Control plane protection using Link Management Protocol (LMP) in the ASON/GMPLS CARISMA network

Jordi Perelló, Eduard Escalona, Salvatore Spadaro, Fernando Agraz, Jaume Comellas, Gabriel Junyent

Optical Communications Group, Signal Theory and Communications Dept., Universitat Politècnica de Catalunya, C. Nord D4-S107, Jordi Girona, 1-3, E-08034, Barcelona (Spain)
{jperello, escalona, spadaro, agraz, comellas, junyent}@ tsc.upc.edu

Abstract. In the ITU-T ASON architecture, the control plane is responsible for providing intelligence to the network. The GMPLS paradigm pleads for a separation between the control plane and the forwarding plane. If the control plane is deployed disjoint from the forwarding plane, recovery mechanisms to ensure its proper operation are required. In this paper, on one hand, a quasi-associated mode backup control channel proposal is compared with a traditional associated 1:1 protection. On the other hand, extensions to LMP defined by the IETF are presented and evaluated to address both control channel and nodal failure recovery. The merits of the proposals are assessed by experimental results.

Keywords: ASON, GMPLS, LMP, protection

1. Introduction

The increasing utilization of the Internet, the emerging applications that require large bandwidth, in conjunction with the nowadays high speed access networks, have put the current transport infrastructure in a tight spot. While designed to support circuit-based traffic, it shows its inefficiencies, mainly due to its static bandwidth provisioning when carrying IP and Ethernet based data traffic. It is proved that Automatic Switched Optical Network (ASON) architecture [1], defined by the International Telecommunications Union (ITU-T), has become a hopeful possibility to support the current data traffic explosion, which requires a high degree of flexibility. Among multiple functionalities, the ASON architecture accomplishes the requirement of fast, dynamic and flexible end-to-end bandwidth provisioning. This architecture relies on three well separated planes: an all-optical transport plane over which the light paths/connections are established, a control plane responsible for creating, maintaining and deleting the requested connections, and a management plane with a whole network view, capable for requesting, supervising and tearing-down light paths as well as for managing both the transport and control plane. To meet the above mentioned requirements, the key entity in the ASON architecture is the control plane, which provides the necessary intelligence to the network. It supports the required routing and signaling information for dynamically creating the requested connections. The Generalized Multi-Protocol Label Switching (GMPLS) protocol set [2], defined by the Internet Engineering Task Force (IETF) in concordance with the ITU-T, arose as the preferred technology to implement control plane functions. GMPLS is an extension of the set of protocols designed for the MPLS technology and encompasses time-division (e.g. SONET/SDH, PDH, G.709), wavelength as well as spatial switching. The main GMPLS protocols are: Resource Reservation Protocol with Traffic Engineering Extensions (RSVP-TE) [3], Open Shortest Path First with Traffic Engineering Extensions (OSPF-TE) [4] and Link Management Protocol (LMP) [5][6]. Connection signaling tasks are performed by RSVP-TE. The OSPF-TE protocol is used to flood the state of all the node outgoing data links to all the other network nodes. In order to establish a connection, foremost a route to reach the destination node is calculated using the link state information provided by OSPF-TE. Afterwards, the request is forwarded by RSVP-TE to all the nodes involved in that connection. Upon reception of RSVP-TE Path/Resv messages, the required resources are reserved.

The whole of the control channels constitute the Data Communications Network (DCN) [1], whereby routing and signaling information is transmitted. LMP is defined by the IETF (hereafter standard LMP) as a new protocol with multiple functionalities, such as the management of the control channels between

The work reported in this paper was supported in part by the Spanish Science Ministry through the Project "Red Inteligente GMPLS/ASON con Integración de Nodos Reconfigurables (RINGING)", (TEC2005-08051-C03-02)

neighbors. Other functionalities of LMP are the correlation of the logical resources mapped over the physical existing resources between neighbors, link discovery, and fault isolation procedures.

This paper presents an enhanced control channel protection scheme in order to optimize the required control network resources. Moreover, some extensions to standard LMP are presented to properly perform control channel and nodal failure protection. Our proposals have been implemented and evaluated in the ASON/GMPLS CARISMA network, which relies on an out-of-fiber control plane.

The remainder of the paper is organized as follows: firstly the CARISMA network is described in Section 2 and the LMP protocol is overviewed in Section 3. Section 4 presents the proposed protection scheme and the results of its performance compared to the standard LMP. In Section 5, some extensions to the standard LMP are discussed in order to successfully perform control channel and nodal failure protection while minimizing the required control network resources. Finally, Section 6 concludes the paper.

2. The ASON/GMPLS CARISMA network

The CARISMA project [7] was initiated in 2002 as an initiative to build a high performance Wavelength Division Multiplexing (WDM) based network to be used as a field-trial for the integration and evaluation of the current emerging innovative technologies. It is intended to provision bandwidth on demand while ensuring Quality of Service (QoS) between IP networks.

The CARISMA network (Fig. 1) implements the ASON architecture. Its transport plane is formed by three OADM capable optical nodes. These nodes are connected through two unidirectional fibers (working and protection fibers respectively) forming a dual ring topology. Each OADM is able to insert up to four WDM channels by means of transceivers (two of them tunable and two fixed) and to extract four channels of the twelve available in the ring. Each WDM channel is transparent to 2.5 Gbit/s and at least three of them to 10Gbit/s. The distance between nodes is about 35Km far, so the total ring length is more than 100Km.

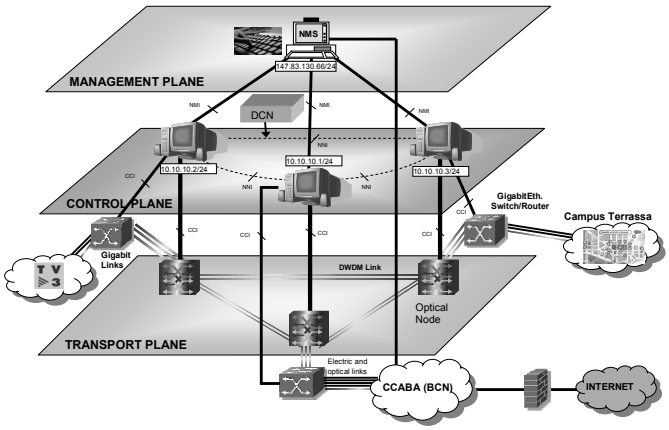

Fig. 1. The CARISMA network Architecture

The CARISMA network control plane has been deployed as an IP out-of-fiber network. Specifically, three Optical Connection Controllers (OCCs) are implemented using Linux-based routers, which run the GMPLS protocols. The three deployed OCCs are interconnected through Ethernet point-to-point links. The GMPLS paradigm intends to clearly separate the control plane from the data plane. Control channels can be in-fiber in-band, in-fiber out-of-band or out-of-fiber out-of-band. In the latter case, as in the CARISMA network, control plane liveliness is not associated to the data plane one, so control plane communication can be maintained alive upon a data plane failure and vice versa. Therefore, fault detection and protection mechanisms in addition to those existing in the data plane must be implemented also in the control plane. Since there is no association between the control channels and the data channels (e.g., as in MPLS), control channel protection can not be resolved using data plane protection mechanisms. In the CARISMA network we have implemented the LMP protocol to maintain both control channel connectivity and link properties between two data plane adjacent neighbors.

Finally, the CARISMA network management plane is formed by the Network Management System (NMS), implemented as a web application facilitating network administration through Internet.

3. Link Management Protocol overview

In order to enable the communication between nodes for signaling, routing and link management purposes, control channels must be established between any pair of nodes. The LMP protocol has been defined to fulfill control channel management and also to perform additional functionalities. The four functionalities proposed to be done by LMP are: control channel management, link property correlation, link connectivity verification and fault management. The first two are mandatory when implementing LMP, whereas the rest are optional.

Control channel management is related to two procedures, namely the control channel establishment and the maintenance between LMP neighbors. Specifically, a hello-based keep-alive mechanism is used to maintain control channel connectivity. These procedures begin with a negotiation phase, where the control channel is established and the keep-alive mechanism intervals negotiated. The use of a keep-alive mechanism takes crucial importance when lower-layer mechanisms are not able to detect control channel failures (e.g., out-of-fiber control plane). Hello messages are transmitted every *HelloInterval*. If no hellos are received in a *HelloDeadInterval*, control channel connectivity is declared lost.

Link property correlation deals with the synchronization of the properties of the defined TE links between adjacent neighbors, where a TE link is a logical aggregation of various data links defined between neighbors. In fact, those properties include both the TE link local and remote identifiers and the characteristics of all the data links contained in that TE link. This process is achieved by sending *LinkSummary* messages to a neighbor, which contain all the properties referent to a TE link towards that neighbor. The information contained in a *LinkSummary* message can be agreed or disagreed by responding with a *LinkSummaryAck* or a *LinkSummaryNack*. If the information comprised in a *LinkSummary* message is set to non negotiable, it is forced to be accepted. Link property correlation procedures must be done before a TE link is considered ready to transmit traffic. Furthermore it can be periodically performed.

Link connectivity verification is required to test the physical connectivity of the data links, and also to dynamically learn the TE link and data link ID associations, so it can be used for link discovery purposes. On the other hand, fault management is intended to be used to isolate data link and TE link failures. It becomes greatly useful if physical circuits are established upon an all-optical transport plane. In such environment, Loss of Light (LoL) fault detection mechanisms do not apply properly, since LoL alarms are detected by all the downstream nodes from the point where the failure has occurred. These two mechanisms are out of the scope of this paper.

4. Protection scheme for control channel failure

The standard LMP does not focus on any control channel protection schemes. In fact, in [5], control channel protection schemes such as 1+N (i.e., sending signaling and routing information through one or more control channels towards the same neighbor at the same time) or dedicated 1:1 (i.e. having a standby control channel waiting to become active upon working control channel failure (Fig. 2a)) are just mentioned. In this Section, we focus on how it is possible to take advantage of the LMP control channel management procedure for the out-of-fiber control plane protection. Generally speaking, according to [8], three kinds of control channels can be established: associated, quasi-associated and non-associated. Associated control channels directly interconnect two physically adjacent neighbors, whereas quasi-associated and non-associated control channels indirectly interconnect two physically adjacent nodes through a pre-determined route or following an undetermined route respectively.

The out-of-fiber CARISMA network control plane is based on a bidirectional ring topology. Therefore an alternative route which avoids a determined failed control channel is always available. In this context, establishing a disjoint quasi-associated backup control channel that surrounds the affected control channel is really advantageous. This scheme (Fig. 2b) is better than the associated mode possibilities (e.g., 1+1 protection, 1:1 dedicated protection) in terms of both resource and computational cost savings. In fact, on one hand, if more than one active control channel between neighbors is used, redundant control traffic packets have to be sent over the redundant control channels. Moreover, the redundant packets have to be discarded upon reception. Although this solution presents a nearly zero switching time, its required packet overhead, which increments nodal computational cost, makes it inappropriate to be applied to protect the DCN, where traffic recovery times are not as restrictive as in the transport plane.

On the other hand, dedicated 1:1 protection scheme does not require the computational cost as the previous option. However, it increments the number of resources required to implement control plane

protection in contrast with the quasi-associated scheme which is proposed in this paper. Fig. 2 presents the two evaluated protection schemes.

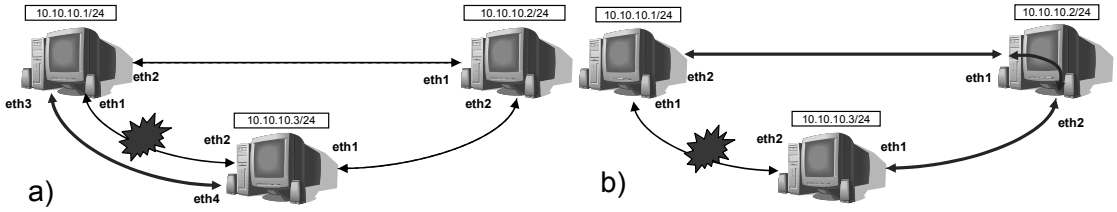

Fig. 2. Control channels protection schemes: a) Dedicated 1:1 backup control channel, b) Proposed control channel scheme: the backup channel is established through the alternative disjoint to the failure route.

This paper compares the here proposed quasi-associated protection scheme and the dedicated 1:1 one. The comparison has been done in terms of protection switching time, which is the time required to set up the backup control channel after the working control channel failure detection. It includes the backup control channel negotiation phase and the working control channel interface shutdown. Such protection switching time results have been experimentally obtained by using different LMP hello intervals. Choosing the correct hello intervals depends on the control plane protocols which are running. Upon a control channel failure, too high hello intervals can result in blocked connection requests, outdated node Traffic Engineering Databases (TED) or even RSVP-TE state loss. *OSPF Hello* messages and RSVP-TE retransmission times are in the order of tens of seconds, so a 5 second hello interval seems to be sufficient. Nevertheless, in order to minimize blocked connection requests and Link State Advertisement (LSA) losses, lower hello requests are required. Fig. 3 illustrates the obtained control traffic switching times function of the used HelloInterval value. Each point has been obtained as the mean of a statistical relevant number of results.

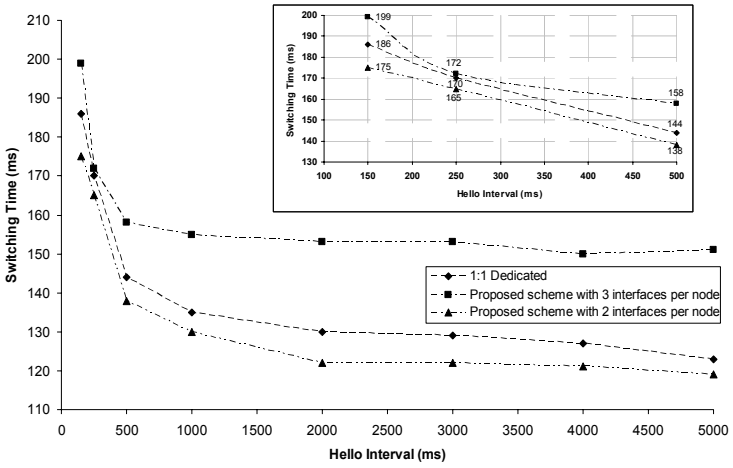

Fig. 3. Obtained control traffic switching times.

It can be seen that control channel traffic switching times increase as hello intervals decrease. The same happens for node overload. The proposed protection scheme has lower control channel switching times than the dedicated 1:1 one. This is due to the fact that the nodes which implement dedicated 1:1 protection have to maintain an additional interface, causing an increased computational cost. To verify it, control channel switching times obtained with the proposed scheme but in a scenario where each node maintain not two but three interfaces have been also included in Fig. 3. The LMP standardization proposes 150 ms for *HelloInterval* and 500 ms for *HelloDeadInterval*. However, values of 500 ms and 1500 ms for them are used in the CARISMA network. While avoiding signaling and routing operation to be disrupted, they involve lower CPU costs, reduce the traffic over the DCN and they are totally applicable in a metropolitan environment with low incoming call volumes.

Using the proposed intervals, we obtain significantly lower control plane restoration times compared with the tens of seconds of IP dynamic routing, and even with in-fiber management solutions, such as the one proposed in [9]. This is due to the fact that, since the control plane is decoupled from the transport plane, no transport plane configurations should be modified upon a control plane failure, which spares TED updates.

5. LMP extensions for control channel and neighbor node controller failures

In this Section, some extensions to the standard LMP to be used for both control channel and neighbor node failure situations are presented and evaluated. The standard LMP does not provide specific mechanisms for control plane network protection. On one hand, as above discussed, redundant control channel connectivity has to be avoided. On the other hand, criterions to be applied prior to LMP graceful restart procedures to properly distinguish nodal from control channel failures, take crucial importance.

According to the standard LMP, when no hello messages are received in a *HelloDeadInterval*, the node assumes that the working control channel is down. The node then set both the working and the backup control channel, which was in a standby state, to the negotiation state (i.e., the node sends *Config* messages towards its neighbors). Maintaining the failed working control channel in the negotiation state has its pros and cons. In fact, allowing the automatic control channel re-establishment, once the failure has been repaired, entails, for one hand unnecessary node controller computational costs and, on the other hand, redundant control channel connectivity.

To overcome both the computational cost and redundant connectivity problems, we propose a mechanism, which is an extension of the standard LMP, in order to tear down the backup control channel upon working control channel re-establishment. It uses the *ControlChannelDown* flag functionalities [5], which actually allow to gracefully take down a certain control channel. Fig. 4 shows the state diagram of our proposal. We consider three control channel states: Up, Going Down and Down [5]. The Up state is the operational state, wherein the hello-based keep-alive mechanism is performed. On the contrary, the Down state is the initial state, where the control channel is on standby and no attempts to bring up the control channel are made. A control channel passes from the Up to the Going Down state when an administratively control channel tear down is desired. In this state, the node sets the *ControlChannelDown* flag to 1 in all messages it sends.

When the working control channel connectivity is re-established after the failure is repaired, the event which indicates the first received valid hello message (*evHelloRcvd*), meaning a successful negotiation phase, is caught. Upon this event, the backup control channel passes from the Up state to the Going Down state, since its functionality is no more needed due to the fact that the working control channel is again fully operating. If a message is received with the *ControlChannelDown* flag set to 1 or no messages with this flag set to 1 are received in a *HelloDeadInterval*, the backup control channel is considered down. Fig. 4 shows that proposed functionality.

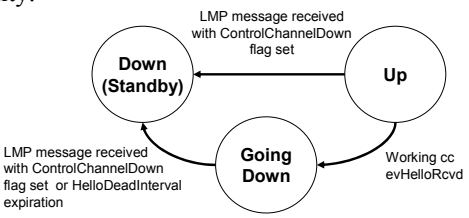

Fig. 4. Proposed backup control channel performance state diagram to avoid redundant control channel connectivity

The remainder of this section deals with LMP session recovery after a node failure. It can be possible that the failed node takes long to restart, making useless the negotiation state of the backup control channel, since the working control channel is healthy and has to be re-established once the failure is overcome. When a neighbor node controller failure occurs, LMP graceful restart mechanisms [5] should be applied to re-synchronize the properties of the defined TE links between adjacent neighbors.

In order to avoid this excess of computation costs related to maintaining the backup control channel into a negotiation state upon a neighbor node controller failure, we propose a method so the neighbor nodes detect if the failure is a link failure or a controller failure prior to LMP graceful restart procedures. When the node initiates the backup control channel negotiations, a timer set to three times *ConfigRetryInterval* is initiated. If no *Config*, *ConfigAck* or *ConfigNack* messages are received before this timer expiration, a neighbor controller failure is considered. Then, the backup control channel is set again to standby and the *Config* messages will just be sent through the working control channel. This is done until the reception of a *Config* message from the failed neighbor with the Restart flag set to 1. In that way, this proposed mechanism allows nodal failure distinction from control channel failure before the restart and reduces nodal computational costs by removing the backup control channel negotiation state. The proposed performance is shown in Fig. 5. The *Config* retry interval and *ConfigDeadInterval* used in the CARISMA network have been 500 ms and 1500 ms respectively.

To complete the LMP session restoration upon a neighbor controller failure, a *LinkSummary* message with no negotiable information is sent towards the restarted node for every TE link established between them. Graceful restart procedures avoid TE link and data link parameter misconfigurations once the node

is restarted. LMP session restoration times function of the defined data links (i.e lambdas) between neighbors are also depicted in Fig. 5, where the obtained points are the mean of a statistical relevant number of results. They are measured since the neighbor node controller is restored, to the ending of LMP graceful restart procedures, obtaining values close to 1,5 s.

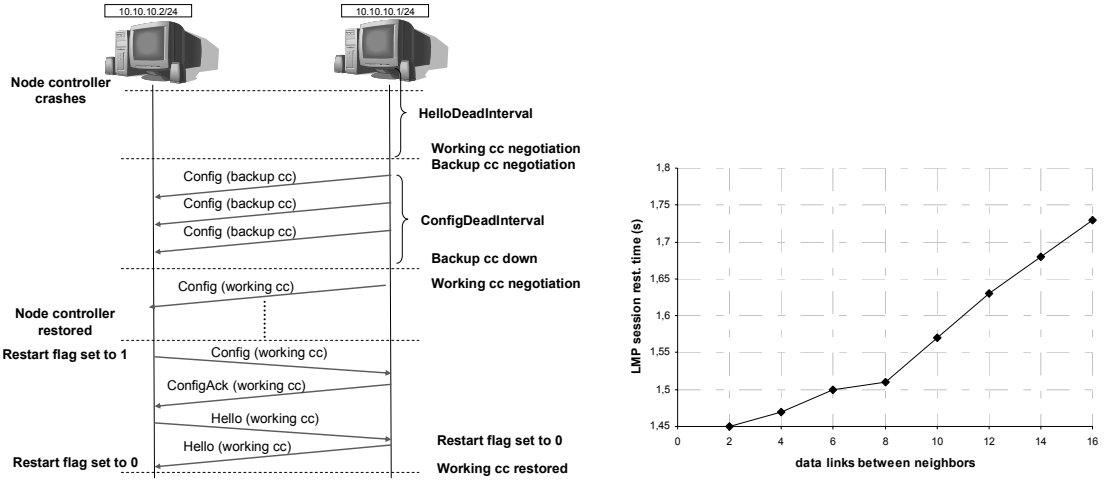

Fig. 5. Proposed method for LMP session restoration and LMP session restoration times function of the established data links between neighbors.

6. Conclusions

This paper presents an alternative backup control channel protection scheme in order to minimize the required resources to perform DCN protection in the out-of-fiber control plane of the ASON-based CARISMA network. Its functionality in terms of control traffic switching time has been evaluated and compared with the 1:1 dedicated protection. The obtained results not only show the feasibility of our proposal, but also reflect improved control traffic switching times in a metropolitan environment (with a low number of nodes and links) compared with the dedicated option.

Some extensions to the standard LMP have been proposed in order to properly perform both control channel failure recovery and nodal LMP session recovery optimization. Such extensions allow, on one hand, to avoid redundant control channel connectivity and, on the other hand, to differentiate between control channel and node controller failures. The latter implies the reduction of the computational costs for the node controllers. Upon nodal failure recovery, LMP graceful restart procedures have to be performed to re-synchronize the state of the TE links defined between neighbors.

The experimental results demonstrates that LMP with some extensions can be a useful way for providing control plane protection in ASON based out-of-fiber control plane environments.

References

1. ITU-T Recommendation G.8080: Architecture for the Automatically Switched Optical Network (ASON), 2001.
2. Mannie, E. (ed.): Generalized Multi-Protocol Label Switching Architecture. IETF RFC 3945, 2004.
3. Berger, L. (ed.): Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions. IETF RFC 3473, 2003.
4. Katz,D., Kompella, K., Yeung,D.: Traffic Engineering (TE) Extensions to OSPF Version 2. IETF RFC 3630, 2003.
5. Lang, J. (ed.): Link Management Protocol (LMP). IETF RFC 4204, 2005.
6. Fredette, A.(ed.), Lang, J. (ed.): Link Management Protocol (LMP) for Dense Wavelength Division Multiplexing (DWDM) Optical Line Systems. IETF RFC 4209, 2005.
7. CARISMA (Conexión y acceso a RedIRIS2 mediante anillo óptico multicanal) Project, http://carisma.ccaba.upc.edu
8. Young, K.: Requirements for the Resilience of Control Plane. IETF Internet draft draft-kim-ccamp-cpr-reqts-01.txt, 2005.
9. Muñoz, R., et al. : Experimental GMPLS fault management for OULSR transport networks, OSA/IEEE Optical Fiber Communications/SPIE National Fiber Optic Engineers Conference (OFC/NFOEC 2005).