

An Experimental GMPLS-Controlled Network Test-bed Enabling Sub-Wavelength Connection Provisioning

F. Agraz, L. Velasco, J. Perelló, M. Ruiz, S. Spadaro, G. Junyent and J. Comellas

Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC)
Jordi Girona 1-3, 08034 Barcelona (Spain), e-mail: comellas@tsc.upc.edu

ABSTRACT

The efficient accommodation of sub-wavelength client flows on optical channels is a current challenge for resources use optimization in GMPLS controlled optical networks. While optical channels capacity usually exceeds 10 Gbit/s, client connections requests show finer granularity. This paper concentrates on the design and implementation of a GMPLS-controlled grooming capable transport infrastructure on the ASON/GMPLS CARISMA test-bed. Through the paper, the operation of a GMPLS-controlled multi-layer network architecture is introduced, subsequently highlighting implementation issues that come into light. The contributions proposed in the paper are supported by illustrative experimental results.

Keywords: GMPLS, multi-layer, test-bed

1. INTRODUCTION

Wavelength-routed optical networks have received increasing attention as a promising approach to deploy end-to-end transparent networks in a cost-effective way. The definition of the ITU-T Automatically Switched Optical Network (ASON) architecture [1], allows wavelength-routed optical networks to include dynamic connection capability. This capability is accomplished by means of a control plane entity, responsible for the establishment, maintenance and release of connections over the optical transport plane. The IETF has standardized Generalized Multi-Protocol Label Switching (GMPLS, [2]) as a set of protocols to implement a common control plane, able to manage several switching regions in an integrated way. GMPLS is the most widely accepted solution to implement the control plane functionalities in the ASON architecture. These ASON networks with a GMPLS-capable control plane will be hereafter referred as ASON/GMPLS networks.

The role of IP as a convergent technology has triggered the development of a wide range of new multimedia services, like HDTV, video conference, telemedicine applications or Internet telephony, having each one different Quality of Service (QoS) requirements. This huge, heterogeneous and predominantly bursty generated traffic poses new challenges to network operators to provide a cost-effective data transmission. Because the bandwidth granularity of wavelength-routed optical networks is very coarse, typically a whole wavelength supporting 10 or even 40 Gbps, these networks lack the flexibility to support sub-wavelength traffic demands, which lead to a poor bandwidth usage.

In this context, the term *traffic grooming* identifies the process of packing several low-speed traffic streams into higher-speed streams, trying to maximize optical channels bandwidth usage in WDM meshed transport networks [3][4].

From the GMPLS point of view, the grooming problem is translated into merging several higher-order Label Switched Paths (LSPs) into a lower-order LSP (e.g., grooming packet LSPs carrying IP traffic into a λ -LSP). Such an LSP aggregation in GMPLS is accomplished by advertising newly created lower-order LSPs as Forwarding Adjacency LSPs (FA-LSPs, [5][6]), for instance, by means of the OSPF-TE protocol[7]. In this way, conventional data-links along with the previously advertised FA-LSPs can indistinctly enter the path computation process. Supposing that a valid route would be found, resource reservation would then be performed by RSVP-TE [8].

The goal of this paper is to introduce the design and implementation of the grooming-capable ASON/GMPLS CARISMA test-bed. To this end, the implementation of the FA-LSP functionality in the ASON/GMPLS CARISMA test-bed is reported. Later on, the FA-LSP performance in a 9-Node meshed scenario is assessed by experimental results.

2. FA-LSP CONCEPT AND IMPLEMENTATION ON THE ASON/GMPLS CARISMA TEST-BED

From an architectural point of view, the control plane in multi-layer networks can follow three differentiated models namely overlay, augmented and peer [9]. In the traditional overlay model different control plane instances run on each layer. Thus, they are independently controlled. Alternatively, in the augmented approach, different control plane instances run on each layer but some information is exchanged amongst them aiming at improved bandwidth allocation in the network. Finally, in the peer approach all layers are controlled by a unified control plane, and decisions are taken considering whole network information.

The enhanced TE protocols introduced in GMPLS pave the way to peer multi-layer network architectures, controlled by means of a GMPLS-enabled common control plane. The enabling entity to this goal is the FA. In

GMPLS, those already established lower layer LSPs (e.g., λ -LSPs) are advertised as FA-LSPs, which can be used to transport new client LSPs. In this way, lower layer resources can be more effectively utilized.

A two-layered network peer architecture is assumed in this work, that is, an optical server layer and a client aggregation layer on top (e.g., SONET/SDH, MPLS, GbE etc.). This layer allows the mapping of the client traffic to be transported over the DWDM physical layer. At the bottom, optical nodes provide network ports as well as client access ports, used to inject an aggregated client flow to the network. The incoming signal would be afterwards adapted, switched to a network port, multiplexed into a DWDM bundle and finally transmitted to the next optical node. On top, the client aggregation layer includes generic nodes providing electrical switching, flow aggregation and other features. Client nodes are connected to optical nodes through the client access ports.

2.1 Definition of the FA-LSP routing metric and establishment cost function

In current GMPLS standardization [5] there is an intrinsic association between the signaling of new client LSPs and the creation of the required λ -LSP to support them. As will be later detailed, a route from source to destination is computed upon client LSP request, which may be constituted of both unallocated data links and already existent FA-LSPs. In the case that no FA-LSP is comprised along the route, a new λ -LSP is typically set-up from source to destination to support the incoming request. Otherwise, λ -LSPs are set-up to provide connectivity on those route segments where no FA-LSP is yet established. This operation, however, may lead to resource waste in the network. Notice that long FA-LSPs connecting far-off nodes are limited to be only reused by incoming LSP requests between remote endpoints. Hence, it appears more appropriate to separate the signaling functionality from the λ -LSP creation, so that λ -LSP placement can be decided based on network characteristics.

In the present paper, we set the routing metric of the already established FA-LSPs to be $\max(I, FA-LSP \text{ hops} - I)$, as described in [5]. Besides, the routing metric assigned to the unallocated data links spanning one single physical hop is set to 1. Aiming at better resource utilization, however, we dissociate λ -LSP establishment from network signaling functionality in the following way. Once the route from source to destination is calculated, the heuristic cost function $C_{FA}(H) = H [(1-p_H) + h / H]$ is applied to the route segments where connectivity is not yet existent, standing H for the number of hops of the yet to be created λ -LSP and p_H for the probability that any incoming demand in the network has a certain number of hops H .

The cost function provides us with the most appropriate λ -LSP configuration to optically connect the yet uncovered route segment. As will be later depicted by example, the term $(1 - p_H)$ encourage those λ -LSP lengths close to the average network distance, thus being more likely to be reused. The term $1/H$ identifies the use of O/E port pairs per hop, so that the larger the λ -LSP, the lower the use of O/E ports to connect its endpoints. The tunable h parameter fosters/penalizes the use of O/E ports in the network. Finally, the total cost is multiplied by H as long λ -LSPs need a higher number of unallocated data links. In this context, let us imagine that a new λ -LSP, which will afterwards act as FA-LSP, has to be established between a node-pair distancing 4 hops. Supposing that a 2 hops client LSP length is the most likely in the network, the combination $C_{FA}(2)+C_{FA}(2)$, that is, two λ -LSPs spanning each one two hops, could have lower cost than $C_{FA}(4)$, meaning one single end-to-end λ -LSP. Very short FA-LSP establishment (e.g., 1 hop) is also penalized $C_{FA}(H)$, due to the large amount of required expensive O/E ports, as well as the large amount of bypass traffic to be electrically processed.

2.2 Implementation work

The implementation of the FA-LSP functionality in the CARISMA test-bed has been made according to GMPLS standards. A new RSVP-TE object `LSP_TUNNEL_INTERFACE_ID` was proposed to be used when signaling a new FA-LSP in the PATH and RESV RSVP-TE messages [8]. The object contains two fields, that is, the FA-LSP identifier and the router ID.

In the event of a new λ -LSP to be set up, the head-end OCC must allocate an identifier for the interface associated to the yet to be created FA-LSP. Next, it originates an RSVP-TE PATH message containing a `LSP_TUNNEL_INTERFACE_ID` object filled with the selected local interface identifier, along with the local optical node identifier. When the PATH message arrives to the destination, the tail-end OCC must allocate an identifier for that FA-LSP end. This is called the remote FA-LSP interface identifier, which is reported back to the head-end within the RSVP-TE RESV message. As soon as the λ -LSP has been created, the head-end OCC advertises it as a forwarding adjacency by means of OSPF-TE. Being the FA-LSP bidirectional, it is also advertised by the tail-end OCC. All OCCs receiving the FA advertisement update its link state database adding a new link between the involved nodes. λ -LSP length may be limited to maximize resource re-utilization in the network. To achieve such purposes, λ -LSP establishment must be dissociated from the client LSP setup procedure, allowing in this way the establishment of several underlying λ -LSPs while signaling only one client LSP request. To permit this separation between client LSP and λ -LSP set-up, the head-end OCC evaluates the accumulated optical length, using *loose* hops in the ERO RSVP-TE object [10] when it decides to divide the whole route segment into two or more underlying λ -LSPs. The same mechanism is used also when an intermediate λ -LSP must be created. Every intermediate OCC receiving an RSVP-TE PATH message with a next hop set as *loose*, it must compute the next route segment possibly signaling a new λ -LSP.

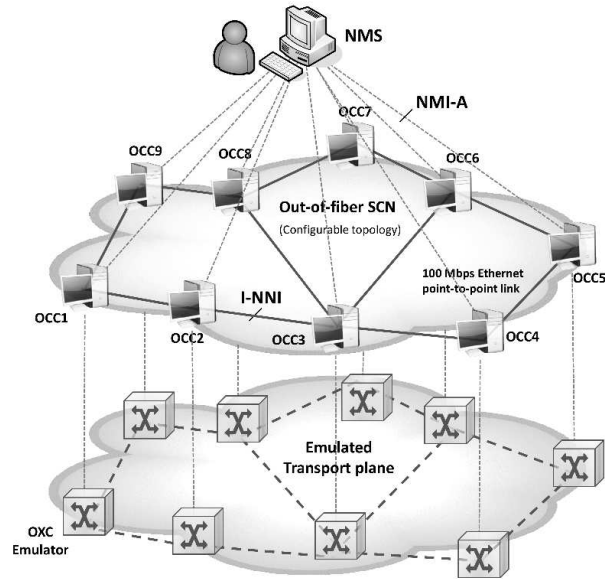


Figure 1: The ASON/GMPLS CARISMA Test-bed

3. EXPERIMENTAL RESULTS AND DISCUSSION

The FA-LSP experimental evaluation has been carried out over the ASON/GMPLS CARISMA test-bed [11], a configurable multi-topology Signaling Communications Network (SCN) running over Optical Cross-Connect (OXC) emulators, showing an out-of-fiber control plane architecture associated to the underlying transport plane. In the test-bed, OCCs are deployed by means of Pentium IV Linux-based routers at 2-GHz, so that each OCC implements the full GMPLS protocol set: RSVP-TE for signaling, OSPF-TE for routing and information advertisement and, finally, LMP for resource discovery and management [12]. In this work, we particularly build a 9-Node meshed SCN configuration to evaluate the proposed contributions as depicted in Figure 1. In such a scenario, each link carries 8 bidirectional wavelengths.

For the traffic characteristics, we consider that uniformly distributed client LSP requests arrive to the network following a Poisson process with mean Inter-Arrival Time (IAT) equal to $1/\lambda$. Besides, connection duration follows an exponential distribution with mean Holding Time (HT) set to $1/\mu$. In particular, the requested BW of all incoming client LSP requests is considered to be $1/4$ of the total wavelength capacity.

Figure 2 illustrates the obtained $C(H)$ function for the scenario under study. The bar graph plots the probability that an incoming client LSP request has a certain number of hops, assuming availability of resources through the shortest path. As seen, there is a 40% probability that an incoming request traverses 2 hops from source to destination. In contrast, only 5% of the incoming requests would traverse 4 hops. This validates our assumption in subsection 2.1, where we stated that by splitting very long FA-LSPs into shorter ones resources are much more likely to be reused. Values greater than the network diameter (i.e., 4 hops in our scenario) have $p_H = 0.0$ and are not depicted in the figure. To finally obtain $C_{FA}(H)$ we fix $h = 0.5$, as it provided the best network performance, while fulfilling our design criteria: $C_{FA}(2) + C_{FA}(2) < C_{FA}(4)$ and $C_{FA}(1) + C_{FA}(2) < C_{FA}(3)$.

Figure 3 plots the obtained client LSP blocking probability as a function of the offered load to the network ($\lambda/\mu = HT/IAT$) normalized to a value of 200. The curve *w/o FA-LSP* indicates the situation where no FA-LSP capabilities exist in the network, thus allocating a whole wavelength capacity for the incoming connection requests (even though they only request $1/4$ of the total wavelength capacity). Conversely, the curve *with FA-LSP* evaluates the improvement obtained when enhanced sub-wavelength provisioning flexibility is provided to the network by implementing GMPLS-controlled grooming actions.

As expected, significantly better resource usage is achieved when implementing FA-LSP capabilities in the network. For instance, if 1% client LSP blocking probability would have to be assured, a maximum load $L=0.05$ could be offered to a pure wavelength-routed optical network (i.e, *w/o FA-LSP*). Conversely, it could be further increased to $L=0.55$ when FA-LSPs are implemented. This $\Delta L = 0.5$ experimentally assesses the implementation of FA-LSP capabilities to automatically manage grooming actions in future transport networks, given the lack of pure wavelength-routed optical networks to allocate incoming sub-lambda client LSP requests. In fact, as a whole wavelength is allocated in per $1/4$ wavelength capacity client LSP request in the *w/o FA-LSP* situation, $3/4$ of the total network capacity is directly wasted. Qualitatively speaking, this approximately results in four times less carried traffic by the network.




Figure 2: $C_{FA}(H)$ function in the scenario under study

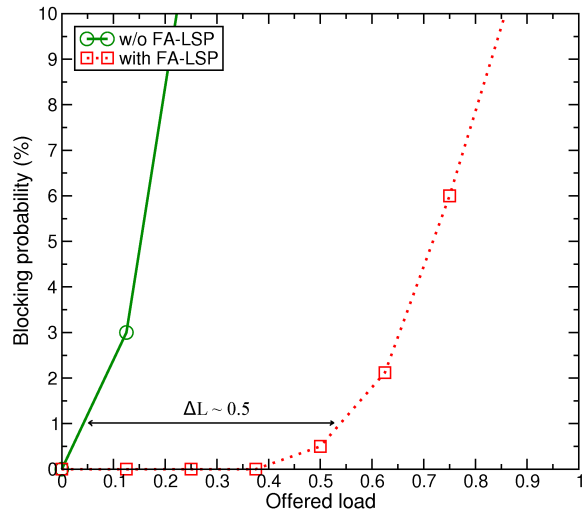


Figure 3: Client LSP blocking probability

4. CONCLUSIONS AND FUTURE WORK

This paper reported the implementation and experimental evaluation of the CARISMA Test-bed, a GMPLS-controlled grooming-capable network. To start with, the paper reviewed the Forwarding Adjacency (FA) entity, further elaborating on a FA creation cost function $C_{FA}(H)$ that dissociates the λ -LSP set up from the network signaling process, enhancing in this way the reuse of already established FA-LSPs. This function was afterwards particularized for the 9-Node scenario under study. This opened the presentation of the experimental results, which validated GMPLS-controlled grooming actions, drastically improving client LSP blocking probability compared to a pure wavelength-routed optical network scenario.

The evaluation presented in this paper concerns a single 9-Node ASON domain. Further work will extend the implementation of the FA-LSP functionality in larger network scenarios, including signaling and routing interworking issues in multi-domain multi-layer network environments. Particularly, it is our goal to assess not only the performance, but also the scalability of the GMPLS-controlled grooming and the proposed re-allocation approach as the network gets larger, even spanning more than a single domain.

ACKNOWLEDGEMENTS

The work presented in this paper has been partially supported by the Spanish Science Ministry through the project ENGINE (TEC2008-02634).

REFERENCES

- [1] ITU-T Rec. G.8080/Y.1304, "Architecture for the Automatically Switched Optical Networks", Nov. 2001.
- [2] E. Mannie, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, Oct. 2004.
- [3] K. Zhu, B. Mukherjee, "Traffic grooming in an Optical WDM Meshed Network", IEEE Journal on Selected Areas in Communications, Jan. 2002.
- [4] R. Dutta, G. Rouskas, "Traffic grooming in WDM networks: past and future", IEEE Network, Nov. 2002.
- [5] K. Kompella, Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", IETF RFC 4206, Oct. 2005.
- [6] J. Comellas, R. Martínez, J. Prat, V. Sales, G. Junyent, "Integrated IP/WDM routing in GMPLS-based optical networks", IEEE Network, Mar. 2003.
- [7] D. Katz, K. Kompella, D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", IETF RFC 3630, Sep. 2003.
- [8] L. Berger, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", IETF RFC 3473, Jan. 2003.
- [9] B. Rajagopalan, J. Luciani, D. Awduche, "IP over Optical Networks: A Framework", IETF RFC 3717, Mar. 2004.
- [10] K. Kompella, Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", IETF RFC 3477, Jan. 2003.
- [11] J. Perelló, E. Escalona, S. Spadaro, J. Comellas, G. Junyent, "Resource Discovery in ASON/GMPLS Transport Networks", IEEE Communications Magazine, Aug. 2007.
- [12] J. Lang, "Link Management Protocol (LMP)", IETF RFC 4204, Oct. 2005.