

# Policy-based Grooming in Optical Networks

Fábio L. Verdi<sup>1</sup>, Cláudio Carvalho<sup>2</sup>, Maurício Magalhães<sup>1</sup> and Edmundo Madeira<sup>2</sup>

<sup>1</sup> Department of Computer Engineering and Industrial Automation (DCA)  
School of Electrical and Computer Engineering (FEEC)

State University of Campinas (Unicamp)  
13083-970, Campinas-SP, Brazil.

<sup>2</sup> Institute of Computing (IC)

State University of Campinas (Unicamp)  
13084-971, Campinas-SP, Brazil.

{verdi,mauricio}@dca.fee.unicamp.br,  
{claudio.carvalho,edmundo}@ic.unicamp.br

**Abstract.** In this work, we are taking into account the scenario with IP/MPLS client networks over an optical GMPLS network. Based on that, aggregating lower traffic flows (e.g. packet-based LSPs) within higher traffic flows (e.g. lambda-based LSPs) is considered an effective way to maximize the use of the optical network resources. Furthermore, there is a clear necessity to control the flows entering in the optical layer taking into account Service Level Agreements and differentiated services. In this paper we describe some policies and an architecture for policy-based grooming responsible for managing the installation and aggregation of packet-based LSPs within lightpaths (optical LSPs). By defining some policies the monitoring and the network planning can be specified in a higher level. The policies we have defined intend to reduce the optical network overhead to remove and reroute lower priority LSPs. The architecture presented here is composed of an Admission Control and a Policy Manager and the performed simulations indicated that the number of LSP removals is smaller when policies are applied.

## 1 Introduction

An optical network typically has ten to thousands of Gb of available bandwidth. These networks will likely consist of elements such as routers, switches, Dense Wavelength Division Multiplexing (DWDM) systems, Add-Drop Multiplexers (ADMs), photonic cross-connects (PXC) and optical cross-connects (OXC) [6]. Due to the advent of Generalized Multiprotocol Label Switching (GMPLS), there is a common sense to use it to dynamically provide resources and perform the routing and signaling functions of the control plane.

The idea of label switching in MPLS (Multiprotocol Label Switching) only considers the packet domain basically taking into account IP networks. However, due to the growing of new optical technologies and its high bandwidth, it is expected that many packet-based network flows will be nested within lightpaths

<sup>3</sup> to cross the optical domain and reach their destination. Lightpaths are seen as LSPs (Label Switched Paths) or optical LSPs (from now on optical LSP and lightpath will be used interchangeably) and because of technologies like DWDM it is now possible to have a very large number of parallel links between two directly adjacent nodes (hundreds of wavelengths, or even thousands of wavelengths if multiple fibers are used). Then, GMPLS has emerged as a solution to act not only in the packet or cell based networks but also in the time division networks (using Time Division Multiplexing-TDM), where the switching decision is based on time slots, and optical networks, where the switching decision is based on wavelengths or physical ports.

Although GMPLS considers all the above kinds of data forwarding, the one that is emerging is IP over DWDM networks. In this context, the overlay model is very indicated for service providers (e.g. Telecom companies) since they are the major part interested in acting as transport networks for client IP networks. A very typical and promising scenario is to have MPLS client networks with their packet-based LSPs asking for an optical resource (typically an optical LSP) in order to cross the optical domain and get their destination.

Depending on how the aggregation of packet-based flows within lightpaths is done, the use of the network bandwidth can be maximized or wasted. It is clear that if some rules are followed, the optimization of the network resources is increased and more traffic may be accepted. At the same time, considering that packet-based flows can be divided into different classes (e.g. DiffServ), it is necessary to give better treatment for higher classes. By defining the problem, we realized that some simple policies could efficiently improve the use of the network bandwidth. In this paper we present such policies, the architecture to apply the policies and the results of the simulations we have been doing.

This paper proposes a set of policies to manage the installation and aggregation of packet-based LSPs within optical LSPs assuming that there are several lightpaths between two end nodes. In a previous work [8] we shortly showed some results and introduced the architecture. In this work we detailed the policies and their modelling. Most of the recent works [5, 2] have dealt with heuristics for off-line routing, multilayer networks and multihop traffic grooming. They have not considered the use of policies to accommodate new on-line (dynamic) requests in the overlay model taking into account QoS and the overhead to remove and reroute packet-based LSPs. In this work, the packet-based LSPs are divided into two classes: Low Priority (LP) LSPs and High Priority (HP) LSPs. HP LSPs have higher priority than LP LSPs and, therefore, if an HP LSP needs to be installed in an optical LSP, some LP LSPs from that optical LSP may be removed to attend the higher requirement in the case of not having enough bandwidth in that lightpath. These removals are very expensive for optical networks due to the overhead to tear down one or more LP LSPs and reroute them in another optical LSP. Because of this overhead, reducing the number of removals is a very critical and important point to be considered in optical networks. The architecture pro-

---

<sup>3</sup> The aggregation of lower order LSPs within higher order LSPs is well known as traffic grooming problem.

posed in this work consists of an Admission Control responsible for verifying the Service Level Agreement (SLA) <sup>4</sup> and a Policy Manager responsible for applying the policies and finding an optical LSP to accommodate the packet-based LSPs. The policies we have defined are mainly intended to maximize the use of the network resources and reduce the number of LP LSP removals in the optical domain. Some representative situations in order to highlight the actual advantages of using policy-based grooming will be depicted. All the simulations have been made using the GLASS Simulator from NIST [4].

This paper is organized as follows: In the next section we shortly describe some basic concepts and the reference scenario. Section 3 presents our proposed architecture, its modules and the defined policies. Section 4 is dedicated to depict the results of our simulations. Finally, Section 5 presents the conclusion and future works.

## 2 Network Scenario and Technical Background

The concept of nesting LSPs (LSP within LSP) is already used in MPLS networks. In that domain, lower bandwidth LSPs whose destination node is the same are aggregated within another higher bandwidth LSP by means of label stacking. Considering that in GMPLS optical networks the bandwidth of an optical LSP is very high (from 1 Gb/s to hundreds of Gb/s), more lower bandwidth LSPs can be nested within the lightpath.


The concept of elasticity is also very important since it allows a packet-based LSP to increase or decrease its bandwidth as needed. Because transport networks will charge the clients for the used bandwidth, having a mechanism to give bandwidth to the connection only for the time it is actually requested is very important. When bandwidth is reduced, the portion that is released is left at the disposal of the network. The possibility of increasing and decreasing bandwidth raises a problem when classes of priority are defined. The problem occurs when a higher priority packet-based LSP wishes to increase its bandwidth. If there is available bandwidth in the tunnel, no LP LSPs need to be removed. On the other hand, if there is not available bandwidth, one or more LP LSPs will have to be removed to accommodate that new demand.

Fig. 1 shows an optical network and its physical topology. There are three optical LSPs between OXC A and OXC G. These LSPs represent a logical topology and are considered as Forwarding Adjacencies (FA) since each one is advertised as a Traffic Engineering (TE) link into IS-IS (Intermediate System - Intermediate System) or OSPF (Open Shortest Path First) protocols allowing other LSRs (Label Switched Routers) to use the FAs for their path computation [6]. This is a clear way to maximize the use of the optical network resources.

Fig. 1 also shows a typical scenario with MPLS networks acting as client networks and the optical network acting as the transport network (IP over optical network integration). This integration between IP/MPLS and optical layers is

---

<sup>4</sup> In [3] an example of an SLA applied to the optical domain is presented together with its parameters as well as their values for four classes of services.



**Fig. 1.** The overlay reference scenario.

done by means of a standard user-network interface (UNI) [7]. In this case, the IP/MPLS asks for a connection and the optical layer tries to set it up according to the SLA. The three optical LSPs shown in Fig. 1 are seen by client networks as tunnels from the ingress node 1 to egress node 2 and, therefore, from the point of view of the client networks, node 1 and node 2 are peers. Without applying any policy, the packet-based LSPs will be aggregated taking into account only the available bandwidth in each tunnel. Moreover, due to the elasticity concept presented above, each packet-based LSP needs to inform its maximum bandwidth that will eventually be needed.

### 3 The Proposed Policies and the Architecture for Policy-based Grooming

Our framework is responsible for receiving the request from the client, verifying the SLA between the client and the provider and applying the policies we have defined in order to find an optical LSP to be used [8]. All the information about the client (client ID, SLA, etc.) and the LSP (traffic parameters) being installed is carried by the PATH message of the RSVP (Resource Reservation Protocol). In this work we have only used RSVP as the signaling protocol and some few modifications were made on it to deviate the PATH message to reach our framework and to allow the interaction between the framework and the simulator.

The policies we have created try to separate HP LSPs from LP LSPs. The main point is to dynamically find a tunnel to allocate each packet-based LSP in order to minimize the number of LP LSP removals in case of increasing the bandwidth of HP LSPs. This decision is taken each time a new packet-based LSP ( $L_i$ ) is being installed. The PATH message carries the information about the LSP and each LSP has the following data:

- Requisition bandwidth: This is the quantity of bandwidth required by the LSP at the moment of its installation. After installing it, this value can be increased up to the maximum bandwidth and decreased down to the minimum bandwidth;
- Maximum bandwidth: This is the maximum bandwidth an LSP can ask for, based on what was agreed in the SLA;

- Minimum bandwidth also based on what was agreed in the SLA.

The policies are applied following the sequence presented below. A given policy will be applied only if its precedent failed, e.g., policy 2 is applied only if policy 1 fails and so forth. The allocation of an HP LSP is done taking into account only the maximum bandwidth in order to allow the HP LSPs ask for more bandwidth after the installation. Thus, in case of being an HP LSP, the verification of the bandwidth available in a given tunnel should consider the maximum bandwidth of all already installed HP LSPs in that tunnel (an example is showed when policy 2 is presented below). The requisition bandwidth of the HP LSPs is only considered when an LP LSP is being installed. While the HP LSPs do not ask for more bandwidth, the bandwidth that is left (the difference between the maximum bandwidth and the bandwidth that HP LSPs are actually using) can be used by LP LSPs. Therefore, when an LP LSP is to be allocated, not only the requisition bandwidth of all LP LSPs already installed needs to be verified but also the requisition bandwidth of the HP LSPs also needs to be considered since the sum of all of them represents the quantity of bandwidth that is actually being used in that moment in a given tunnel. Based on that sum, it is possible to know if there is enough bandwidth to accommodate the LP LSP. Below the proposed policies are presented.

### 3.1 Detailing the Policies

Let  $T_j$  be the  $j$ th tunnel between an ingress node and an egress node.

**if  $L_i$  is HP** (the allocation is done based on the maximum bandwidth of all already installed HP LSPs):

- **Policy 1:** Accommodate  $L_i$  in  $T_j$  if  $T_j$  is not empty and has only HP LSPs. This policy is always the first one to be applied and it assumes that if a tunnel has only HP LSPs, when they request more bandwidth no LP LSP removals will be needed;
- **Policy 2:** Accommodate  $L_i$  in  $T_j$  if  $T_j$  is not empty and has both kinds of LSPs. This policy has two restrictions and both must be matched:
  - a. Accommodate  $L_i$  in  $T_j$  if no LP LSP removals are needed at this moment, i.e., at the time the LSP is being installed;
  - b. Accommodate  $L_i$  in  $T_j$  if the sum of the maximum bandwidth of all HP LSPs and the sum of the requisition bandwidth of all LP LSPs do not exceed the bandwidth of  $T_j$ . Thus, if in the worst case all HP LSPs in the tunnel need to increase their bandwidth, no LP LSPs will be removed. For example, suppose that a new request to install the following LSP arrives: requisition bandwidth is 100 Mb/s and maximum bandwidth is 200 Mb/s, and suppose that the bandwidth of  $T_j$  is 1 Gb and has the following already installed LSPs:
    - HP LSP 1: Requisition bandwidth 100 Mb/s and Maximum bandwidth 200 Mb/s;
    - HP LSP 2: Requisition bandwidth 200 Mb/s and Maximum bandwidth 300 Mb/s;

- **LP LSP 3:** Requisition bandwidth 100 Mb/s and Maximum bandwidth 200 Mb/s;

The sum of the maximum bandwidth of all HP LSPs (including the one that is being installed) is 700 Mb/s and the sum of the requisition bandwidth of all LP LSPs is 100 Mb/s. Since the total sum is 800 Mb/s and the tunnel has 1 Gb, the LSP can be installed in  $T_j$ . On the other hand, suppose that the maximum bandwidth of the HP LSP 1 is 400 Mb/s and the maximum bandwidth of the LSP being installed is 300 Mb/s, the LSP would not be allocated in tunnel  $T_j$  since the total sum would be 1.1 Gb.

- **Policy 3:** Accommodate  $L_i$  in  $T_j$  if  $T_j$  is not empty and has both kinds of LSPs. This policy does not have the policy 2's restriction (b);
- **Policy 4:** Accommodate  $L_i$  in  $T_j$  if  $T_j$  is not empty and has both kinds of LSPs. This policy is the same as policy 2 without restrictions (a) and (b);
- **Policy 5:** Accommodate  $L_i$  in  $T_j$  if  $T_j$  is empty.

if  $L_i$  is **LP** (the allocation is done based on the sum of the requisition bandwidth of all already installed LP LSPs and HP LSPs):

- **Policy 6:** Accommodate  $L_i$  in  $T_j$  if  $T_j$  is empty;
- **Policy 7:** Accommodate  $L_i$  in  $T_j$  if  $T_j$  is not empty and has only LP LSPs;
- **Policy 8:** Accommodate  $L_i$  in  $T_j$  if  $T_j$  is not empty and has both kinds of LSPs;

In case of applying all the policies and not finding a tunnel, the LSP will be blocked and not installed. Note that an LP LSP installation does not cause any removal since there is only one level of priority among the LSPs of this type. Moreover, an LP LSP is allowed to increase its bandwidth only if there is the quantity it is requiring in the tunnel. When applying policy 4, there can be more than one available tunnel to allocate the LSP. In this case and considering that there will be LP LSP removals in each tunnel, we choose the tunnel whose number of removals is the smallest, thus reducing the overhead to tear down and reroute such LP LSPs.

Fig. 2 shows the partial class diagram used to model the policies described above. This class diagram is based on the Common Information Model (CIM) specification [1], which defines an approach on how to apply policies using actions and conditions. In Fig. 2 the conditions as well as actions can be combined (compound) using the classes *CompoundPolicyCondition* and *CompoundPolicyAction* in order to attend specific policies. When removals are needed, the actions *PolicyActionRemoveLSP* and *PolicyActionInstallLSP* are combined to support the policies that triggered the LP LSP removals. Other compositions can be done based on each specific policy. Furthermore, policies can be grouped and organized in sets to represent specific rules.

It is important to mention that there are some different ways to aggregate packet-based LSPs within lightpaths without using policies. Basically in this work

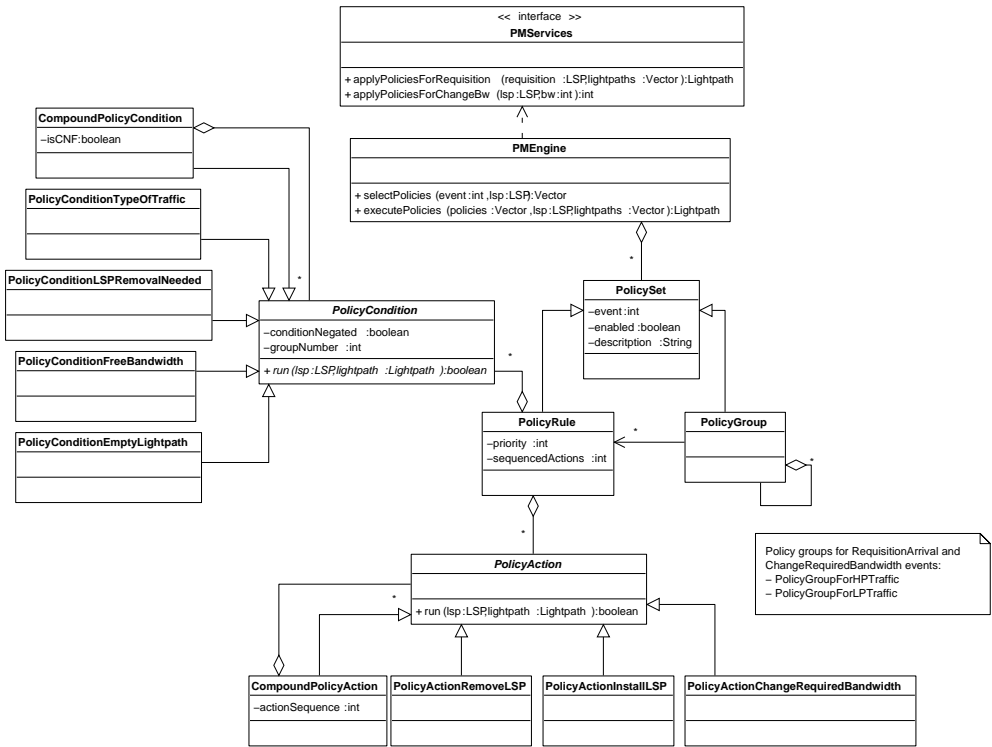


Fig. 2. Policy information model and its association with the system classes.

we assumed that when not considering policies, the aggregation is done by sequentially looking for a tunnel until find one that has enough bandwidth to accommodate a given LSP.

The architecture we propose for verifying the SLA and applying the above policies is shown in Fig.3. The Admission Control and the Policy Manager belong to the Bandwidth Broker responsible for managing the domain as a whole and coordinating the tasks of each inner module. Using the scenario presented in Fig. 1 we now added the following modules:

- **Admission Control-AC:** It receives the PATH message (step 1 in Fig. 3) and gets the information carried by it in order to verify the SLA between the client and the transport network provider. This verification is done when the client requests the installation of a new LSP and when an already installed LSP asks for increasing or decreasing its bandwidth. After verifying the SLA and in the case of agreement, the AC calls the PM (Policy Manager, see below) to apply the policies (step 2). Considering that LP LSP removals are necessary, the AC is in charge of interacting with RSVP to send a tear down message for each LP LSP that needs to be removed (step 3). After removing

the LP LSPs, the new HP LSP will be installed and thus the PATH message will follow its path (step 4). At the last (step 5), for each LP LSP to be removed a RESVTear message needs to be sent upstream.

- **Policy Manager-PM:** Responsible for applying the defined policies to find a tunnel to allocate a given packet-based LSP. If a tunnel is found and the LSP is an HP LSP, the PM verifies whether one or more LP LSPs need to be removed to accommodate such LSP. On the other hand, an LP LSP is only installed if there is a tunnel with available bandwidth to satisfy the request.

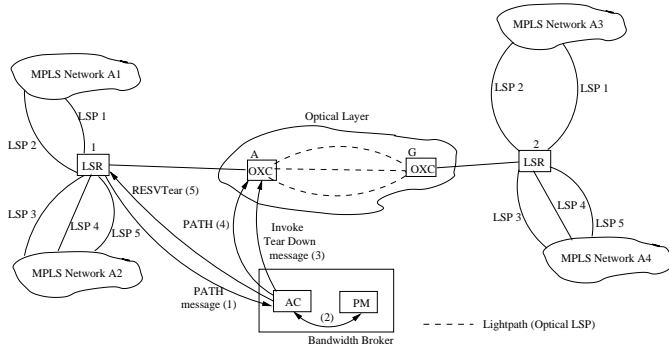


Fig. 3. The proposed architecture.

Although we have done our tests using a simulator, our proposed architecture is independent of it and can be applied in a real scenario just doing some small modifications in the RSVP to deviate the PATH message from the ingress node to the AC.

### 4 Implementation and Results

To simulate an optical network with IP/MPLS clients asking for connections (optical LSPs) we have used the GLASS Simulator from NIST [4]. GLASS is a discrete network event simulator for GMPLS-based Optical Internet and by using it, we created a typical scenario with IP over DWDM. The modules of our architecture - AC and PM - were implemented in Java and there is a very clear separation between the management plane and the control plane of the network being managed.

In order to test our policies we have created many different scenarios. The sequence used to apply the policies is the one presented in Section 3, i.e., from policy 1 to policy 5 for HP LSPs and from policy 6 to policy 8 for LP LSPs. The results showed in this section are based on that sequence. We assumed that the bandwidth of each tunnel is 1 Gb/s, the number of LSPs to be installed is 200 and about 50% of LSPs will ask for increasing their bandwidth. The LSPs are randomly created and their minimum bandwidth is 50 Mb/s and their maximum bandwidth is 400 Mb/s. As a consequence, the requisition bandwidth (see Section 3) varies from 50 Mb/s to 400 Mb/s. The number of available tunnels, and consequently the quantity of available bandwidth between two end nodes varies from



10 to 40. We randomly generated 200 LSPs to be installed and this was done 300 times. The average is obtained after repeating those 300 loops.

In the first simulations we created more HP LSPs ( $\approx 66\%$ ) than LP LSPs ( $\approx 33\%$ ). This case represents a situation in which there are more high priority LSPs than low priority LSPs. The requisition bandwidth average for HP LSPs is 16 Gb. The maximum bandwidth average for HP LSPs is 35 Gb and the requisition bandwidth average for LP LSPs is 8 Gb.

Fig 4 depicts the number of installed LSPs as the bandwidth increases. With 40 Gb all the LSPs are installed:  $\approx 133$  HP LSPs and  $\approx 66$  LP LSPs (out of 200) regardless of using policies or not. Fig. 4 intends to show that the number of installed LSPs is the same when we apply the policies. The difference is when the HP LSPs ask for increasing their bandwidth. In this case the number of removals is quite smaller when applying the defined policies as we can see in Fig. 5 (increasing the bandwidth in 50%).

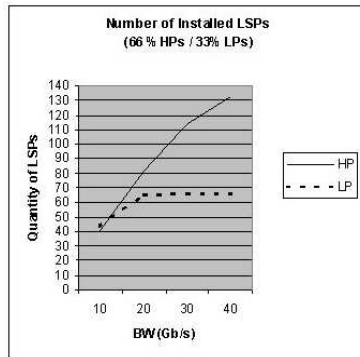


Fig. 4. Installed LSPs:  $\approx 66\%$  HPs and  $\approx 33\%$  LPs.




Fig. 5. LP LSPs removed after increasing the bandwidth in 50%.

While the number of removals without using policies increases up to 19 and keeps this behavior until the end, the number of removals when applying the policies decreases and with 30 Gb there is no more removals. Having 20 Gb of bandwidth available in the optical layer, the number of removals with policies is 6 and without the policies it is 19, a difference of 68%. The point is that in the case of not using policies, the more LSPs are accepted since the network bandwidth increases, the more the number of removals. However, this is not the case when applying the policies because we are separating HPs from LPs. Fig. 6 shows the same situation except that we are now increasing the bandwidth to the maximum allowed. Because we are requiring more bandwidth, the number of removals increases in both situations. With 20 Gb there are 35 removals using the policies and 42 without them, a difference of 16%. With 30 Gb there are 10 removals with policies and 52 without using them, a difference of 80%. Finally, with 40 Gb the number of removals is only 1 when applying the policies and 52 without using them.

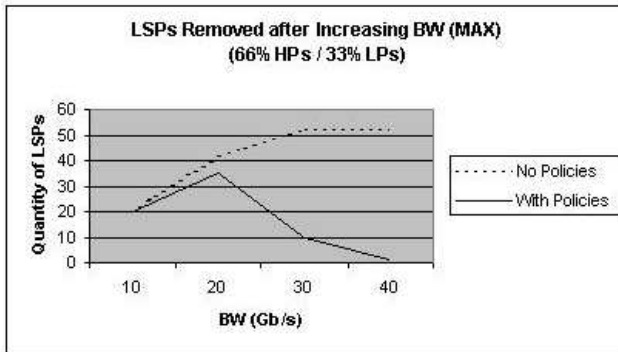


Fig. 6. LP LSPs removed after increasing the bandwidth to the maximum.

The remaining graphs represent the scenario in which the number of HP LSPs and LP LSPs are about the same, 50% for each type. In this case, the requisition bandwidth for HP LSPs is 12 Gb, the maximum bandwidth for HP LSPs is 26 Gb and the requisition bandwidth for LP LSPs is 12 Gb. Fig. 7 and Fig. 8 show respectively the results after increasing the bandwidth in 50% and to the maximum. We can observe that in Fig. 7 the number of LP LSP removals is greater when compared to Fig. 5. With 30 Gb the quantity of removals is 25 without using policies and 3 with policies, a difference of 88%. The result is better with 40 Gb: no removals with policies and 25 without them.

In Fig. 8 we can see that because there is more demand for bandwidth (increasing to the maximum) more removals are necessary when compared to Fig. 7. This is the same case when we compare Fig. 8 to Fig. 6. Since there are more LP LSPs in the network, more LP LSP removals will take place.

The sequence in which the policies are applied can generate different results. The results depicted in this section are based on the sequence presented in Sec-

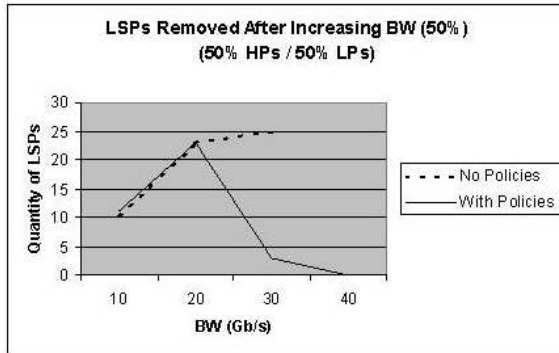


Fig. 7. LP LSPs removed after increasing the bandwidth in 50%.




Fig. 8. LP LSPs removed after increasing the bandwidth to the maximum.

tion 3 and we verified that it is the best sequence to be used. Another possible combination is the one in which policy 5 is applied before policy 2. By applying this sequence, the number of removals after increasing the bandwidth increases up to 10%, but in several cases the quantity of removals is the same as the results presented in this section. Another possible sequence is to apply policy 7 before policy 6. This combination is not feasible since the number of removals during the installation of the LSPs increases in 60% in some cases.

The results presented above represent two scenarios. The first one is when there are 66% HPs and 33% LPs and in the second one there are 50% HPs and 50% LPs. We have tested more scenarios and observed that if there are many LP LSPs to be installed (more than 50% of LPs) the policies will start helping only when the network bandwidth is 20 Gb. On the other hand, if there are less than 30% of LP LSPs to be installed the policies start helping when the network bandwidth is 10 Gb.

We also observed that in scenarios whose traffic flows are mostly composed of high priority flows (e.g. typically in multimedia scenarios) the use of policies is indicated.

## 5 Conclusions and Future Works

In this paper we have presented an architecture to manage the integration between the IP/MPLS layer and the optical network layer taking into account two different classes of services: High Priority (HP) LSPs and Low Priority (LP) LSPs. Some policies intended to reduce the number of LP LSP removals and consequently the necessity to reroute the LP LSPs in the optical layer were described. Also, many scenarios were tested and the results were shown. Such results indicated that the policies proposed in this paper seem to be appropriated since the number of removals by applying them is significantly reduced. Furthermore, the relation between the quantity of bandwidth of the network and the percentage of LP LSPs and HP LSPs used to estimate an upper bound in terms of quantity of traffic to take advantageous of policies is an useful parameter to support the transport network provider in deciding when the usage of policies is recommended.

This architecture makes part of a project in which other modules are being developed. An OVPN (Optical Virtual Private Network) service is being currently defined. A Fault Manager responsible for controlling the faults in the optical layer and applying corrective solutions is also part of the architecture. Finally, although we have tested several scenarios there are others to be created with more levels of priority (HP1, HP2, LP1, LP2, etc.) considering new policies and situations.

## Acknowledgments

The authors would like to thank CNPq and Ericsson Brazil for their support.

## References

1. <http://www.dmtf.org/standards/cim>.
2. R. Sabella et al. A Multilayer Solution for Path Provisioning in New-Generation Optical/MPLS Networks. *IEEE Journal of Lightwave Technology*, Vol. 21, Issue 5:1141–1155, May 2003.
3. W. Fawaz, B. Daheb, O. Audouin, M. Du-Pond, and G. Pujolle. Service Level Agreement and Provisioning in Optical Networks. *IEEE Communications Magazine*, Vol. 42, Issue 1:36–43, January 2004.
4. GLASS Simulator: <http://dns.antd.nist.gov/glass/>.
5. P. Iovanna, M. Setembre, and R. Sabella. A Traffic Engineering System for Multilayer Networks Based on the GMPLS Paradigm. *IEEE Network*, Vol. 17, Issue 2:28–37, March/April 2003.
6. E. Mannie. Generalized Multi-Protocol Label Switching Architecture. *draft-ietf-ccamp-gmpls-architecture-07.txt*, 2003.
7. OIF User Network Interface (UNI) 1.0 Signaling Specification.
8. F. L. Verdi, E. Madeira, and M. Magalhães. Policy-based Admission Control in GMPLS Optical Networks. *First IEEE International Conference on Broadband Networks - Broadnets'04 (formerly OptiComm)*, pages 337–339, San Jose, USA, October 2004.