

# Transport Plane Resource Discovery Mechanisms for ASON/GMPLS Meshed Transport Networks

Jordi Perelló, Eduard Escalona, Salvatore Spadaro, Jaume Comellas,  
and Gabriel Junyent

Optical Communications Group (GCO), Signal Theory and Communications Department  
Universitat Politècnica de Catalunya (UPC), Jordi Girona, 1-3, 08034 Barcelona, Spain  
{jperello, escalona, spadaro, comellas, junyent}@tsc.upc.edu

**Abstract.** The reduction of both deployment costs and operating expenses has emerged as a critical goal of next-generation optical networks. In this context, as each node maintains a large number of input/output ports (i.e. wavelengths, fibers, time slots ...), automation is needed to reduce long and tedious manual node configurations, thus easing network management and minimizing resource misconfiguration probability. In this paper, we present two novel mechanisms which permit GMPLS Link Management Protocol to deal with the resource discovery of all-optical meshed transport networks. The subsequent evaluation demonstrates their applicability over large meshed network topologies with considerably high nodal degrees.

**Keywords:** GMPLS, LMP, Automatic Resource Discovery.

## 1 Introduction

The proliferation of Internet Protocol (IP) technology, coupled with the large amount of bandwidth offered by recent advances such as Wavelength Division Multiplexing (WDM), is moving towards IP over WDM to be the preferred solution to implement next-generation optical networks. This way, supported by extraordinary technology innovations (i.e. ROADMs and OXCs), a migration from legacy SONET/SDH ring architectures to all-optical meshed ones is expected to be carried out, mainly due to their improved bandwidth efficiency and scalability. Moreover, the advent of a distributed control plane, implemented by means of GMPLS related protocols [1], emerges as a promising candidate to provide next-generation networks with the automation required to fulfill new applications requirements [2].

Automatic resource discovery is understood as the process of a node automatically finding the working links towards its transport plane neighboring nodes. This process avoids long and tedious manual node configurations and minimizes resource misconfiguration probability, which can result in high cost savings for network operators. A first approach to the required discovery procedures for the ASON architecture is introduced in [3], where their applicability over either the control or the transport plane is differentiated. It is noteworthy that their description is there carried out in an abstract way. With the introduction of GMPLS, control plane IP adjacency

discovery is delegated to legacy OSPF operation [4], whereas the discovery of the transport plane port mappings is intended to be done by the Link Management Protocol (LMP) [5]. However, to achieve its purposes, LMP requires optical nodes to transmit in-band information, so requesting link termination on each incoming port. This issue, in all-optical networks, implies additional hardware, high complexities or even signal losses.

In this paper, we propose and compare two different resource discovery mechanisms, which enable LMP to cope with the automatic discovery of all-optical meshed networks. These mechanisms have been called *Parallel Discovery* (PD) and *Concurrent Discovery* (CD). Specifically, PD extends the mechanism for ring topologies previously presented and experimentally evaluated over the ASON/GMPLS CARISMA network in [6], to be also applicable over meshed network topologies. In turn, CD is introduced to optimize PD, permitting to concurrently verify all neighboring nodes at a time. Their ease of adoption becomes a key issue, since no extensions on standard LMP messages are required.

The rest of this paper is organized as follows. In the following section, we explain the current LMP standardization regarding link connectivity verification and resource discovery. Section 3 introduces an alternative method to resolve transport plane interface mappings, enabling its applicability over all-optical networks. Section 4 depicts both proposed mechanisms, whose performances are evaluated by means of simulation results in Section 5. Finally, Section 6 draws up some conclusions.

## 2 Verifying Link Connectivity with LMP

Within the GMPLS framework, LMP has been newly introduced to perform several functionalities. Mainly, LMP is responsible for maintaining the connectivity between control plane neighboring nodes (i.e. Optical Connection Controllers in [3], OCCs), for correlating the properties of the existent data links between nodes, for performing resource discovery and even for providing fault isolation capabilities upon lightpath failures. While the implementation of the former two functionalities is mandatory in LMP, the implementation of the remaining is left to be optional. In this section, we provide a brief introduction to standard LMP resource discovery. For further reading and extensive explanation about LMP, the reader can refer to [5].

The TE link concept, understood as an aggregation of data links, is introduced in GMPLS so that scalability of both routing and signaling protocols can be improved. Specifically, LMP provides a mechanism to verify the connectivity of the data links, and to automatically learn TE link and data link interface associations, avoiding in this way long manual configurations. It is worth to point out that, these procedures should be performed when establishing the TE link and afterwards on a periodic basis for all the unallocated data links (i.e. to ensure their liveness). Initially, verification sessions are requested by sending *BeginVerify* messages [5], which can be accepted or rejected by responding with *BeginVerifyAck* or *BeginVerifyNack* messages [5]. Along this exchange, several parameters for the current verification session are set (i.e. verification identifier, verification scope, number of data links to be verified, data link transmission properties ...). Then, once the session has been established, *Test*

messages [5] are sent in-band over the first data link being verified. These messages contain the local interface identifier of that data link so that, whether a node detects the arrival of a *Test* message on an incoming port, it can resolve its remote interface association (i.e the connectivity of that port), which is subsequently notified to the neighboring node by sending a *TestStatusSuccess* message [5]. Conversely, if no *Test* message is received within an observation period (set during session establishment), a *TestStatusFailure* message [5] is returned. In turn *TestStatusAck* messages [5] are sent back, confirming the receipt of either a *TestStatusSuccess* or a *TestStatusFailure*. This operation is sequentially performed for all the data links previously considered under verification. Finally, verification sessions are ended by exchanging *EndVerify* and *EndVerifyAck* messages [5]. It is worth to recall that, since *Test* messages are sent in-band over the transport plane, electrical-optical-electrical conversion capabilities are needed on every data plane port, thus complicating its adoption by all-optical devices (extra hardware might be needed in most situations).

### 3 Resource Discovery using Loss-of-Light

The transmission of light on the emitter side of a link and its detection on the receiver side can be also a possibility to achieve transport plane interface mappings. This way, there is no necessity to send control information in-band (i.e. *Test* messages), so permitting control and transport planes to be decoupled. Moreover, this solution can be easily adopted by all-optical devices.

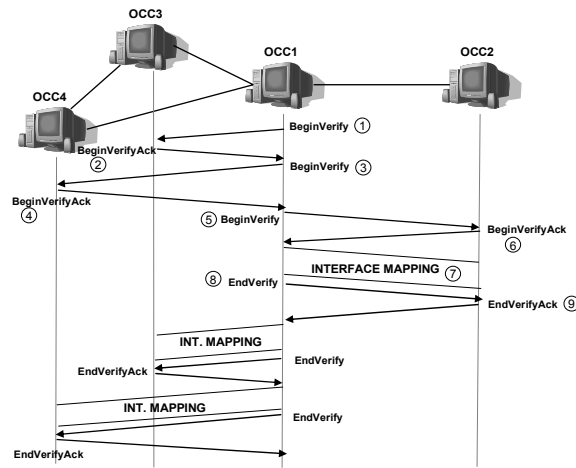
Nevertheless, there are two aspects in this method which deserve special attention. The first one is *time ortogonality versus data traffic*. This means that verification can not be performed while data links are carrying user traffic. However, this is not a big drawback if we take into account that resource discovery is performed before to consider data links ready to carry traffic (before the LMP Link Property Correlation process [5]). The second one is *Loss-of-Light (LoL) verification collision*. Firstly note that a node can only perform verification with one neighbor at a time. In addition, it must be avoided the situation where a node receives light coming from different simultaneous sessions at the same time, since this could suppose an erroneous interface mapping. It is worth to mention that, when a node is acting as an emitter of a verification session, it can send light to any of its neighbors, since it does not know in principle the connectivity of its outgoing data links.

### 4 Parallel and Concurrent Discovery Mechanisms

In this Section, we present both PD and CD mechanisms. These mechanisms, using the LoL to resolve interface mappings, extend the LMP behavior to enable resource discovery (i.e. link connectivity verification) over all-optical meshed networks.

Concerning PD (Fig. 1), when a node wants to perform verification towards a certain neighbor (for instance, OCC1 desires to start a verification towards OCC2), to avoid any verification collision, it has to ensure that none of its neighbors is also performing verification at the moment. Let us suppose the situation where OCC1

starts verification against OCC2 while OCC4 is performing verification towards OCC3. Since OCC1 does not know in advance the connectivity of its outgoing data links, it can emit light on a data link directly connected to OCC3. This situation would lead to a verification collision on OCC3, thus an erroneous interface mapping could happen on OCC3. To avoid such situations, firstly OCC1 tries to reserve in a sequential way the rest of neighbors by sending *BeginVerify* messages (step 1). The reception of a *BeginVerifyAck* message (step 2) from a neighbor indicates that it was not under verification and also that it will not accept any verification until the present one will be concluded. The same is done until all the other neighbors are reserved (steps 3, 4). It is worth to mention that, since *Test* messages are here sent over the control plane, the parameters needed to be negotiated for the verification sessions are drastically reduced compared to [5].



**Fig. 1.** PD mechanism for meshed networks

In Fig. 1, once OCC3 and OCC4 have been reserved, the session towards OCC2 can be initiated (steps 5, 6). Interface mappings (step 7) are resolved by the exchange of *Test*, *TestStatusSuccess* and *TestStatusFailure* messages like in standard LMP [5], but sending here the *Test* messages over the control plane. This way, in order to test an outgoing interface, the emitter turns its interface related laser on and sends a *Test* message to the neighbor, which contains the local identifier of that interface. When it arrives, the receiver checks if there is any presence of light on any of its incoming interfaces. If so, it resolves the interface mapping and returns a *TestStatusSuccess* containing the local and remote interface association. If no light is detected on any interface, a *TestStatusFailure* is returned. Upon their reception, the emitter confirms it by sending back a *TestStatusAck* message, stores interface association in case of a successful verification and it subsequently turns the laser down. This process is performed for all the emitter outgoing interfaces under verification. Once all of them have been tested, the session is ended by exchanging *EndVerify* and *EndVerifyAck* messages (steps 8, 9). In that moment, if any of the reserved neighbors have not been previously verified, interface mappings are sequentially performed towards each one.

Considering that, in order to initiate verification procedures, a node must reserve all its neighboring nodes so ensuring that none of them will accept any verification request, it is worth noting the possibility to concurrently verify all of them at a time. Unlike PD, where neighboring nodes are verified sequentially, the here proposed CD mechanism (Fig. 2) tries to take advantage of such situation. In the CD mechanism, firstly neighboring nodes are reserved as in PD by exchanging *BeginVerify* and *BeginVerifyAck* messages. Then, interface mappings are performed as follows. In order to test an outgoing interface, first of all the emitter turns its interface related laser on (step 1). Next to, *Test* messages containing the related local interface identifier are sent to all neighboring nodes (step 2). In turn, upon reception of a *Test* message, each node checks the presence of light on any of its incoming interfaces, so that whether the checking has been successful, a *TestStatusSuccess* is responded carrying the interface association (step 3). Otherwise, a *TestStatusFailure* is returned (step 4, 5). Once the emitter has received response from all its neighboring nodes, it stores the interface association (if any) and it subsequently turns the laser down (step 6). This process is done for all the emitter outgoing interfaces. Finally, verification sessions are ended as in PD.

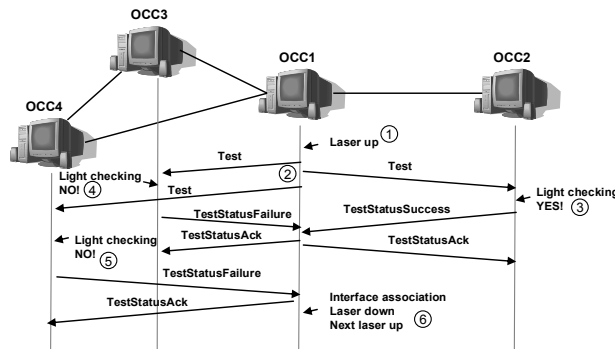


Fig. 2. CD interface mapping

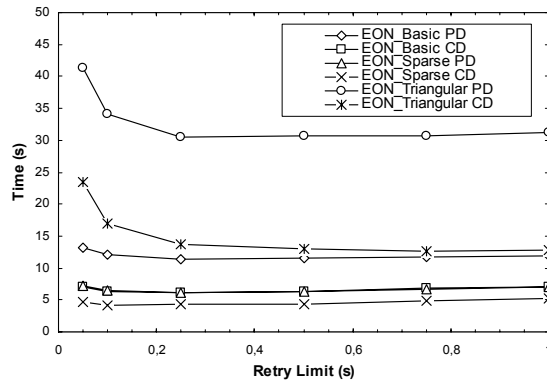
In both PD and CD, if a node, which at a moment is performing verification, receives a *BeginVerify* message, it responds with a *BeginVerifyNack* indicating that it is unwilling to verify. In turn, whether a node receives a *BeginVerifyNack*, it liberates any previously established session by exchanging *EndVerify* and *EndVerifyAck* messages and then waits for a random time (uniformly distributed between 0 and *retry limit*) before re-trying verification. Note that, as in LMP standardization, these mechanisms can be periodically performed for all unallocated data links. In such case, to avoid any traffic disruption, those allocated data links are excluded from verification.

## 5. Performance Evaluation

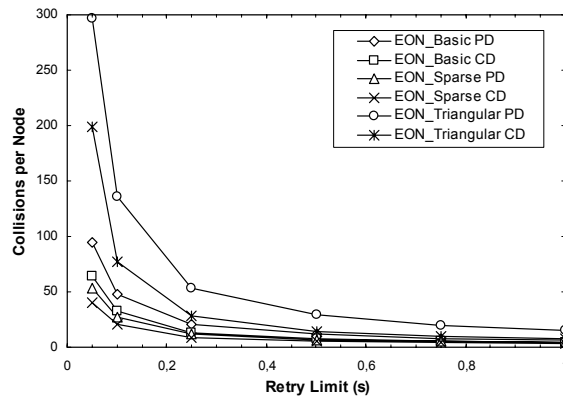
In order to quantify the performance of our proposals, simulation studies using the OPNET network simulator [7] over three 28-node different topologies, each one with different average nodal degree (ND), have been carried out. Specifically, the

European Optical Network (EON) *basic reference* (ND = 2.93), *Sparse Ring* (ND = 2.36) and *Densely Meshed Triangular* (ND = 4.36) topologies have been considered [8]. In this way, the obtained results show also how both PD and CD behave as average nodal degree increases in a meshed topology.

The evaluation of the mechanisms has been done taking the discovery time (i.e. the time required to discover the whole transport plane) as the figure of merit. Some assumptions have been initially adopted: (1) eight wavelengths (four bidirectional data links) are available between nodes; (2) the retry values in the PD and CD mechanisms are generated following a uniform distribution between 0 and *retry limit*; (3) 100  $\mu$ s for laser turn on/off time; (4) 10 ms for the transponder reaction time upon presence/absence of light (i.e. once a laser is turned up, the node waits 10 ms to send the related *Test* message). In this way, we ensure that, upon receiving a *Test* message, the neighbor will have detected the optical power rise on its incoming ports (if connectivity exists between both). Each simulation has been conducted in order to reach steady state results within a 95% confidence interval.



**Fig. 3.** Transport plane discovery times vs. retry limit



**Fig. 4.** Collisions per node vs. retry limit

Fig. 3 plots the behavior of PD and CD function of ND and the *retry limit* value used. Although both mechanisms highly depend on ND, it can be noticed their

feasibility, even in densely meshed networks. It can be also seen notwithstanding, how CD outperforms PD in every evaluated topology. Furthermore, the difference between both drastically increases as nodal degree gets higher. For instance, when considering the Sparse Ring topology (degree 2.36), the differences are about 2 sec, whereas when considering the Triangular one (degree 4.36), the differences can rise up to 20 sec.

To properly choose the best value for the *retry limit*, we have also quantified the introduced overhead (due to the retrying) for both PD and CD in terms of the number of collisions per node versus the *retry limit* (Fig. 4). First of all, it can be seen how CD introduces less overhead than PD. In addition, it can be concluded that the best compromise between the achieved discovery time and the introduced overhead is obtained for a *retry limit* value equal to 500 ms. Although the lowest discovery times are achieved for a retry limit value of 250 ms in almost all situations (except when using CD over the triangular topology), similar performance is obtained for a 500 ms *retry limit* while introducing less overhead. Therefore, according to these obtained results, we suggest the use of this value to implement both mechanisms. It is noteworthy how, lower *retry limit* values than 250 ms are translated into large undesirable overhead increments.

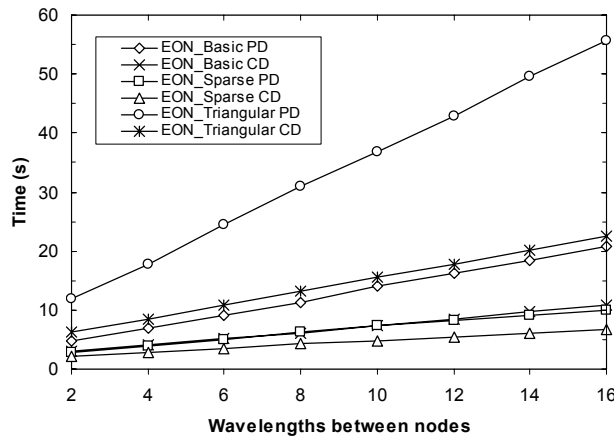


Fig. 5. Transport plane discovery times vs. the number of data links between nodes

So far, both PD and CD have been evaluated taking into account a fixed number of data links between neighbors (8 wavelengths, 4 bidirectional data links). In addition, we have also evaluated how both proposals depend on the number of wavelengths established between nodes. Here, in order to minimize the introduced overhead and according to the previous results, a 500 ms *retry limit* value has been applied. The results are shown in Fig. 5. We can firstly notice the linear dependence of PD and CD on the number of resources between nodes. Moreover, we can observe that this dependence becomes more stressed as ND gets higher, especially for PD. In fact, when considering 16 wavelengths (8 bidirectional data links), we can distinguish differences greater than 30 seconds between PD and CD on the triangular topology.

## 6. Conclusions

In this paper, we have addressed the automatic resource discovery of all-optical meshed transport networks. To begin with, we have introduced the standard LMP link connectivity verification functionalities, emphasizing its electrical-optical-electrical capability requirement, usually not available in all-optical devices. Alternatively, in order to overcome such requirement, we have also presented a method which resolves data link interface mappings without requiring to send in-band *Test* messages. Then, we have proposed two novel discovery mechanisms for meshed networks, whose performance has been evaluated and further compared by means of simulations over three reference networks with different nodal degrees. Mainly, simulation results show the feasibility of both mechanisms. Moreover, improvements can be seen for CD in front of PD in terms of less achieved total discovery times, less introduced overhead and less dependence on the number of existent resources (i.e. wavelengths) between nodes. Furthermore, such improvements drastically increase as the average node degree gets higher.

## Acknowledgements

The work reported in this paper has been partially supported by the Spanish Science Ministry through Project "Red inteligente GMPLS/ASON con integración de nodos reconfigurables (RINGING)", (TEC2005-08051-C03-02), and by the i2CAT Foundation ([www.i2cat.cat](http://www.i2cat.cat)).

## References

1. Mannie, E.: Generalized Multi-Protocol Label Switching (GMPLS) Architecture, IETF RFC 3945, Oct. 2004.
2. Jajszczyk, A.: Automatically switched optical networks: benefits and requirements, IEEE Communications Magazine, Feb. 2005.
3. ITU-T Recommendation G.8080/Y.1304 version 1.0 (and Amendment 2): Architecture for the automatically switched optical network (ASON), Nov. 2001 (and Feb. 2005).
4. Katz,D., Kompella, K., Yeung,D.: Traffic Engineering (TE) Extensions to OSPF Version 2, IETF RFC 3630, Sept. 2003.
5. Lang, J.: Link Management Protocol (LMP), IETF RFC 4204, Oct. 2005.
6. Perelló, J., et al.: Resource Discovery using Link Management Protocol for ASON/GMPLS Ring-based Transport Networks, in Proc. of 32nd European Conference on Optical Communication (ECOC), Sept. 2006.
7. OPNET Technologies Inc., [www.opnet.com](http://www.opnet.com)
8. Inkret, R., Kuchar, A., Mikac, B.: Advanced Infrastructure for Photonic Networks, Extended Final Report of COST Action 266, Sept. 2003.