the development of advanced traffic engineering (TE) and network planning techniques. MPLS TE enables the explicit routing and arbitrary splitting of traffic, which are both desirable characteristics for optimization purposes. However, scalability and robustness may become issues in MPLS TE [1]. TE can also be achieved with pure IP networks. In such IP-based TE networks, the link weights of interior gateway protocols (IGPs) are set in order to meet a given objective [2]. IP TE provides better scalability – but less control over the selection of the individual paths. This in turn requires for the IGP weights to be carefully tweaked in order to match the new conditions. However, altering the IGP weights of IP links is highly undesirable by network operators. As a result, many network operators prefer to assign link weights based on simple rules that are unaffected by the network conditions.

In this work we examine two IGP weight setting policies that are agnostic to the network conditions: the Hop-Based Policy and the Distance-Based Policy. In order to provide a fair comparison, we perform multi-layer planning in order to optimize the IP topology design for each policy - as we find that different IGP policies may lead to different optimized IP topologies for the same traffic and optical layer topology. We show that the Hop-Based Policy is more cost-efficient than the Distance-Based Policy, for networks that do not require significant regeneration. However, these cost savings come at the expense of increased end-to-end latency. For networks with high regeneration requirements the Distance-Based Policy is more cost efficient and yields lower latency. To the best of our knowledge, this is the first time that such trade-offs are examined. Furthermore, we evaluate the impact of Equal Cost Multi-path (ECMP) on the two examined IGP policies. The ECMP mechanism allows for a demand to split among multiple sub-routes with equal IGP link weights. Traffic is evenly split to the next hop routers on these paths – again influencing the IP-optical design.

2. Multi-Layer Network Planning and Optimization: Exploring IGP Policies

We apply a multi-layer network planning and optimization methodology that jointly considers the IP and optical layer cost and includes failure analysis of optical links in the network design process. The optimization objective is set to the minimization of the combined IP layer and the optical cost. The input to the network planning procedure includes the traffic demand, the optical topology, the network equipment characteristics, the failure states under which survivability should be provided, and the IP routing policy. The output of the planning approaches includes the IP topology, the required network equipment in different network layers (e.g. IP ports, transponders, regenerators), as well as the total network cost.

In the failure analysis phase, the considered failure states are sequentially examined and the exact behavior of the IP layer is modeled, resulting in finding the worst case traffic on each IP link. As a result, the IP and optical circuits are dimensioned according to the expected worst case traffic. The resulting total network costs are used to evaluate the multi-layer network design. Thus, the impact of failures is not only considered for the dimensioning of the IP links, but also in the design of the IP topology itself (i.e., which adjacencies are established in the IP layer). Note that we adopt here today's approach to failure mitigation, in which optical link failures are handled by the IP layer (i.e., no optical restoration).

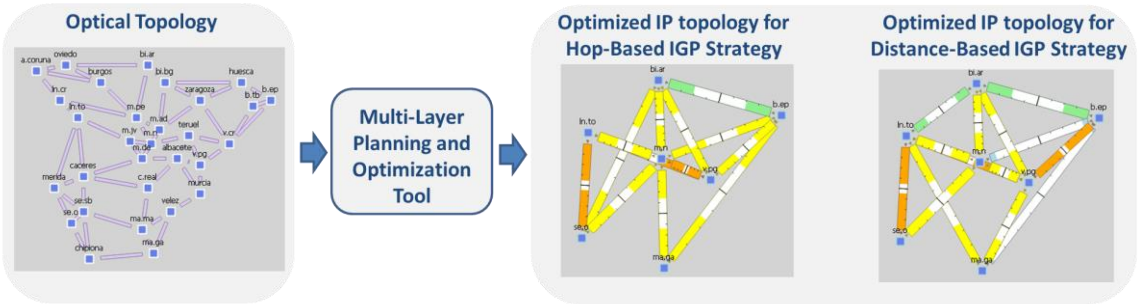

Fig. 1. Different IGP routing policies may lead to different optimized IP layer topologies.

We focus on two IGP routing policies:

- **Hop-Based Policy:** In this policy the IGP metric of all IP links is set to a constant value. For example, all IGP metrics are set to 10.
- **Distance-Based Policy:** In this policy the IGP metric of an IP link is set to be equal to the physical distance on the optical circuit that it corresponds to. For example, if an IP link is carried over an optical circuit with a length of 432 km, then its IGP metric is set to 432.

We apply the described multi-layer methodology to a Spanish national backbone network [3]. As shown in Fig.1, we find that different IGP routing policies may lead to different optimized IP layer topologies. Note that the locations of the IP core routers are provided as input. However, the established IP adjacencies are subject to optimization.

3. Case Study Results

In the following, case study results on the performance of different IGP policies are discussed. In order to evaluate the impact of the traffic demand volume, the corresponding traffic matrix for reference year 2012 is uniformly scaled by a set of traffic demand multipliers. Note that a traffic demand multiplier of 1 corresponds to the traffic demand of the reference year 2012. The optimization objective is set to the minimization of the total network costs - which include IP ports, transponders, and regenerators. The starting point for the optimization is Telefonica IP topology. A capacity of 200 Gb/s (corresponding to the commercial state of the art for 16QAM modulation) is set for the IP ports, the transponders, and the regenerators. The maximum transparent reach of the transponders is set to 800 km. Regenerators are assumed to incur a cost overhead of 80% compared to the transponders. Transponders are assumed to incur half the IP ports costs (again, corresponding to current commercial price points). All single optical link failures are considered in the failure analysis.

Fig. 2 presents the relative performance of the Hop-Based Policy over the Distance-Based Policy in terms of latency and cost as a function of the traffic demand volume. Both average and maximum values are presented with respect to latency (averaging and maximization are applied to the different paths that all demands are routed over). Let us first focus on Fig. 2(a), which corresponds to the presented Spanish national backbone network, and examine the case of traffic demand multiplier 1. We find that costs savings of 16% can be achieved by adopting the Hop-Based Policy as opposed to the Distance-Based Policy. However, these savings come at the expense of 11% higher average end-to-end latency and 17% higher maximum latency compared to the Distance-Based Policy. For networks that span a relatively small geographical area, we observe that the Hop-Based Policy incurs less costs than the Distance-Based Policy (on average 11% cost savings) – but leads to higher latency. While the penalty in terms of increased latency may be expected – as the Distance-Based Policy forces demands to be routed over the path with the shortest physical distance – the gain in terms of cost savings is not as straightforward. In the following we analyze this behavior.

First of all, the Hop-Based Policy captures quite well the current network cost model for networks with low regeneration requirements - as the main contributor to overall network costs are IP ports and transponders that are deployed at the end-points of each IP link. So a demand that is routed on the path with the minimum number of network hops actually consumes the minimum portion of such costly resources. However, this is only part of the picture. Fig. 3(a) presents the IP link utilization of the different IGP policies as a function of the traffic demand volume. We find that the maximum IP link utilization in both policies is approximately 90% - which is the upper limit of the IP link utilization given as input to the multi-layer planning and optimization tool. What is interesting to observe is that the average IP link utilization is significantly lower in the Distance-Based Policy. Thus, better load balancing is achieved for the Hop-Based Policy. We find that this is greatly affected by ECMP. Fig. 3(b) shows the impact of ECMP on different IGP policies as a function of the traffic demand volume. In order to quantify this impact we examine the ratio of the total network cost with ECMP set to be inactive over the total cost with ECMP set to be active. Note that when ECMP is set to be inactive, the maximum number of allowed ECMP paths is set to

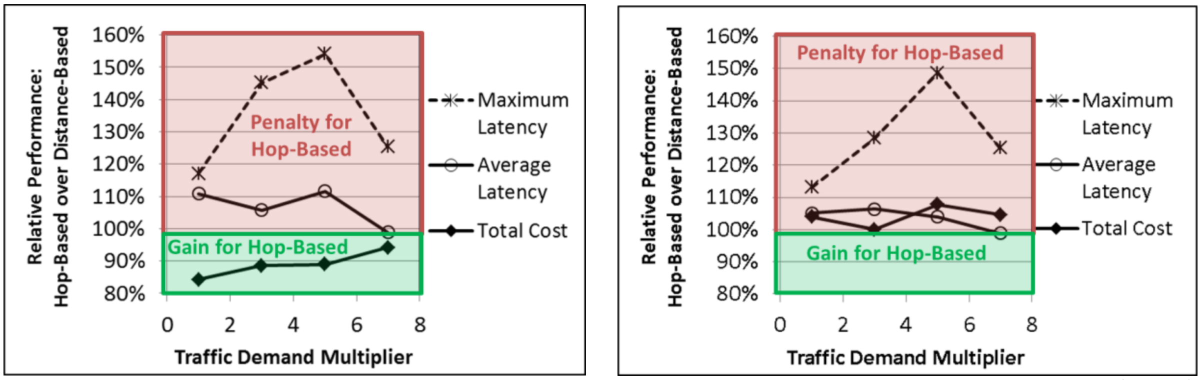

(a) Spanish national backbone network (b) Scaled Spanish national backbone network by a factor of 6

Fig. 2. Relative performance of Hop-Based over Distance-Based IGP policies in terms of latency and cost

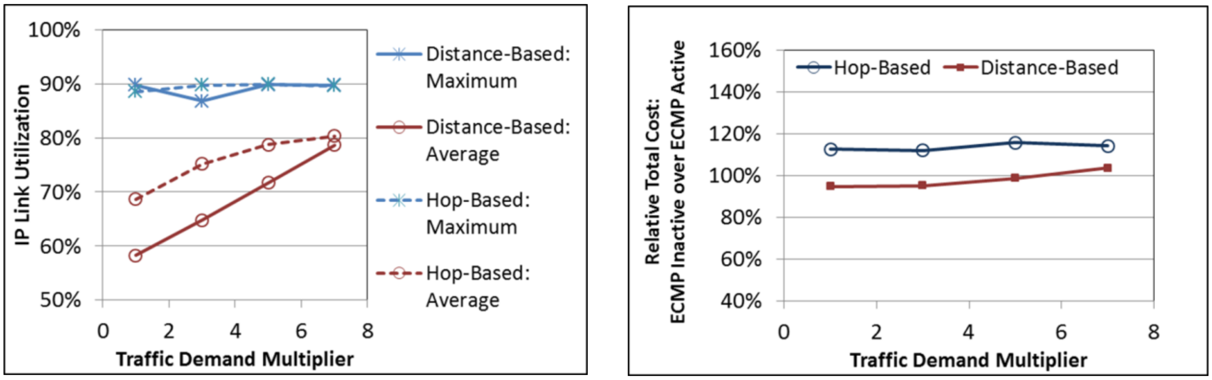

(a) (b)

Fig. 3. Spanish national backbone network: (a) The IP link utilization of different IGP policies; (b) The total network cost with set to be ECMP inactive (i.e., the maximum number of ECMP paths is set to 1) over the total cost with ECMP set to be active (i.e., no restrictions are imposed on the maximum number of ECMP paths) for different IGP policies.

1. When ECMP is set to be active, no restrictions are imposed on the maximum number of ECMP paths. We find that on average 14% cost savings can be achieved when no restrictions are imposed on the maximum number of allowed ECMP paths for the Hop-Based Policy. On the contrary, no such savings are achieved for the Distance-Based Policy (on average a 2% penalty in terms of cost is incurred). This is affected by the nature of the equal cost paths in the Distance-Based Policy. When two paths have the same metric (corresponding to paths with the same physical length) but a different number of hops, there will be no incentive for the path with fewer hops to be used. Thus, the traffic is split over paths with actually unequal network costs. Furthermore, the probability of paths with equal costs to arise in the Hop-Based Policy is greater.

Let us now focus on Fig. 2(b), which corresponds to the Spanish national backbone network scaled by a factor of 6. We find that the Distance-Based Policy is more cost efficient (up to 8% savings) for geographically large networks with high regeneration needs – without sacrificing on latency. In this case the Distance-Based Policy captures better the network cost model.

4. Conclusion

We evaluated the impact of Hop-Based and Distance-Based IGP routing policies in terms of the network cost and latency. Case study results on a Spanish backbone network indicate that the Hop-Based Policy incurs less costs than the Distance-Based Policy (up to 16% cost savings) – but leads to penalties in terms of end-to-end latency (up to 12% with respect to average latency). For networks with significant regeneration requirements the Distance-Based policy yields both lower costs and latency. We find that ECMP, which splits demands among multiple sub-routes with equal IGP metrics, allows significant cost savings to be achieved only for the Hop-Based policy.

This work is partially supported by the CISCO University Research Program and the EC FP7 project n°317999 (IDEALIST)

References

[1] N.Wang et al., "An overview of routing optimization for internet traffic engineering," IEEE Com. Surveys & Tutorials, pp.36-56, 2008.
[2] B. Fortz and M. Thorup, "Robust Optimization of OSPF/IS-IS Weights," In Proc. INOC, 2003.
[3] S. Martinez et al., “Assessing the Performance of Multi-Layer Path Computation Algorithms for different PCE Architectures,” OFC/NFOEC 2013.