Policy-based Provisioning in Hybrid Photonic Networks

Wissam Fawaz[124], Belkacem Daheb[14], Olivier Audouin[3], Bela Berde[3], Ken Chen[2], Guy Pujolle[1]
[1]University Paris 6 - LIP6 Lab, 8, rue du Capitaine Scott, 75015 Paris, France
[2]University Paris 13 - L2TI Lab, 99, Avenue Jean-Baptiste Clément, 93430 Villetaneuse, France
[3]Alcatel Research & Innovation, Route de Nozay, 91460, Marcoussis, France
[4]Institut Supérieur d'Electronique de Paris (ISEP)
{wissam.fawaz, belkacem.daheb}@lip6.fr

Abstract

The aim of this work is to present provisioning strategies for GMPLS-enabled Hybrid Photonic Networks (HPN). Contributing to network flexibility, the paper presents extensions to the management system of these transparent wavelength and switching capable networks that provide the means to leverage their inherent capabilities.

The paper first addresses the motivation and utility of HPN. Building on this information, it goes on with the constraints to be applied in provisioning of those networks. Finally, to achieve the strategic goal of provisioning, a protocol independent Policy-based Management (PBM) approach is proposed with the corresponding policy control framework and relevant policy categories.

Keywords: *Hybrid Photonic Network, SLA, Service Provisioning, Policy-based Management.*

1. Introduction And Background Material

Historical networks providing predominated voice services grew from scratch at less than 10 percent each year. Recently, however networks are witnessing a drastic change in the overall networking picture due to the explosive growth in IP-centric traffic which already surpassed voice traffic and will continue to outpace voice for years to come. This shift has created a demand for capacity and has a profound impact on the network architectures.

The most obvious strategy needed to deal with this new situation consists in boosting transmission capacity. Optical DWDM transmission has thus become a key technology to accommodate the continuing expansion of demand that keeps on fueling the growth of data traffic.

Nonetheless, blindly augmenting network capacity is not the long-term solution. With regard to cost efficiency, although laying multiple fibers may help to reduce the transportation cost, yet it cannot serve the ultimate goal, as the complexity and cost are being shifted to the bottleneck switching and regeneration nodes. Therefore, in order to keep up with the incumbent challenges, next-generation optical carrier networks are expected to support the increasing load by employing advanced transmission (DWDM), and new switching technologies such as hybrid (transparent-opaque) optical crossconnects [1].

The revolutionary DWDM technology increases transmission capacity of fiber links by several orders of magnitude. This huge increase of capacity challenges the switching equipments managing the wavelengths. It is in this context that emerging hybrid hierarchical optical crossconnects [2] become an attractive solution in next-generation optical networks. In addition, the hybrid technology provides significant expenditure savings, since it replaces much expensive opto-electronic fabric with an all-optical one. This potential is augmented by the hierarchical technology merit (i.e., waveband switching), which further reduces capital expenditure since the same optical port can process multiple wavelengths simultaneously. Indeed, hybrid crossconnects are constituted of a transparent waveband switching stage and of a regenerative wavelength switching stage with a partial capacity with regard to the overall node throughput Figure 1. A waveband is formed by a set of wavelengths, and is either switched in the optical domain to another waveband or dynamically directed to the wavelength switching stage where electronic processing is performed. The transition into the electronic stage is only required to regenerate a wavelength, to aggregate traffic into it, or to switch it to another waveband. Such a hybrid switching environment will heretofore be referred to as a Hybrid Photonic Network (HPN).

Configuring network elements in a HPN to provide a specific service requires flexible and simultaneous configuration of more than one network element in the network. The concept of Policy-based Network Management (PBM) [3] addresses that problem and offers solutions. To backup our study with concrete business objectives related to the optical domain, the Optical Service Level Agreement (O-SLA), which we defined in [4], was of great help.

1530-2016/04 $20.00 © 2004 **IEEE**

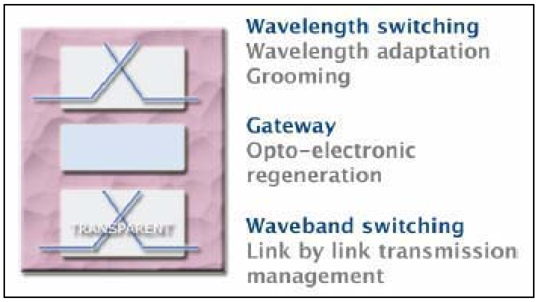

Figure 1. Hybrid crossconnect architecture

This is especially true, since there are no other Service Level Agreements defined in the literature that are adapted to the peculiar needs of optical networks. As a matter of fact, the O-SLA serving as a formal contract was intended to provide the optical operators with guidelines on how to propose different optical services and service classes to their clients. To meet this purpose, the different Service Level Specifications (SLS), embodied in the technical part of the O-SLA, were defined, including the individual metrics and operational data needed to achieve the agreed-upon quality measures for network traffic. In this paper, the different policy categories were deduced of these SLS criteria.

However, since the policy approach is a very general one and has to solve a number of issues simultaneously, it is useful to examine its application to GMPLS-enabled HPN [5], for which the internal topology abstraction consists of Traffic Engineering links (TE links) and the set of advertised Forwarding Adjacency (FA) [6]. They form the topology perceived in the control plane over which the routing algorithm will run. The topology view will differ whether the computation of the explicit route is for a waveband-LSP or for a lambda-LSP. Following the LSP nesting principle, and in the considered HPN context, a waveband-LSP must be established before establishing the lambda-LSP to be nested in. The waveband coverage is a space-time dependent problem and is based on a predictive and approximate statistical traffic analysis. Due to this traffic approximation, and with the intent to keep some flexibility at the waveband level, it is envisaged not to dedicate physical resources to each pre-computed waveband. Two different types of Forwarding Adjacencies are thus built: soft FAs and traditional (or hard) FAs. The difference between these two types relies on the relation between the FA and the physical resources. To a traditional FA, there are physically assigned resources while the association of resource is virtual for the soft FA.

The next section provides some background on specifics to the provisioning in HPNs.

2. Provisioning In HPN Networks

The performance of the provisioning process in any kind of networks influences the network flexibility, that is, the ability to rapidly reconfigure connections to deal with service requirements as well as traffic variations. The main objective is always to avoid a costly over-dimensioning of resources to absorb traffic dynamics.

In order to best meet service requirements stated in the different O-SLAs, the relevant SLS parameters are identified and taken into account during the waveband establishment. In adapting wavebands to accommodate this requirement, there are several options:

- Waveband per class of service
- Waveband encompassing several class of services
- Waveband independent from class of service

Choosing one of these options is left for further investigation. But, of course the first two will necessitate the use of policies during the spatial and spectral routing phases [5] in an attempt to meet the service requirements of the relevant class of service. To wrap up things, the alternative is to choose between:

- Establishing waveband-LSPs irrespective of service needs, and leave the whole complexity of service provisioning into the lambda-LSP establishment phase.
- Account to specific SLS parameters during the waveband-LSP establishment, and divide the complexity with the lambda-LSP establishment phase.

Our decision is heavily directed in favor of the second choice, as it means less complexity at the level of each phase, and a reduced blocking probability. But still, the efficiency of resource utilization relative to this choice must not be neglected, and should be studied carefully.

Once waveband-LSPs have been pre-established, lambda-LSPs are expected to be nested in the advertised waveband FAs. In the routing process, the computation of an explicit route for the lambda-LSP will make effectively use of FAs. Based on what has been stated before, the following questions can be raised:

1) When do new LSPs have to be pinned down, or removed?
2) How can an LSP be established in a way that meets all service objectives?
3) After the establishment of an LSP, how is the right traffic mapped into the LSP?

Answers to these three questions can be given, thanks to well defined rules or policies. The idea is to associate policy rules to each service request. Based on these rules, new LSPs can be established or old ones can be re-used. When the decision is to use a new one, the LSP is established in such a way that service requirements are respected along the path of the LSP.

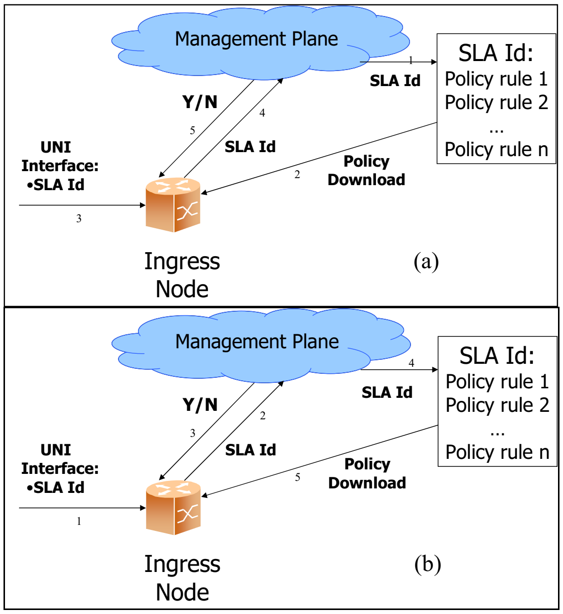

Figure 2. Framework of the policy enforcement: (a) pre-provisioned, (b) per session provisioned policies.

Finally, the user traffic is injected into the established LSP, based on some other directives.

These different policies will be categorized with more explanation in a later section. The next section provides the framework defined for the purpose of policy control and that serves as a frame for the picture depicted through the current section.

3. Policy Control Framework

The overall framework that has been retained for the purposes of controlling the network using policies is shown Figure 2. The rationale behind this framework lies in the assignment of a unique O-SLA Identifier (*SLA-Id*) to a customer, once a service contract has been negotiated with the operator. The *SLA-Id* is needed to connect a client to the service being requested, as it is possible for one customer to contract several services and therefore O-SLAs with the same operator. The *SLA-Id* would serve as a unique attribute based on which a specific service contract is recognized by the management and control planes.

The utility of the *SLA-Id* is augmented by the necessity to associate with each contracted O-SLA a number of policy rules whose role would be to ensure the right enforcement of the service. Such an association is made possible through the use of the *SLA-Id* object, which can provide the needed link to index the different rules connected to a specific O-SLA. As soon as this object has been assigned to the customer, two main events stamp the service provisioning process of the identified service contract.

The first one is the client's session request via the UNI interface, which entails a classical Call Admission Control (CAC) performed by the management plane. In the request, the message is conveying the *SLA-Id* object. The *SLA-Id* is then communicated by the ingress node to the management plane in order to perform the CAC function. The result of the admission control operation would be either to grant or deny the user access to network resources. The decision however is based on three main factors:

- The service contract referred to by the *SLA-Id* object:

For instance, if the customer provides an *SLA-Id* object that the operator will not be able to identify, the service request would be refused.

- The conformance to the *service schedule* stipulated in the O-SLA contracts:

In this regard, if the service request is received in a time period outside of that indicated by the negotiated service schedule parameter, this can result in a service denial.

- The Network State:

The actual state of the network can determine whether it is possible to accommodate the new service request. So, when the network is in a congestion state, it is better to first test the class of service associated with the demand before taking any admission decision. Hence, if the service is of type Premium or Gold for example, it would be more beneficial to preempt lower level services in order to accommodate the new request. Otherwise, the user will not be granted the access to network resources.

Besides the service request procedure, the second event that can be highlighted during the service provisioning process would be policy provisioning. In fact, as stated previously, with each contracted O-SLA is associated a set of policy rules. Some of these policy rules are inferred from the service contract, and intend to make sure that the service is being deployed under good conditions. Other policy rules serve to implement and enforce operator's global objectives. Importantly, during this phase, the management plane downloads the rules to the ingress node. However, in order to identify the group of policies associated to a specific O-SLA, the management plane employs the *SLA-Id* object that serves as an index pointing out the different rules related to the O-SLA.

With regard to the chronological order of the two distinct events addressed in the previous paragraph (represented by the arrows in Figure 2), one event can precede the other depending on the provisioning strategy. In this regard, two scenarios are to be distinguished:

1) Policies are downloaded *a priori* (Figure 2.a). In that case, policy rules pertaining to each O-SLA

are provisioned by the management plane toward the ingress node prior to customers request arrival.

2) Policies are downloaded on a per session basis (Figure 2.b). It is only when a session request is received from the customer's side that policy rules are provisioned into the ingress node.

Choosing between one of these two scenarios must be based on the operator's global objective and need. For example, if the operator tends to pre-provision lightpaths in his network, then the second scenario is not a good fit. However in this case, the first one would be the right choice, as policies can be provisioned beforehand, that is, before the customers request arrives. Then the process of pre-establishing lightpaths could be guided by the policies downloaded *a priori.* The second scenario is useful in cases where the operator establishes dynamic O-SLA with its clients. In other words, a dynamic O-SLA is a contract where a subset of SLS parameters can be changed easily over time. Service negotiation for a dynamic O-SLA is thus not performed manually but rather via a protocol transporting the varying SLS parameters. In the context of the present framework, service negotiation for a dynamic O-SLA could be done using the UNI interface which can be extended to transport objects related to SLS parameters. When doing so, the policies related to an O-SLA are not downloaded unless a service request is received from the customer specifying the desired values for the SLS parameters.

4. Policy Categories

The O-SLA defined in [4] provides guidelines of client expectations regarding service fulfillment. As such, the different policy categories useful for this purpose may be deduced from the corresponding SLS parameters.

The relationship defined by the O-SLA considers a Service Provider to be an optical carrier operator, and a service subscriber to be either an optical client or an IP/MPLS client. An optical client subscribes for services with a granularity equal to a wavelength, waveband, or even a complete fiber. On the other hand, an IP or MPLS client, within the context of our proposal, subscribes services at a granularity that is smaller than a wavelength, and his traffic may undergo a process of grooming or aggregation provided by the optical operator's network. As part of the O-SLA, we included generic parameters applicable to any SLA, such as the *service boundary*, specifying the geographical region over which the service's various QoS parameters are to be enforced, the *service schedule*, indicating the start and end time of a service, and the *Flow Id*, identifying the data flow receiving the service guarantees. In order to specify how long it will take for a service connection to be established once it has been negotiated, the parameter *connection setup time* was introduced. Since service recovery is of great importance in optical networks, parameters related to *service availability and resilience* have been stipulated, indicating the triggering of the resilience mechanism, as well as how long it takes to reestablish a failed connection. Under the category of routing constraints, we identified parameters like *routing stability*, referring to how often traffic trunks can be rerouted, and *confidentiality*, defined as a routing constraint where a confidential connection was considered as a transparent one, and where no grooming or electronic processing is allowed along the route. Cases where clients might request services consisting of two or more LSPs not belonging to the same Shared Risk Link Group [7], or excluding some regions along the route of their connections, are accounted for through the *route differentiation attribute*. Finally, classical *service performance guarantees*, *traffic conformance*, and *excess treatment* were introduced as well in the SLS. Distinction, when needed, between optical and IP clients was presented in the O-SLA for each of the previously-mentioned parameters.

SLS parameters defined in the O-SLA are classified into:

- *Traffic flow related parameters*, such as like Flow Id, traffic conformance, and excess treatment. This is normal since flow id identify the traffic flow for which the service would be provided. While traffic conformance indicates the profile based on which the traffic is classified as either in or out of profile, excess treatment precise the way of treatment for out of profile traffic.
- *Control plane related parameters*, including routing constraints, service performance guarantees, and service availability and resilience are included. These parameters characterize the lightpath that will be setup using the control plane.

Based on the above taxonomy, the following policy categories are identified as crucial to ensure efficient provisioning in GMPLS-enabled HPN:

1) Routing policies
2) LSP life-cycle management policies
3) Flow management policies

A. Routing Policies

These policies support the selection of the path taken by the lightpath and to ensure the requested performance characteristics. The performance of a lightpath is tightly related to the characteristics of the links assigned to it.

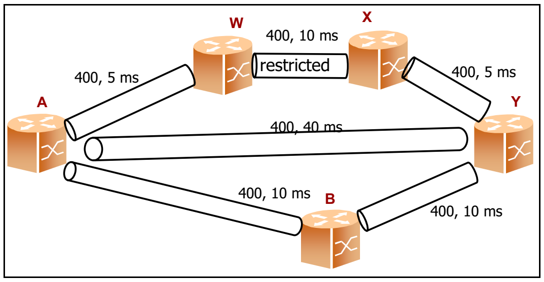

Figure 3. Provisioning Scenario (delay = 30ms, BW = 300)

Hence, the route calculation is an important step during the lightpath creation. In GMPLS the path computation is performed by a Constraint-based Routing function (CBR), using a Constraint-based Shortest Path First (CSPF) algorithm [6], which uses the following information as input:

- •SLS parameters characterizing the lightpath. For example performance guarantees parameters (bandwidth, delay)
- •Attributes associated to resources, i.e. FA attributes, indicating resource availability in the network
- •Other topology information

Based on this information, a CBR process located on each node computes explicit routes for lightpaths originating from that node. The explicit route is a specification of a path that satisfies the requirements expressed in the O-SLA, subject to constraints imposed by resource availability and other topology state information. The SLS parameters characterizing a lightpath are twofold:

- – Quantitative parameters: such as performance guarantees parameters (bandwidth, delay).
- – Qualitative parameters: such as route differentiation, and confidentiality attributes.

Therefore, the impact of routing policies on route calculation would be according to two axes. The first one is related to quantitative parameters, i.e. bandwidth and delay for example. In this case, the role of routing policies would be to prune from the internal topology abstraction, over which the routing algorithm is running, the links that do not meet the service exigencies. For instance, if a connection request necessitates 300 bandwidth units, and a delay of 30 ms, the routing algorithm, due to routing policies, would be able during the path computation process to exclude and remove all FAs presenting bandwidth less than 300, or a delay more than 30 ms. However, while pruning FAs not having enough bandwidth is enough during the computation process, it is not the case for the delay parameter. In other words, it would not be sufficient to prune FAs with excessive delay since such a parameter is a cumulative one and we must make sure that the final selected path verifies the end-to-end delay specified in the O-SLA. This leads us to another policy rule ensuring that the selected path does comply with the end-to-end delay; otherwise the selected explicit route is not retained. Then based on the traffic priority new FAs are built or preemption of other LSPs is performed, in an attempt to find another adequate route. It is important to note in this regard that the bandwidth as well as the delay offered by an FA are parts of the FA attributes flooded throughout the network by the routing protocol, and thus accessible by the node performing path computation.

Moreover, through routing policies it would possible to handle qualitative attributes during the explicit route determination, which leads to the second axis of routing policies impact. In other words, one can ask: how can we build a lightpath satisfying the confidentiality attribute? And how can we avoid some links from being associated to the lightpath in order to satisfy a certain route differentiation requirement? Answers to these questions are likely to be found when deploying routing policies. In fact, resources are administratively colored in such a way that resources with the same color belong to the same class.

This color concept is already defined as an attribute to TE links and FAs and flooded in the network by the routing protocol [8]. Following this idea, a color is assigned to each TE link or FA using routing policies. Next, the path is explicitly restricted to specific subsets of resources identified by a common color. For example, considering the Route differentiation SLS parameter of the O-SLA, the lightpath may not be supposed to pass through a certain country X. Then, based on routing policies, the links falling within this country would be provided a certain color Y by the management plane. At the same time, the path computation process would be informed to exclude color Y during route calculation. It is finally important to note that both FA coloring and lambda-LSP path restriction are based on routing policies.

In order to illustrate the merit of routing policies, let us consider the following scenario. An optical client is requesting a connection of 300 bandwidth units, and a delay of 30 ms between ingress point A and egress point Y (see Figure 3). Furthermore he asked that the connection does not pass through the region situated between W and X as part of the route differentiation requirement. Figure 3 depicts the network topology abstraction, where the couple of values situated besides each FA represent respectively the available bandwidth and the delay through the designated FA. It is obvious that the FA between W and X has been colored as restricted during the pre-provisioning process.

The set of routing policies that would derive from this service request would be the following:

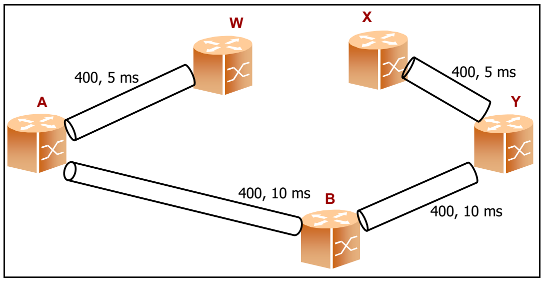

Figure 4. Provisioning scenario: resulting virtual topology

- *If ((FA bandwidth < 300) or (FA delay > 30 ms)) then prune the designated FA from the virtual topology during computation process*
- *If (FA is restricted) then prune from virtual topology*

These two policies will guide the CSPF operation for path computation. First, based on the first set of policies, CSPF will prune links that do not satisfy the quantitative parameters constraints. Afterwards, links not satisfying the qualitative attribute are removed. The previous two filtering operations will result in the virtual topology depicted in Figure 4. Finally, CSPF runs a Dijikstra algorithm over the remaining virtual topology in an attempt to obtain the explicit route for the service request. Within this scenario, the resulting route includes the FAs along the path A, B, and Y. Finally, the policy rule, examining the end-to-end delay of the selected path, verifies whether the required delay is met along the chosen route, which is the case here.

B. LSP Life-cycle Management Policies

These policies deal with crossconnect configuration covering initiation, maintenance, and removing of lightpaths. The rationale behind such policy category is to determine through events triggering the creation, rerouting, or the deletion of LSPs. For example, policies belonging to the LSP creation policies could state that when the connection request is received from an optical client, then new lambda-LSPs must be pinned down. Another example could concern the distinction between hard and soft FA's. The process of hardening a soft FA introduces a certain delay during provisioning. Consequently, a rule can privilege the use of hard FAs for service requests with stringent connection setup time in order to avoid the delay introduced by hardening a soft one. In a same way, always in the concern of meeting strictly reduced connection setup time, life-cycle management policies at ingress of soft FA could envisage hardening a soft FA, if all hard FAs going to the same destination approach full load.

On the other hand, policies relevant to LSP deletion provide the guidelines on when to remove an existing LSP based on SLS parameters such as *service schedule*, *routing stability*, along with LSP traffic load. For instance, when a service request transported along a certain LSP reaches a time period outside the *service schedule*, the corresponding LSP must be deleted. In the same regard, when the load of the LSP attains a certain threshold for a certain amount of time, and if the *routing stability* of the service making use of this LSP allows for rerouting, then traffic is rerouted and the underutilized path is removed.

To recapitulate the process, LSP life-cycle management policies include three main types of policy rules. The first one is related to LSP creation, while the second is pertaining to LSP deletion. Finally the third is concerned with modification of parameters of established LSPs. The first two types are exemplified next, based on the different cases presented previously. An example of LSP creation policies can be:

- *If (optical client) then create new LSPs*
- *If (connection setup time < threshold) then privilege the use of hard FAs*
- *If (reserved bandwidth of hard FAs) > threshold then harden a soft FA*

The following examples can be enclosed under the category of LSP deletion policies:

- *If (Time is outside service schedule) then delete LSP*
- *If (LSP load < threshold) and (routing stability allows for rerouting) then reroute traffic and delete underutilized LSP*

C. Flow Management Policies

These policies tackle classification directives for mapping data flows onto lightpaths. It is important to filter flows that will use network resources based on the Flow Id directive defined in the O-SLA. Furthermore, once the traffic flow is identified, policing and shaping are applied based on the *traffic conformance*, and *excess treatment* directives.

The basic idea is to first make sure that the right user is being served by testing his identity against the Flow Id parameter specified in the service contract. Then later on, his traffic profile must be examined against the *traffic conformance parameter* to recognize whether the traffic is in or out of profile, in order to identify the type of treatment to apply, such as shaping or degrading.

Two policy rule examples are provided next for both "in profile" and "out of profile" traffic treatment.

a) *In profile case:* For the in profile case, no shaping or degrading action is performed on the designated traffic, and the following policy can be put under this category:

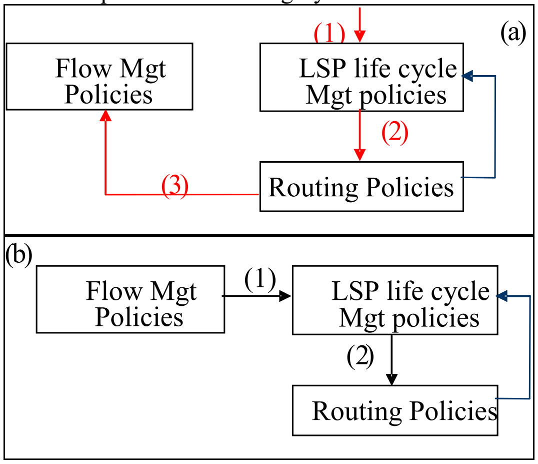

Figure 5. Policy rule categories Finite State Machine: (a) pre-provisioning, (b) per session provisioning.

- *If ((Client Id = Flow Id) and (time within service schedule) and (traffic profile = traffic conformance)) then transmit into corresponding LSP*

b) *Out of profile case:* When traffic flow is out of profile, then it is either shaped or degraded according to what is stipulated in the service contract through the Excess Treatment parameter. The following policies depict the two cases of shaping and degrading traffic in the action part:

- *If ((client Id = Flow Id) and (time within service schedule) and (traffic profile != traffic conformance) and (Excess treatment = shaping)) then shape and transmit into an LSP conform to the O-SLA*
- *If ((client Id = Flow Id) and (time within service schedule) and (traffic profile != traffic conformance) and (Excess treatment = degrading)) then degrade though the transmission into a non-conform LSP (degradation)*

Each one of the policy categories discussed before, can be viewed as a state in a Finite State Machine Figure 5. At each state, the policies related to the corresponding policy category would be activated and thus enforced. However, the sequence of the different states depends on the adopted provisioning strategy previously described in section III.

If the operator chooses to pre-establish the LSPs beforehand, even before the service request is received from the client side, then the LSP life-cycle management policies would be activated first, see slanted numbered red arrows in Figure 5.a. Next, once a decision of establishing LSPs has been taken, routing policies are applied to determine the right path for these LSPs. Finally, at the client traffic receipt, flow management policies are put into action for filtering, policing, shaping, and mapping the flow to an existing LSP.

However, if the network operator decides the establishment of LSPs after the receipt of user's UNI request, then flow management policies will be the one to be activated first. Next, the LSP life-cycle management policies are used in order to take the decision whether to establish a new LSP or use an existing one, see black arrows in Figure 5.b. Finally, it would be the turn of routing policies that make sure that the LSP being established remains conform to what is stated in the service contract.

But, it is important to note that there are other possible provisioning scenarios where the mentioned policy categories are combined together in a different manner. For instance, the arrow starting at the routing policies and terminating at the LSP life-cycle management policies (in both figures 5.a and 5.b) indicates a case where an LSP is not found during the routing phase. In other words, if the routing phase would not be able to find a suitable LSP on the existing virtual topology, the LSP life-cycle management state would be activated to create new FAs.

5. Conclusions

In parallel to the adapting of GMPLS to accommodate the characteristics of innovative hybrid optical technologies, this paper aims at contributing to the purpose of proposing and enhancing PBM solutions to the provisioning problem in GMPLS-enabled Hybrid Photonic Networks. The GMPLS control plane on its own cannot guarantee directly the fulfilment of all the service objectives specified in an O-SLA, along with objectives coming from network performance and business rules. The proposed solution clearly should lead to network operational and exploitation gains to be derived by deploying such a provisioning system, and especially in efficient provisioning of services with an enhanced accuracy. The implementation of the corresponding PBM system is already under way in a simulator environment.

6. References

[1] R. Izmailov and al., *Hybrid Hierarchical Optical Networks*, IEEE Communication Magazine, November 2002.

[2] J.-P. Faure and al., *A Scalable Transparent Waveband-based Optical Metropolitan Network* ECOC'2001, Amsterdam, The Nederlands, October 2001.

[3] A. Westerinen and al., *Terminology for Policy Based Management*, RFC 3198, 2001.

[4] W. Fawaz, B. Daheb, M. Du-Pond, G. Pujolle, O. Audouin, B. Berde, M. Vigoureux, *SLA and Provisioning in Optical Networks*, IEEE Communication Magazine, January 2004.

[5] M. Vigoureux and al., *GMPLS Architectural Considerations for (Hybrid) Photonic Networks*, IETF Draft, draft-vigoureux-ccampgmpls-architecture-hpn-00.txt, June 2002.

[6] E. Mannie and al., *Generalized Multi-Protocol Label Switching (GMPLS) Architecture*, IETF Draft, draft-ietf-ccamp-gmpls-architecture-07.txt, May 2003.

[7] D. Papadimitriou and al., *Inference of Shared Risk Link Groups*, IETF Draft, draft-many-inference-srlg-01.txt, 2001.

[8] D. Katz, D. Yeung, K. Kompella, *Traffic Engineering Extensions to OSPF version 2*, IETF Draft, draft-katz-yeung-ospf-traffic-09.txt, October 2002.