A Comparison of In-Fiber and Out-Of-Fiber GMPLS-Based Control Plane Configurations: Benefits, Drawbacks and Solutions

Jordi Perelló, Luis Velasco, Fernando Agraz, Salvatore Spadaro, Gabriel Junyent, Jaume Comellas
Optical Communications Group, Signal Theory and Communications Department
Universitat Politècnica de Catalunya, C. Nord D4-S107, Jordi Girona, 1-3, E-08034, Barcelona, Spain
e-mail: {jperello, luis.velasco, agraz, spadaro, junyent, comellas}@tsc.upc.edu

ABSTRACT
Current GMPLS standardization allows control plane deployment to be either physically or only logically separated from the data plane, describing out-of-fiber or in-fiber control plane configurations. It is widely accepted, that the out-of-fiber strategy provides enhanced flexibility, which derives from the natural separation of both planes. Nonetheless, such a configuration poses new challenges to protocol designers, aiming to provide the control plane with the resilience and automation required for evolving optical networks. In this paper, we discuss benefits and drawbacks of an out-of-fiber control plane configuration. Furthermore, we evaluate key GMPLS functionalities by experimental results over the ASON/GMPLS CARISMA Test-bed, thus giving insight into the applicability of GMPLS to future optical networks. The results on obtained performance are subsequently disclosed.
Keywords: GMPLS, control plane, in-fiber, out-of-fiber

1. INTRODUCTION

Traditionally, a one-to-one association exists between control and data channels in packet-switched networks, such as MPLS or IP networks. In MPLS-based networks, control and data packets are transmitted over the same medium, following the same LSP. More pronounced, in IP networks, there is even no explicit separation between control and data networks, since control packets are processed in the Layer-3 identically as end-user packets. This does not happen, however, in emerging intelligent optical networks, built following the ITU-T ASON architecture [1] and controlled by a GMPLS control plane.

As defined in [2], the GMPLS control plane can be separated from the data plane, even describing a different topology than the latter. This appears as fundamental to provide clean separation between control and data planes, so that the health of each one does not correlate to the health of the other. For instance, the control plane could be transmitted through the same physical fiber as the data channels but on a dedicated wavelength. Alternatively, it could be supported by a separated IP network connecting the remote endpoints of the fiber links under maintenance. In both situations, control and data plane forwarding states would be no more linked.

Because in packet-switched networks control and data packets share the same medium, and thus the same fate, protection/restoration of control channels can be handled over the protection/restoration of data channels. In this context, protection schemes as 1+1, 1:1 or 1:N commonly implemented in the data plane jointly deal with the protection of the control plane. Then, it makes no sense to distinguish between control and data plane recovery. Switchover times are, therefore, constrained to the most stringent recovery requirements (usually those of the data plane). Conversely, in the case that control and data plane become decoupled, they share no more the same resilience, as the control plane may fail independently from the data plane and vice versa. On the one hand, this poses new challenges to provide the control plane with the required resilience to fulfill emerging services resilience requirements. On the other hand, efficient protection and restoration can be provided to the control plane, as the applied recovery schemes are no more associated to those used in the data plane.

Besides the mentioned resilience considerations, further issues stem from the separation of both planes. First, automatic resource discovery procedures, which emerge as necessary in next-generation networks, get far more complicated. As long as the control and data planes share the same physical links, control and data plane topologies coincide. However, a physical separation of both planes may lead to asymmetrical topologies, where control channels may be established through intermediate nodes or IP networks. This may require extra work for protocol designers to achieve automatic bootstrap of control channels, as well as to achieve automatic data plane resource discovery. Contrarily, the physical separation of control and data planes eases fault localization procedures, critical in transparent optical networks where Loss-of-Light (LoL) alarms are propagated downstream (in terms of data flow) from the failure point on. Particularly, in the GMPLS protocol set, Link Management Protocol (LMP, [3]) takes advantage of the fact that, in the event of data plane failure, the control channel between endpoints of the failed link is still alive, enabling simple fault localization procedures.

This work is devoted to highlight such benefits and drawbacks of emerging control plane configurations. To this end, we give insight into key issues that come into light, supporting our conclusions on experimental results over the ASON/GMPLS CARISMA test-bed. Section 2 defines architectural concepts of the GMPLS control plane. Then, section 3 provides comparison of GMPLS in-fiber and out-of-fiber control plane configurations, as well as evaluates standard GMPLS fault localization against alternative in-fiber solutions.

This work was partially supported by both the Spanish Science Ministry through the project RINGING (TEC2005-08051-C03-02) and the i2CAT foundation through the project TRILOGY.

978-1-4244-2626-3/08/$25.00 ©2008 IEEE

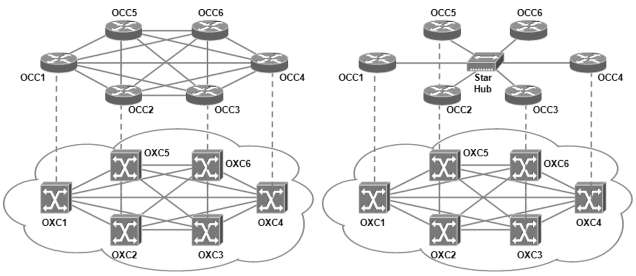

Figure 1. Symmetrical(left) and asymmetrical (right) control plane topologies.

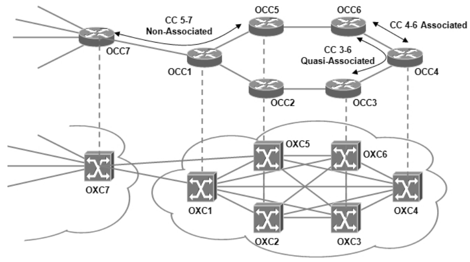

Figure 2. Associated, quasi-associated and non-associated control channels.

2. GMPLS CONTROL PLANE ARCHITECTURE

The purpose of the control plane in intelligent optical networks is to automate connection setup, as well as to provide resource discovery. In addition, some data plane protection and restoration mechanisms may be also supported on the control plane. Three different communication networks to support control and management functionalities are defined in ITU-T recommendation G.7712 [4]. The Management Communications Network (MCN) carries management data from network elements to management elements. The Signaling Communications Network (SCN) communicates control plane components for the transport of signaling messages. Both the MCN and the SCN could be understood as logical networks, which are supported over the Data Communications Network (DCN), that is, the IP-based optical control plane infrastructure. As a matter of fact, the MCN and the SCN can be either deployed over separate DCNs or over the same DCN. From now on, with the term "control plane" in ASON/GMPLS architectures we refer to the DCN supporting the SCN.

There are two possible control plane configurations. Control channels can be established either *in-fiber*, so that control and data planes share the same transmission medium, or *out-of-fiber* thus introducing a physical separation in between. In the former, two sub-configurations could be also differentiated, depending on whether the control channels are transmitted out-of-band, over a separated wavelength than the end-user traffic, or in-band along with the data channels (e.g., as in MPLS). Note, however, that the *in-fiber in-band* solution is not feasible in transparent optical networks, as lightpaths optically bypass intermediate nodes from source to destination. Hence, communication may not be possible between neighboring control plane nodes. Instead, in the *in-fiber out-of-band* configuration, point-to-point control communication is deployed on a different wavelength channel. Therefore, in this work we compare the *out-of-fiber* solution against the *in-fiber out-of-band* one.

As discussed in [5], from a topological perspective, the control plane can be deployed congruent with the corresponding data plane, thus following a *symmetrical* topology or, on the contrary, it can be deployed following an *asymmetrical* topology, describing control and data planes different topologies (Figure 1). Note, that the *in-fiber* configuration typically represents a *symmetrical* topology, as control channels are constrained to the same path as data channels. Contrarily, the freedom introduced with the *out-of-fiber* configuration fosters the deployment of *asymmetrical* topologies. Alternatively, ITU-T recommendation G.807 [6] classifies control channels depending on the way they are associated to data channels. Control channels directly connecting control plane nodes are named *associated* (CC 4-6 in Figure 2). Control channels connecting adjacent control plane nodes of the same network through intermediate control channels are called *quasi-associated* (CC 3-6 in Figure 2). Finally, control channels connecting control plane nodes of different networks through intermediate control channels are presented as *non-associated* (CC 5-7 in Figure 2). In this context, *symmetrical* control plane topologies are prone to be composed of *associated* control channels. Conversely, *asymmetrical* topologies usually lead to *quasi-associated* (or *non-associated*) control channels, where each control link supports several control channels.

3. PERFORMANCE COMPARISON OF IN-FIBER AND OUT-OF-FIBER CONTROL PLANES

In this section, we break down and further discuss major issues that appear when moving from *in-fiber* GMPLS control plane configurations to *out-of-fiber* ones.

3.1 Control plane failure detection and recovery

As long as control and data planes become decoupled (*in-fiber out-of-band* or *out-of-fiber*), additional fault detection and recovery mechanisms for the control plane must be provided. Failure detection may follow different approaches depending on whether the control plane is configured *in-fiber out-of-band* or *out-of-fiber*. The former may benefit from fault detection mechanisms used in the data plane (e.g., LoL alarms). However, for the latter alternative mechanisms must be provided. To this end, LMP implements a Hello-based keep-alive protocol to maintain the connectivity of the control channels [3]. In the case that no *Hello* message is received along a pre-determined time interval, control channel is declared down. While fault detection in the data plane is typically fast (lower than 50 ms), *Hello* intervals in LMP should be tightly dimensioned to avoid major protocol disruptions due to control plane failures, which could even affect data plane connections [7].

Because the GMPLS control plane is IP-based, IP layer rerouting is the natural recovery mechanism. Therefore, as long as the failure has been detected, the failure state should be notified to the control plane routing protocol (e.g., OSPF), which starts IP re-convergence actions to avoid the failed link.

3.2 Resource discovery and link connectivity verification

A procedure is introduced in LMP to verify the connectivity of the data links and to perform data plane resource discovery between any pair of nodes (initiator and receiver). Basically, it consists of sending in-band *Test* messages [3] over the data link whose connectivity is under verification. These messages contain the Node ID of the verification initiator, as well as the interface ID over which the messages are sent. As a result, whether the receiver detects *Test* messages on any incoming port, it firstly deduces that connectivity exists to the initiator through such incoming port. Furthermore, it concludes on the remote interface ID associated to the incoming port. Next, the receipt of the *Test* messages is acknowledged to the initiator by means of a *TestStatusSuccess* [3] message, which contains the receiver Node ID along with the interface ID on which *Test* messages have been received. This permits the initiator to learn interface associations and the remote node to which the data link is connected. Nonetheless, if the receiver does not receive any *Test* on any incoming port during a time interval, it returns a *TestStatusFailure* [3], reporting that no connectivity exists between both nodes on that data link.

Although control channels can be established *out-of-fiber*, *Test* messages must be transmitted *in-fiber in-band* to allow connectivity verification. Therefore, O/E/O capabilities must exist on each incoming interface, requirement not easily supported by transparent optical networks. In [8], we proposed an alternative connectivity verification mechanism to facilitate its adoption by transparent optical networks and to completely decouple control and data planes. In such a mechanism, *Test* messages are transmitted *out-of-fiber* (or *in-fiber out-of-band*), and connectivity is detected by sending and detecting light in the remote endpoints of the data links. Particularly, if the initiator desires to verify the connectivity of a given data link, it starts sending light over it. Then, it subsequently sends the *Test* message to the receiver over the control plane. Upon *Test* reception, the receiver checks the existence of light on any incoming port. If so, connectivity exists to the initiator through that incoming port. Otherwise, no connectivity exists between them on the data link under verification.

A similar problematic appears when trying to automatically set up (i.e., bootstrap) control channels in *out-of-fiber* control plane configurations. Specifically, the problem arises when both control and data plane adjacencies must be automatically configured. In this context, a natural discovery sequence would stay as follows. First, data plane adjacencies would be discovered. Next, control channels would be established between the related controllers of data plane neighbors, enabling the maintenance of the resources in between. Unfortunately, to allow LMP connectivity verification procedures, control channels must be previously established between data plane neighbor pairs. As a matter of fact, LMP does not propose any method to automatically learn control plane adjacencies, letting them to be manually configured. Some works [9] tried to join control and data plane resource discovery by sending bootstrap messages *in-fiber in-band*. As said before, this may be prohibitive for transparent optical networks. Taking advantage that in *symmetrical out-of-fiber* control plane configurations both control and transport planes coincide, in [8] we dealt with control channel bootstrap in such simplified scenario by multicasting *Config* messages through the control interfaces, with the IP Time-to-Live (TTL) field set to 1. Alternatively, we also proposed in [10] a possible solution to tackle control channel bootstrap in *asymmetric out-of-fiber* configurations, which was based on sending and detecting light in the data links endpoints.

3.3 Data plane fault localization and management

In the event of a failure, the optical transparency in the data plane yields to downstream propagation of LoL alarms, as all downstream nodes from the failure point on detect LoL on its incoming ports. To suppress multiple alarms stemming from the same failure, localization procedures should be implemented in transparent optical networks. One of the main benefits of the *out-of-fiber* control plane configuration is that control channels remain alive despite data plane failures. Taking advantage of such particularity, LMP implements a simple fault localization procedure. Note, however, that the applicability of such procedure is restricted to the *out-of-fiber* configuration. On the contrary, fault localization in *in-fiber out-of-band* configurations usually relies on hardware-based solutions.

Aiming to validate the applicability of LMP as a fault localization protocol, we experimentally compare failure localization times obtained with standard *out-of-fiber* LMP localization procedures against a hardware-based solution that fits either *in-fiber out-of-band* or *out-of-fiber* control plane configurations. Specifically, the hardware-based solution is based on sending and detecting a pilot tone on the endpoints of each link, thus having whole link granularity. Contrariwise, LMP provides wavelength granularity, serving for path protection/restoration purposes. Evaluation has been conducted on the ASON/GMPLS CARISMA test-bed [10].

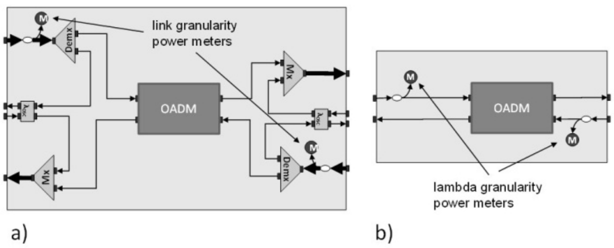

Figure 3. Additional hardware needed for the Optical pilot tone (a) and for LMP-based fault localization (b).

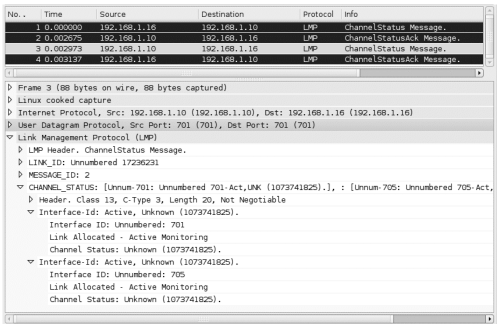

Figure 4. LMP Fault Localization capture

The optical pilot-tone based procedure uses an additional optical transponder to send and receive an optical tone (e.g. in 1300 nm) which is multiplexed with the bundle of DWDM-multiplexed signals (Figure 3a). An optical power meter monitors the incoming pilot tone power level. Whether a link is not affected by a failure, optical power must always be received at each end of the link. In the event of receiving out-of-bound power levels, the node sends a LoL notification through the Connection Controller Interface (CCI, [1]), meaning that a link failure has occurred. Let us denote t_{config} as the time needed to send a message from the Optical Add Drop Multiplexer (OADM) to its Optical Connection Controller (OCC) and $t_{localiza}$ as the time from the out-of-bound power level detection in the OADM until the failure is localized. Using this procedure, we can express this time as $t_{config} = t_{localiza}$. In our implementation t_{config} = 1 ms. Then, in the optical pilot tone based procedure, the time to localize a failure in the control plane remains constant to 1 ms, independent of the network topology. Note that optical power meters in this solution have monitoring sweep times of 10 μs, thus assuring very fast detection.

The LMP-based localization procedure also uses optical power meters on each input port to monitor the incoming optical power level (Figure 3b). In this case, in order to obtain better granularity, we use an arrayed power meter which is able to monitor every single wavelength in parallel. Upon the reception of out-of-bound power level in either one or multiple wavelengths, the node sends a LoL notification for these wavelengths though the CCI interface. Upon receiving the LoL notification, the OCC should determine whether the failure is in the local link or in any upstream link (Figure 4). To this end, it sends a *ChannelStatus* message [3] to its upstream neighbor with the list of individual failed data links (or the complete link if all data links have failed). Upon reception of a *ChannelStatus* message, the neighboring LRM process checks the status of the data links associated to the failed ones through a local lightpath connection. If the reception of the associated data links is OK, the failure has been localized in this link; on the contrary, another upstream link may be responsible for the failure. Finally, a *ChannelStatus* message is sent to the downstream neighbor with the status of the data links. Let us denote t_{link} as the propagation delay in each link, and t_{LMP} as the time to process a single LMP message. Using this procedure, localization time stays as $t_{localiza} = t_{config} + 2 * (t_{link} + t_{LMP})$. In our implementation t_{LMP} = 0.2 ms and t_{config} = 1 ms. However, t_{link} depends on the length of the links (L). For example, if we consider metropolitan networks with L = 30km, t_{link} = 0.15 ms, whereas in larger networks with L = 100 km, t_{link} = 0.5 ms. As a result, $t_{localiza}$ ranges from 1.7 ms to 2.4 ms in such scenarios. Besides, the arrayed optical power meters used with this solution have a higher monitoring sweep time of 1 ms.

As seen, LMP detection and localization times are generally higher than the ones achieved by the optical pilot-tone based procedure. However, even in large network topologies, their contribution to the total protection time (e.g., 50 ms) could be assumed as marginal. This strongly leverages the applicability of LMP. Note, that the pilot tone-based procedure becomes essentially limited to link recovery (due to its coarse granularity). Nonetheless, the finer granularity of LMP makes it applicable not only to link recovery but also to path recovery, thus giving support for a broader protection/restoration scope.

REFERENCES

[1] ITU-T G.8080: Architecture for the automatically switched optical network, 2003.
[2] E. Mannie: Generalized Multi-Protocol Label Switching Architecture, *IETF RFC 3945*, 2004
[3] J. Lang: Link Management Protocol (LMP), *IETF RFC 4204*, 2005.
[4] ITU-T G.7712/Y.1703: Architecture and Specification of Data Communications Network, 2003.
[5] K. Shiomoto: Control Plane Architecture in GMPLS Networks, *IETF draft*, 2004.
[6] ITU-T G.807: Requirements for Automatic Switched Transport Networks (ASTN), 2001.
[7] J. Perelló: An Analytical Study of Control Plane Failures Impact in GMPLS Ring Optical Networks, *IEEE Communications Letters*, 2007.
[8] J. Perelló: Resource Discovery in ASON/GMPLS Transport Networks, *IEEE Comm. Magazine*, 2007.
[9] J. Lang: Control Channel Bootstrap for Link Management Protocol, *IETF draft*, 2003.
[10] S. Spadaro: The CARISMA ASON/GMPLS Network: Overview and Open Issues, *ICTON 2007*, 2007.