

Signaling Protocols for Lightpath Provisioning

Srinivasan Seetharaman
The Ohio State University
seetharaman.6@osu.edu

Arjan Durresi
The Ohio State University
durresi@cis.ohio-state.edu

Raj Jain
Nayna Networks
raj@nayna.com

Abstract

Challenges presented by the exponential growth of the Internet have resulted in the intense demand for broadband services. Providing bandwidth and connectivity on demand has risen to be an important goal for service providers. To realize this, signaling in the optical network seems to be a critical component. Provisioning involves establishing a circuit from one client end system to another, through the optical backbone. With this general notion, this paper discusses in brief the role of signaling in all-optical networks, which employ DWDM for its lower layer transport. The architectural choices pertaining to lightpath services are outlined. Signaling procedures for ensuring dynamic end-to-end lightpath setup have also been presented. Path provisioning comprises a string of operations like service & neighbor discovery, route computation, signaling requests, and path setup. This paper tries to summarize the various strategies, designed for each step, for use in the all-optical networks.

1. Introduction

In satisfying the increasing demand for bandwidth, optical network technologies represent an unique opportunity because of their almost unlimited potential bandwidth. Recent developments in wavelength-division multiplexing (WDM) technology have dramatically increased the traffic capacities of optical networks.

The general acceptance of IP Multi-Protocol Label Switching (MPLS) framework as a common control plane for data and optical layers has spawned lots of research in this field. Research is ongoing to introduce more intelligence in the control plane of the optical transport systems. The aim is to make optical transport systems more survivable, flexible, controllable and open for traffic engineering. Some of the essential desirable attributes of optical transport networks include real-time provisioning of lightpaths, providing capabilities that enhance network survivability, providing interoperability characteristics between vendor-specific optical sub-networks, and enabling survivability capabilities in operational contexts.

Each client needs to connect to its peer via the optical network. This action requires an appropriate interface at each point. IP over MPLS tries to achieve that with the minimum possible complexity. MPLS brings two main advantages. First, it creates frontiers for traffic engineering. Second, it binds well to WDM by using wavelengths as labels. This extension of the MPLS is called the Multi-protocol lambda switching (MPLambdaS)[1] and it switches data based on the label (wavelength). The transparency of the optical network enhances multiprotocol usage of the resources.

In the MPL(ambda)S architecture, there exist separate clouds of IP networks and of WDM networks. Transfer of data from a source IP router to a destination is the main objective. Signaling lightpaths means requesting optical network elements for a path allocation. The procedures for establishing the connection dynamically, once the optimal path has been determined, is the main topic of discussion. This paper starts off with a description of the optical network architecture in Section 2. Section 3 describes the need for signaling and where it fits in the IP over MPLS architecture. Section 4 briefs the protocol related aspects of lightpath provisioning. Section 5 suggests certain enhancements in the current proposals for providing better services. The last section presents certain drawbacks and concludes the article. Since industry standards are still evolving, the protocols discussed here are not industry wide standards and should not be taken as one.

2. Optical Network Architecture

There are proposals underway in the industry to employ an IP-centric control plane, utilizing IP-based protocols for dynamic provisioning and perhaps restoration of lightpaths within and across optical sub-networks. These proposals are based on the practical view that signaling and routing mechanisms developed for IP traffic engineering applications could be re-used in optical networks so as to support seamless connections. Thus, every node in the optical network may consist of an IP module and a reconfigurable optical crossconnect switch. The IP module is responsible for all non-local management functions, like management of optical resources, provisioning etc. This brings in a constraint

that the node should have an IP address and must be IP-aware.




Figure 1. Generic model of Optical networks

The optical network model considered in this paper consists of multiple Optical Crossconnects (OXCs) interconnected by optical links in a general topology (referred to as an "optical mesh network") as shown in Figure 1. The signaling and routing interface between the client and optical networks is referred to as the *User-Network Interface* (UNI). If the optical network consists of multiple subnetworks, the interaction between the elements of these subnetworks is called the *Network-Network interface* (NNI). In Figure 1, interface A refers to UNI and the interface B represent NNI. The component links are responsible for data transport.

The services defined at these interfaces to a large degree determine the type and amount of control flow across them. It is possible to have a unified service definition across both these interfaces such that there is virtually no difference in the control flow across them, as seen in the *Unified services* model. This paper deals with the optical network service model where A & B are distinct, also known as the *Domain services* model.

The UNI & NNI interfaces are realized by having a control channel between each entity to one another (Direct) or to a common controller (Indirect). The control channel should be bi-directional and point-to-point so as to have a generic protocol on each node. The control channel can be in-band or out-of-band. The bottom line is that there should be one default previously established communication channel to support signaling or control.

There are different views on the model for interaction between the optical network and client networks, such as IP networks. The first model is called the *Peer model* where the IP/MPLS layers act as peers of the optical transport network. This stresses that the optical network elements exchange topology information with the client elements. The second model, called the *Overlay model* separates the IP/MPLS routing, topology distribution, and signaling protocols from those in the optical network. In the third model, called the *Augmented model*, there are actually separate routing instances in the IP and optical domains, but information from one routing instance is

passed through the other routing instance[1]. This paper deals predominantly with the peer/overlay model, though all that is stated can be related to the augmented model too.

3. Signaling Background

Providing bandwidth on demand is an important goal of the service providers. This dynamic resource allocation or provisioning is conducted via signaling operations, which ride over the control channel. In communications, signaling is the exchange of information between involved points in the network that sets up, controls, and terminates each connection. A lightpath from one end to another is established by setting up suitable crossconnects at the entry point of the optical network, the exit point and en route such that a continuous physical path exists all the way through.

As noted from the previous section, the architecture proposed has the tracks laid already for getting control messages across—the Control Channel. This bi-directional channel is essential for implementing the IP-centric control plane. The methodology of realizing this channel is left to the network operators. The control channel provides router-to-router connectivity over this link. Hence considered to reflect (or is similar to) the physical network topology. The neighbors of any particular node can be identified by inspecting its control channel connections.

The MPL(ambda)S architecture uses wavelengths (λ s) as labels. Therefore, all packets sharing a lightpath have the same wavelength at each hop. The key advantage of this label encoding is that optical devices can easily switch light signals by looking at the wavelength or port numbers (in case of a single wavelength per port). There are two options for MPL(ambda)S-based signaling – *Resource reSerVation Protocol (RSVP)* and *Constraint Routed Label Distribution Protocol (CR-LDP)*. The precise purpose of signaling and the messages carried by the signaling protocols are to be discussed in Section 4. There are some basic differences between the two protocols, but both essentially allow hop-by-hop signaling from a source to a destination node and in the reverse direction. Each of these protocols are capable of providing quality of service (QoS) and traffic engineering. To work satisfactorily, certain new features must be introduced – support for bi-directional paths, support for switches without wavelength conversion, support for establishing shared backup paths, and fault tolerance.

Automated establishment of lightpaths involves setting up the crossconnect table entries in the appropriate OXCs in a coordinated manner such that the desired physical path is realized. The request to establish a connection must identify the endpoints of the path. In addition, it may

also optionally specify the input and output ports, wavelengths, and TDM channels. The request may also include some optical parameters. On receipt of the request, the ingress node computes a suitable route for the requested path, following applicable policies and constraints, and sets up the path.

3.1 Constraint Routed Label Distribution Protocol (CR-LDP)

Label Distribution Protocol (LDP) has been defined for distribution of labels inside one MPLS domain. CR-LDP is the constraint-based extension of LDP. One of the most important services that may be offered using MPLS in general and CR-LDP in particular is support for constraint-based routing of lightpaths across the optical network. Constraint based routing is a mechanism used to meet traffic-engineering requirements.

In optical networks, label mapping corresponds to the assignment of input or output ports and wavelengths for paths by optical switches and the communication of this information to the appropriate neighbors. A routing protocol or service discovery protocol is used to formulate that. A *Label Request* message may be used by the upstream OXC to request a port (and wavelength) assignment from the downstream OXC for the lightpath being established. A *Label Mapping* message is used by the downstream LSR (or egress OXC) to distribute label-mapping information to its upstream peer (or ingress OXC). These are two basic operations supported by CR-LDP.

To incorporate the optical constraints, some extensions to the current version of CR-LDP have been proposed [2]. Prominent ones are:

- Inclusion of Signaling Port ID
- Signaling Optical Switched Path Identifier
- Signaling the two end points of the path being set up
- Signaling requirements for both span and path protection
- Inclusion of a probe for recording the precise route of the path being established

3.2 Resource reSerVation Protocol (RSVP)

RSVP is an unicast and multicast signaling protocol, designed to install and maintain reservation state information at each routing engine along a path. The key characteristics of RSVP are that it is simplex, unidirectional, receiver-oriented and soft. The receiver of a data flow generally initiates and maintains the resource reservation used for that flow. It maintains "soft" state in routing engines, which means that the connection has a definite expiration time. The *Path* messages are propagated from the source towards potential recipients.

The receivers interested in communicating with the source send back the *Resv* messages.

The lightpath established will require small setup latency, support for bi-directionality, and rapid restoration in case of network failures. The adoption of RSVP for signaling path setup messages mandates the addition of some more functionality in the protocol [3]:

- Reduction of trail establishment latency by allowing resources to be configured in the downstream direction.
- Establishment of bi-directional lightpaths as a single process instead of establishing two uni-directional lightpaths, one in each direction.
- Fast failure notification to a node responsible for lightpath restoration so that restoration can be quickly initiated.
- Setting up protection paths. The extensions help signal protection requirements for all protection modes.
- Declaring a lightpath as being permanent. Most lightpaths are long-lived and should not be torn down unless explicitly requested.

The protocol extensions serve as the binding format for transporting the lightpath setup messages as well as for traffic engineering (TE). Therefore, these extensions are referred to as "RSVP-TE." The protocols discussed so far operate in the application layer and tend to regulate the path setup process. The complete protocol stack of an optical node that supports Lambda switching is illustrated in Figure 2.

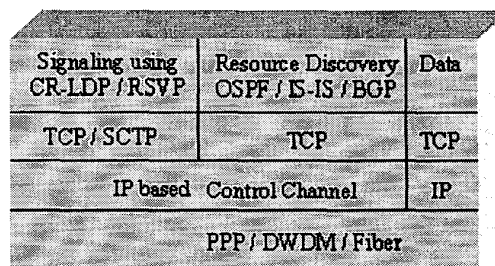


Figure 2. Summary of protocols functioning in a node (any OXC / border Router)

There is further discussion on which transport layer protocol to use for the signaling messages encapsulated in CR-LDP / RSVP-TE. The IETF Sigtran working group came up with designs for a new protocol called Stream Control Transmission Protocol (SCTP), which could be used in lieu of TCP, and is designed especially for signaling purposes [4]. Like TCP it runs directly over IP but offers some signaling tailored features.

4. Path Provisioning

Lightpath is the symbol of optical layer connectivity between two end points and is either unidirectional or bi-directional. Path provisioning means to establish an end-to-end connection as a cumulative effort on the part of the client network and the optical network. A complete path might contain the two endpoints and an array of intermediate OXCs for transport across the optical network. This section describes the handshake used for establishment of lightpaths in the network.

The procedure adopted can be categorized into two – UNI, NNI - based on the interface over which it is used. Control protocols are used at these interfaces to perform dynamic provisioning of paths. The connectivity over the client–optical interface is ensured by the UNI operations. Provisioning an end-to-end optical path across multiple sub-networks involves the establishment of path segments in each sub-network sequentially. Inside the optical domain, a path segment is established from the source OXC to a border OXC in the source sub-network. From this border OXC, signaling across the NNI is performed to establish a path segment to a border OXC in the next sub-network. Provisioning continues this way until the destination OXC is reached. Figure 3 tries to summarize the complete set of operations happening inside the network when a IP host requests a connection to another host. To automate this process, there are certain initiation procedures so as to determine the route for each segment (viz. IP host – IP border router , IP border router - border OXC, between border OXCs).

4.1 Resource Discovery

Routing within the optical network relies on knowledge of network topology and resource availability. First step towards network-wide link state determination is the discovery of the status of local links to all neighbors by each OXC. Each OXC must determine the aliveness of each optical link, the bandwidth and other parameters of the link, and the identity of the remote end of the link. On boot, each network node goes through neighbor discovery after which it creates an inventory of local resources and resource hierarchies, namely: channels, channel capacity, wavelengths, and links. The end result is that each OXC creates a port state database. All of the required information is obtained via routing protocol updates, and is propagated throughout the network. This procedure requires OSPF / IS-IS with appropriate extensions to handle optical parameters as well.

4.2 Route calculation

Different mechanisms for routing exist [1]. In the *integrated routing* approach the IP and optical networks

are assumed to run the same instance of an IP routing protocol, e.g., OSPF with suitable "optical" extensions. These extensions must capture optical link parameters, and any constraints specific to optical networks. This permits a router to compute an end-to-end path to another router across the optical network. In *domain specific routing*, the routing protocol allows the separation of routing within the optical and IP domains, with a standard routing protocol running between domains. The routing information exchange within the domain is performed using an interior gateway protocol (like OSPF or IS-IS, with appropriate traffic engineering and optical extensions) and between the domains using a protocol like Border Gateway Protocol (BGP).

The route computation, after receiving all network parameters in the form of link state packets, reduces to a mathematical problem. It involves solving a problem of Routing and Wavelength Assignment (RWA) for the new connection. The problem is simplified if there exists a wavelength converter at every hop in the optical network. Suitable solutions already exist to the RWA problem that makes optical routing a practical possibility [6].

4.3 End-to-end path setup

The link state information is used to compute the routes for lightpaths being established. It is assumed that a request to establish a lightpath may originate from an IP router (over the UNI), a border node (over the NNI) , or a management system. This request is carries all required parameters. After computing the route, the actual path establishment commences. Figure 3 tries to summarize the overall mechanism. However, once path setup is complete the data transfer happens passively and is straightforward without much intervention from the control plane. The connection needs to be maintained as per the service level agreements. The performance monitoring operation happens on a continual basis to note the network characteristics. This verifies the quality of service provided to the user. In case of problems, an alert is sent to the network devices to initiate diagnosis and healing operations [7].

Ongoing work at the Optical Interworking Forum (OIF), the Optical Domain Service Interconnect (ODSI) coalition and the Internet Engineering Task Force (IETF) reflects most of the work in this area. The OIF & ODSI deal with the Optical UNI related issues (assuming an overlay model). Since some of the solutions proposed by these organizations use protocols developed at IETF, the final consolidation happens at the IETF. Also, IETF is also working on the peer model along with the overlay model.

4.3.1 UNI Path Provisioning. The real handshake between the client network and the optical backbone happens after performing the initial service & neighbor discovery. The continued operation of the system requires that client systems constantly register with the optical network. The registration procedure aids in verifying local port connectivity between the optical and client devices, and allows each device to learn the IP address of the other to establish a UNI control channel. The following procedures may be made available over the UNI:

1. Client Registration: This service allows a client to register its address(es) and user group identifier(s) with the optical network.
2. Client De-Registration: This service allows a client to withdraw its address(es) and user group identifier(s) from the optical network.

The optical network primarily offers discrete capacity, high bandwidth connectivity in the form of lightpaths. The properties of the lightpaths are defined by the attributes specified during lightpath establishment or via acceptable modification requests. To ensure operation of the domain services model, the following actions need to be supported at the UNI so as to offer all essential lightpath services. The UNI signaling messages are structured as *requests* and *responses* [8].

1. Lightpath creation: This action allows a lightpath with the specified attributes to be created between a pair of termination points. Each lightpath is assigned a unique identifier by the optical network, called the lightpath ID. Lightpath creation may be subject to network-defined policies and security procedures.
2. Lightpath deletion: This action, originating from either end, allows an existing lightpath (referenced by its ID) to be deleted.
3. Lightpath modification: This action allows certain parameters of the lightpath (referenced by its ID) to be modified. Lightpath modification must not result in the interruption of the original lightpath.
4. Lightpath status enquiry: This service allows the status of certain parameters of the lightpath (referenced by its ID) to be queried.
5. Notification: This action sends an autonomous message from the optical network to client to indicate a change in the status of the lightpath (e.g., non-restorable lightpath failure).

Thus, the above actions provision both edges of the overall connection, while NNI provisioning builds the central portion of the setup

4.3.2 NNI Path Provisioning. The model for provisioning an optical path across optical sub-networks is as follows. A provisioning request may be received by a source OXC from the client border IP router (or from a

management system), specifying the source and destination end-points. The source end-point is implicit and the destination endpoint is identified by the IP address. In both cases, the routing of an optical path inside the optical backbone is done as follows[9]:

- a) The source OXC looks up its routing information corresponding to the specified destination IP address. If the destination is an OXC in the source sub-network, a path maybe directly computed to it. If the destination is an external address, the routing information will indicate a border OXC that would terminate the path in the source sub-network. A path is computed to the border OXC. This part of the signaling is called Intra-domain interface or IaDI.
- b) The computed path is signaled from the source to the destination OXC within the source sub-network. The destination OXC in the source sub-network determines if it is the ultimate destination of the path. If it is, then it completes the path set-up process. Otherwise, it determines the address of a border OXC in an adjacent sub-network that leads to the final destination. The next OXC then acts as the source for the path and the same steps are repeated. The signaling between adjacent sub-networks is called "Inter-Domain Interface" or IrDI.

Thus, NNI provisioning involves looking up in the routing table computed by various schemes mentioned previously and performing path setup within an optical sub-network [10].

Inside the sub-network, a lightpath request from a source is received by the first-hop optical network element (ONE). The first-hop ONE creates a lightpath setup message and sends it towards the destination of the lightpath, which would be the last-hop ONE. Every ONE in the path processes that message and allocates a channel for the lightpath on the downstream link. When a channel has been allocated at a particular node, the control elements at that node communicates with the switching matrix to provide the desired connectivity. After processing the setup, the destination returns an acknowledgement to the source. If no channel is available on some link, the setup fails, and a message is returned to the first-hop ONE informing it that the lightpath cannot be established.

When wavelength converters are not available, then a common wavelength must be located on each link along the entire route, which requires some degree of coordination between different nodes in choosing an appropriate wavelength. A probe message can be used to determine available wavelengths along wavelength continuous routes. A wavelength availability vector, containing list of candidate wavelengths, is forwarded down the desired route. If a wavelength on a link is not available or does not exist, then this is noted in the vector.

Once the entire route has been traversed, the vector will denote the viable wavelengths for the route. After wavelength selection is complete, setup happens in the above-described manner.

The construction of a bi-directional lightpath differs from the construction of a uni-directional lightpath above only in that upon receiving the setup request, the last-hop ONE returns the setup message using the reverse of the explicit route of the forward path. CR-LDP & RSVP have extensions to support the same. A lightpath must be removed when it is no longer required. To achieve this, an explicit release request is sent by the first-hop ONE along the lightpath route prompting each intermediate ONE to release the resources allocated to the lightpath.

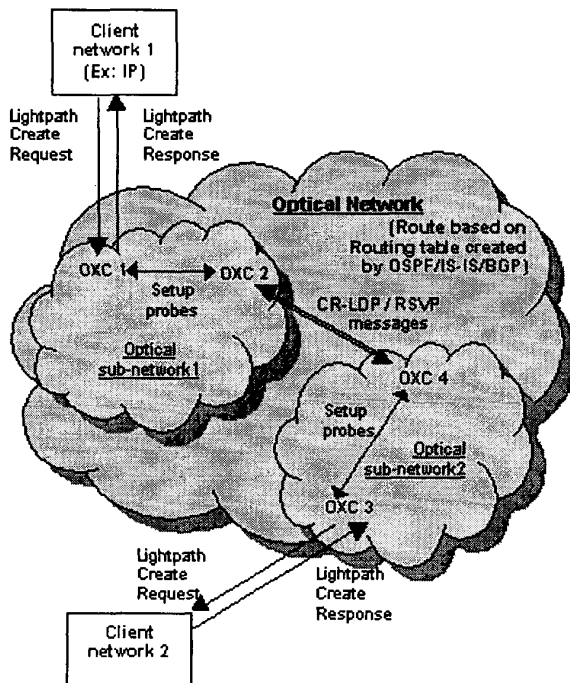


Figure 3. Summary of message exchanges in complete path provisioning

5. Miscellaneous Concepts

The architecture supports additional features and services that fill in most of the gap. Some adaptations to the basic design have also been proposed. This section tries to highlight such secondary enhancements.

5.1 Link Management Protocol (LMP)

A unique feature of a link is that the control channel and the associated component links are not required to be

transmitted along the same physical medium. Therefore, new mechanisms must be developed to manage links, both in terms of link provisioning and fault isolation. LMP, which runs between the neighboring nodes, does exactly that[11]. LMP will be used to maintain control channel connectivity, verify component link connectivity, and isolate link, fiber, or channel failures within the network. LMP is employed to maintain perpetual operation of the control channel by way of fault detection, and provisioning of alternate control channels.

5.2 Optical Service Level Agreements

Lightpath requests issued by a boundary router are only accepted within the constraints of the Optical Service Level Agreement (O-SLA). This can be regulated at UNI and NNI interfaces. The technical part of an O-SLA may include capacity (bandwidth) constraints, service performance parameters (other networking options), constraints on the scope of the lightpath request (which implies restricted destinations), authentication type[12]. Some authentication may be included in the request in order to verify that the rightful IP service provider issued the request. Confidentiality of UNI messages is also likely to be desirable, in which case encryption might be performed.

5.3 Link bundling

In an optical mesh network, two adjacent OXCs may have multiple, parallel, discrete bandwidth links connecting them. When a link state routing protocol is used, these links may all be bundled together and advertised in a single link state advertisement (LSA). The proposed method aims to reduce the number of routing adjacencies between neighboring nodes, and works in conjunction with the link management protocol. The purpose of link bundling is to improve routing scalability by reducing the amount of information that has to be handled by OSPF and/or IS-IS. The bundling procedure consists of forming "link groups"[13] which have a locally significant link group identifier.

5.4 Generalized MPLS (GMPLS)

The Multiprotocol Lambda Switching architecture has recently been extended to include devices whose forwarding plane recognizes neither packet, nor cell boundaries, and therefore, can't forward data based on the information carried in either packet or cell headers. Specifically, such devices include those in which the forwarding decision is based on time slots (e.g., SONET/SDH Muxes), wavelengths, or physical ports. GMPLS differs from traditional MPLS in that it supports

multiple types of switching, i.e., the addition of support for TDM, lambda, and fiber (port) switching. The support for the additional types of switching has driven generalized MPLS to extend certain base functions of traditional MPLS [14].

5.4 Signaling for Survivability

Service level agreements include protection/ restoration options to let the user choose between various schemes, over the UNI, for an individual connection. *Protection* schemes bring redundancy into the network resources so as to perform automatic switching onto the backup path (at the physical layer). *Restoration* is a collection of reactive techniques that work at the higher connection level. It provides alternate end-to-end paths. The backup paths can be pre-calculated during the time of route computation or can be calculated on the fly when a fault is detected. A parameter included in the signaling message indicates the type of protection desired for the lightpath. There are also provisions for signaling alternate paths in the event of a failure. Thus, signaling plays an important role in survivability of optical networks.

6. Conclusion

This article highlights some of the procedures and options built for dynamic path provisioning in optical networks. The signaling strategies quoted so far are not complete and various changes are still happening. But, the increasing demand for bandwidth has hastened the process and some of these innovations are being implemented in real life backbones. The control plane present in each node enables successful operation of the all-optical network. Practical implementations of the optical network adopt the domain services model and define provisioning for both the interfaces.

Path provisioning has many precursors like service discovery, neighbor discovery, route computation etc. before the core functions – UNI and NNI path setup can be performed. CR-LDP and RSVP-TE are being developed to be the signaling protocols. Various proposals developed at the OIF, ODSI and IETF try to achieve an automated provisioning having each node as an independent entity. Though most of the existing proposals are oriented towards IP, they are being extended for the non-IP community too. The transparency of the optical network and interoperability is under doubt and needs to be mended. Standardization is the next step towards industry-wide acceptance.

7. References

- [1] B. Rajagopalan, J. Luciani, D. Awduche, B. Cain, and B. Jamoussi, "IP Over Optical Networks - A Framework," Internet Draft draft-many-ip-optical-framework-01.txt, Work in Progress, July 2000.
- [2] Z.B. Tang *et al.* "Extensions to CR-LDP for Path Establishment in Optical Networks," Internet Draft draft-tang-crlldp-optical-00.txt, Work in progress, March 2000.
- [3] J.P. Lang, K. Mitra, and J. Drake,, "Extensions to RSVP for Optical Networking," Internet Draft draft-lang-mpls-rsvp-oxc-00.txt, Work in Progress, March 2000.
- [4] R.R.Stewart *et al.*, "Stream Control Transmission Protocol," RFC 2960, October 2000.
- [5] N.Chandhok *et al.*, "IP over MPLS – A Summary of Issues," Internet Draft draft-osu-ipo-mpls-issues-00.txt, Work in progress, July 2000.
- [6] K. Lee, S. park, K. Choe, and C. Park, "Routing and Wavelength Assignment in WDM All-Optical Networks," Electronic Letters, Vol.36 No.11, May 2000.
- [7] L. Ceuppens *et al.*, "Performance Monitoring in Photonic Networks in support of MPL(ambda)S," Internet Draft draft-ceuppens-mpls-optical-00.txt, Work in progress, March 2000.
- [8] O. S. Aboul-Magd, "Signaling Requirements at the Optical UNI," Internet Draft draft-bala-mpls-optical-uni-signaling-00.txt, Work in Progress, July 2000.
- [9] D. Pendarakis, B. Rajagopalan, and D. Saha, "Routing Information Exchange in Optical Networks," Internet Draft draft-prs-optical-routing-00.txt, Work in progress, April 2000.
- [10] S. Chaudhuri *et al.*, "Control of Lightpaths in an Optical Network," Internet Draft draft-chaudhuri-ip-olxc-control-00.txt, Work in progress, February 2000.
- [11] J.P. Lang, "Link Management Protocol (LMP)," Internet Draft draft-lang-mpls-lmp-01.txt, Work in progress, July 2000.
- [12] O. Duroyon, R. Hoebeke, and H. De Neve, "Triggering and advertising lightpaths in an IP over optical network," Internet Draft draft-duroyon-te-ip-optical-00.txt, Work in progress, July 2000.
- [13] B. Rajagopalan, and D. Saha., "Link Bundling Considerations in Optical Networks," Internet Draft draft-rs-optical-bundling-00.txt, Work in progress, July 2000.
- P. Ashwood-Smith *et al.*, "Generalized MPLS - Signaling Functional Description", Internet Draft draft-ashwood-generalized-mpls-signaling-00.txt, Work in progress, June 2000.