Multi-protocol lambda switching for IP over optical networks

Thomas M. Chen* and Willie Wu†

*Dept. of Electrical Engineering, SMU, Dallas, TX 75275
†Vitesse Semiconductor Corp., Richardson, TX 75080

ABSTRACT

This paper examines the multi-protocol lambda switching approach to support IP over WDM networks. The MPL(ambda)S approach extends a common control plane to IP and optical domains based on multiprotocol label switching (MPLS) with extensions for the unique characteristics of the optical network. We investigate some practical issues related to signaling, routing, differentiated services, and survivability.

Keywords: lambda switching, multiprotocol label switching, IP over WDM, optical transport network

1. INTRODUCTION

The exponential growth of Internet traffic has drawn extensive attention to optical networking technologies to multiply capacity by overcoming the electronic bottlenecks in the backbone network. Current network architectures are moving in the direction towards an all-optical core network exploiting wavelength division multiplexing (WDM) and wavelength routing where electronic processing and traffic control are relegated to the network edges.[1] WDM allows multiple optical signals at different wavelengths to be multiplexed in a single optic fiber. Current dense WDM systems offer 2.5 Gb/s (SONET OC-48 rate) or 10 Gb/s (OC-192 rates) per wavelength and approximately 16 to 128 channels per fiber.

Although IP routers can be interconnected by point-to-point WDM links, a reconfigurable optical network has the flexibility to respond to changing traffic patterns or traffic demands. ITU-T Recommendation G.872 describes long-haul digital optical transport networks (OTNs) as a mesh of optical add/drop multiplexers (OADMs) and optical cross-connects (OXCs) with multiple lightpaths between source-destination pairs.[2] OXCs can be rearranged for flexible routing of optical channels and fault restoration, typically through network management (not signaling). It has been recognized that current control of optical networks through network management is characteristically slow and manual, and therefore a considerable advantage can be realized by automating the control of the optical network, designated an automatic switched optical network (ASON).[3] A scenario with an ASON serving as an Internet backbone network is shown in Figure 1. A critical element of an ASON is a programmable OXC with enhanced capabilities to rearrange optical channels in real time by means of signaling. The potential advantages of an automated optical network include dynamic allocation of network resources, fast restoration and recovery from faults, and responsiveness to changing traffic patterns (presumably resulting in greater network efficiency).

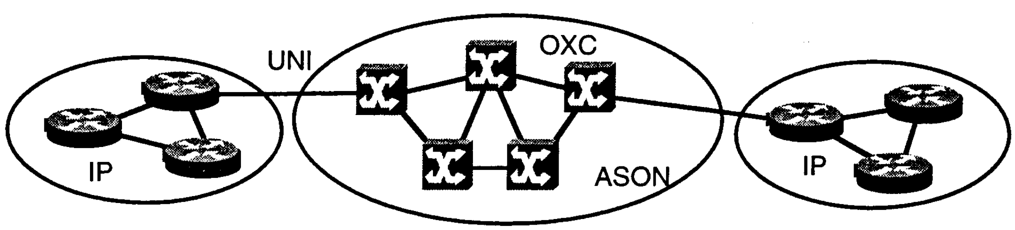

Figure 1. ASON serving as Internet backbone network.

In addition to a signaling protocol for resource reservation, an ASON needs a routing protocol to discover network-wide resource availability; a route selection algorithm; connection admission control; a global addressing scheme; and means

* Correspondence: T. Chen, Dept of Electrical Engineering, SMU, PO Box 750338, Dallas, TX 75275 USA; tel: 214-768-8541; email: tchen@seas.smu.edu. W. Wu, Vitesse Semiconductor Corp., 997 Hampshire Lane, Suite #200, Richardson, TX 75080 USA; tel: 972-479-0162; email: wu@vitesse.com.

to protect or restore connections in the event of faults. For practicality and interoperability, an ASON should make the most use of existing network protocols, in particular Internet protocols.[4] Unfortunately, Internet protocols cannot be used directly because of the physical differences in the optical network. For example, IP routers forward packets on the basis of their headers while OXCs switch on the basis of wavelengths or optical channels. Whereas packet networks have arbitrary granularity (individual packets or packet flows), an optical network has a "coarser" granularity (wavelengths). Also, the optical network is unable to perform queueing-based traffic control, such as packet scheduling, policing, and traffic shaping.

If an administrative boundary exists between the optical network and IP network (at the user-network interface in Figure 1), a so-called "overlay" approach may be reasonable.[3,5] IP routers would request optical connections through a well-defined user-network interface, and the optical network would fulfill requests by means of its own control protocols and algorithms. This approach is characterized by standardized interfaces for the optical network; simple connectivity services provided by the optical network; separate control planes for the IP and optical networks; and limited topology information exchanged between the IP and optical networks. Besides an administrative boundary, there are other reasons for this overlay approach. First, the optical network may serve other "clients" besides the IP network, thus it provides only a well-defined interface but does not reveal complete topology or resource information for direct control by the IP network. Second, the control system in the optical network can be optimized for the unique characteristics of the optical network without being strictly tied to Internet protocols. A clean separation of control planes leads to straightforward design of both networks. Third, the optical network and IP network can evolve independently according to their different needs.

An optical network serving various clients through a well-defined interface is the view of the ODSI (Optical Domain Service Initiative), a coalition of network equipment vendors and service providers promoting open interfaces for IP routers and other networking devices to dynamically request connections from an optical network.[4] The ODSI interface for optical networks allows three functions. First, by service discovery, a networking device (e.g., IP router) can detect that it is connected to an optical network capable of dynamic bandwidth allocation.[6] The service discovery is performed by a bidirectional exchange of device parameters through an in-band control channel on the port connected to the optical network. Second, by address registration, the networking device provides an IP address for the port connected to the optical network to enable a point-to-point optical channel trail to be created across the optical network to another IP address. Third, the networking device can participate in signaling to establish an optical channel trail to another device across the optical network.[7] Signaling can also query a connection's status, modify an existing connection, or terminate a connection. Signaling is viewed as "third party" meaning that the two endpoints of the connection do not necessarily have to initiate the connection request, e.g., it would be possible for a network manager to request a connection between two networking devices. When the optical network receives a connection request, it will validate it, find whether a suitable path exists, and negotiate with the two connection endpoints.

The ODSI specification does not mandate a specific control plane for the optical network, and does not preclude a control plane based on MPLS (multiprotocol label switching) protocols.[8-10] In contrast to the overlay approach, multi-protocol lambda switching or MPL(ambda)S is a "peer model" approach that treats the IP and optical domains as a single integrated network controlled by MPLS protocols.[5,11] The OXCs participate in routing and signaling protocols in the same manner as MPLS label switching routers (LSRs), where MPLS protocols are modified to accommodate the unique characteristics of the optical network. Because analogies can be observed between OXCs and LSRs, the MPL(ambda)S approach extending the MPLS control plane to the optical network may seem natural from an IP-centric viewpoint.[11,12] An OXC associates an input port and incoming optical channel with an output port and outgoing optical channel, analogous to an LSR that associates an input port and incoming label with an output port and outgoing label. By extending the MPLS control plane to the optical network, it may be possible to operate the IP and optical networks seamlessly and eliminate the need to develop new dynamic control protocols for the optical network. However, MPLS protocols must be extended to encompass the physical differences of the optical network, and LSRs must become aware of the optical network.

This paper examines some practical issues related to the MPL(ambda)S approach to IP over WDM integration. Section 2 describes extensions to MPLS signaling and label distribution protocols to allow connections through the optical network. Section 3 describes extensions to MPLS routing protocols to advertise optical link status information. MPL(ambda)S support for the differentiated services architecture is explored in Section 4. Finally, implications of MPL(abmda)S on network survivability is investigated in Section 5.

2. SIGNALING

In MPLS, label switched paths (LSPs) are established and managed through a signaling protocol such as RSVP-TE (resource reservation protocol with extensions) or CR-LDP (label distribution protocol with constraint-based routing).[13-16] The signaling protocol performs label assignment distribution, resource reservation, and path termination. In RSVP-TE, the sender first transmits a Path message to the receiver with a description of traffic characteristics. In response, the receiver returns a Resv message to request resources for the flow. Each LSR along the route has an opportunity to accept or reject the Resv message. If the request is rejected, the node will send an error message to the receiver to terminate the signaling process. If the request is accepted, the relevant flow state information will be installed to reserve resources at each node. To work with MPLS, the first LSR inserts a Label_Request object into the Path message to request a label binding. If an Explicit_Route object is added to the Path message, the Path message will be forwarded along a specific route. The last LSR will return a Resv message including a Label object. As the Resv message travels upstream (i.e., upstream relative to the direction of data in the LSP being established), each LSR will receive a label and record it in the forwarding table for the new LSP, and forward a chosen label to the next upstream LSR. By piggybacking label binding information on RSVP messages, resources can be reserved for an LSP at the same time as label assignments.

CR-LDP is an alternative signaling protocol designed for MPLS. Two LSRs must begin a bidirectional LDP session to exchange label information as LDP peers. During an LDP session, label bindings are assigned by the downstream LSR and label assignments are distributed from the downstream LSR to the upstream LSR (label assignment information flows in the opposite direction as the data packets). This can be done automatically by a downstream LSR advertising to its neighbors via a Label Mapping message without a label request (called downstream unsolicited distribution), or on demand when an upstream LSR requests a label assignment from a downstream LSR via a Label Request message (called downstream-on-demand label distribution). To reserve resources along an explicit path as well as distribute label information, a Label Request message is sent from the ingress LSR through a specified sequence of LSRs to the egress LSR. In addition to a class of service (CoS) request, the message may include traffic parameters such as peak rate, peak burst size, and committed rate. If an LSR can accomodate the new connection, it will reserve the corresponding resources and forward the Label Request message to the next LSR. Upon receiving the message, the egress LSR will return a Label Mapping message in the upstream direction to the ingress LSR.

MPL(ambda)S seeks to extend MPLS signaling and label distribution protocols to the optical network.[17-19] The extensions should allow OXCs to establish connections between an input port and incoming wavelength with an appropriate output port and outgoing wavelength, but this operation is analogous to the associations made by an LSR between an incoming label and outgoing label. Thus, existing MPLS signaling protocols do not have to be modified drastically. Extensions to MPLS signaling called generalized MPLS have been proposed to encompass wavelength routing and other types of switching.[17] The first modification involves broadening the label concept to recognize that labels may be encoded as a packet header field (MPLS), time slot (TDM), or wavelength (OXCs). This generalized label concept does not cause confusion if the two endpoints of a link have a mutual understanding of the meaning of the label encoding. Thus LSR interfaces can be different types including the usual packet switch capable (PSC) interface which can recognize labels and forward packets based on them, TDM capable interfaces, and lambda switch capable (LSC) interfaces for OXCs that operate at the level of wavelengths.

The second modification involves generalizing the label request in RSVP-TE or CR-LDP.[17] The generalized label request includes the desired link protection (e.g., dedicated or shared protection), encoding type (physical framing), and payload type. The type of label is not explicitly contained in label request messages because the type is assumed to be understood by the LSRs participating in label distribution. A Path (RSVP-TE) or Label Request (CR-LDP) message is sent from ingress node to egress node. Any node that rejects the request can generate an error notification message. If the request is successful, a Resv/Label Mapping message is returned in the upstream direction with a generalized label.

The third modification involves a "suggested label" option where an upstream node can express a label preference to a downstream node.[17,19] A suggested label has the same format as a generalized label and is carried in Path/Label Request messages. The purpose is to allow the upstream node to start its configuration early with the proposed label before the label is communicated by the downstream node. Early configuration is intended to help OXCs that can take significant time to establish a label. Reducing the label setup time by this option might be important, for example, if fast LSP rerouting by signaling is needed for fault restoration. However, if a downstream node sends a different label upstream, it must be accepted by the upstream node; downstream control of labels always has precedence over upstream suggestions.

Another proposal suggests several new types of objects for RSVP-TE and new TLVs (type-length-value encodings) for CR-LDP in support of optical LSPs.[18] The new TLVs allow an optical LDP session to be established, an optical channel trail to be described, and optical labels to be set. Similarly, the new objects for RSVP-TE include an optical label request, optical interface type, optical channel trail descriptor, and optical label.

The physical difference between LSPs and optical channel trails raises a question of how LSPs are mapped to optical channel trails. Forwarding data as packets, MPLS routers can support a large number of LSPs with arbitrary bandwidth granularities. On the other hand, OXCs can support a relatively smaller number of optical channel trails, whose bandwidth has much coarser granularity, e.g., OC-12, OC-48, or OC-192. Clearly, if a low-bandwidth LSP is mapped to one optical channel trail, the resulting network utilization could be low. "Nesting LSPs" has been proposed as a method to map a number of low-bandwidth LSPs into an optical channel trail.[11] It depends on pushing labels on several LSPs at the entry to the optical channel trail, and then popping the labels at the end of the trail. Since the method depends on label pushing/popping which cannot be done by OXCs, it must be done at the border LSRs on the edge of the optical network.

3. ROUTING

MPLS signaling protocols work in combination with a constraint-based routing protocol such as OSPF or IS-IS with extensions to carry additional link constraint information in link state advertisements.[20,21] OSPF and IS-IS are dynamic link-state routing protocols where each router broadcasts its status information to all other routers. Link state advertisements traditionally include reachability information and administrative costs for the links. Both routing protocols have been extended to carry link constraint information to make the route selection sensitive to QoS requirements. Additional information could include maximum link bandwidth, maximum reservable bandwidth, current bandwidth reservation at each of eight priority levels, and the resource class of the link.

MPL(ambda)S presumes that the same routing protocol will operate over the IP and optical networks, which might be an extension of constraint-based OSPF or IS-IS with additional extensions for the optical network topology state, available bandwidth, and available optical channels. Extensions to OSPF and IS-IS to include optical link status information have been proposed.[22-24] There is no obvious reason against the use of these proven routing protocols within the optical network for resource discovery, especially in the absence of existing alternatives. A question is what optical link status information is most useful for traffic engineering. The information should represent the state of each fiber and the state of individual optical channel trails within each fiber. Natural choices for attributes are maximum and minimum reservable bandwidth. Other attributes might include the type of protection, e.g., 1+1 (dedicated protection) or 1:N (one disjoint protection channel for N working channels). Optional attributes might include the type of physical layer framing (SONET, SDH, digital wrapper, gigabit ethernet, or none); restoration speed; preemption priority (whether it can be preempted by a higher priority connection); propagation delay; diversity (whether it shares common facilities with a given set of other circuits); unidirectional or bidirectional; resource class attributes allowing implementation of policies; types of links (e.g., normal links or non-packet links).[4,11,22,23] A proposed concept is "shared risk link group" (SRLG) which is a set of links that share a resource and will be collectively effected if the resource fails. A link may belong to multiple SRLGs.[22,23] It may be possible to bundle multiple links together to advertise as a single virtual link, in which case optional attributes may be physical or logical optical link, total available bandwidth, and number of wavelengths.[24]

If the same routing protocol is used within the IP and optical domains, the LSRs and OXCs exchange routing information as peers. In this case, it appears the LSRs are more likely to benefit because they might be able to exploit knowledge about optical links (e.g., reliability or quality of links) to select routes to other LSRs across the optical network. Once established, the path across the optical network might be treated as a single virtual link in further OSPF advertisements. On the other hand, there is not an obvious benefit to the optical network to gain routing information knowledge about the LSRs.

It is a misconception that a network running OSPF or IS-IS will keep identical link-state databases at every node. Both routing protocols allow a two-level hierarchy. The network may be divided into contiguous areas, each area running a separate copy of the link-state routing protocol. The topology of an area is invisible to the outside, and routers within each area are unaware of topology details outside of its area. Hence two routers in different areas will keep different link-state databases. The optical network would be a logical choice for an area, and the IP network could be one or more areas. Thus LSRs would not have the burden of maintaining complete topology information about the optical network, and the overall routing traffic is reduced compared to treating the entire network as a single link-state domain. This has the advantage that the optical network might change physical routes in the optical layer without unnecessarily involving the LSRs. In OSPF, routing information between separate areas is distributed through a special area called the backbone area. The backbone area must be contiguous, and all areas must attach directly to the backbone area. Likewise, in IS-IS, level 1 areas must be directly connected to a single level 2 subdomain (the top level).

The link state information is used to compute routes for optical channel trails being established. It is assumed that a request to establish an optical channel trail specifies the source and destination ports, bandwidth required, restoration parameters, and othe constraints on the path. This request is sent to the OXC that contains the source port, which is then responsible for computing the route and establishing the path. Route computation with constraints may be accomplished by any number of algorithms.[25] Typically, the constraint-based path selection procedure will involve pruning the database of candidate routes to remove links that are ineligible, due to either insufficient bandwidth or policy constraints. A shortest path algorithm (Dijkstra) is then run on the pruned topology to find a route that satisfies the required criteria. The path selection is more complicated if backup routes must be selected and established as well (discussed later).

4. DIFFERENTIATED SERVICES

The differentiated services (diffserv) architecture has been proposed to implement service classes in the Internet without the scalability difficulties of the integrated service (intserv) architecture.[26] Intserv makes use of a resource reservation protocol to request network resources for individual traffic flows satisfying specified QoS requirements. The scaling difficulties arise from the state information maintained in each router which increases proportionally with the number of traffic flows. Diffserv overcomes the scalability problem by avoiding per-flow state at each node. Instead, traffic is handled according to broad service classes (traffic aggregates) which are recognized by the diffserv codepoint (DSCP) in the packet header. The diffserv codepoints of packets are mapped to prescribed per-hop behaviors (PHBs) which are essentially packet handling instructions programmed in each node. Per-hop behavior is a general term encompassing any buffer management and packet scheduling actions, and hence the diffserv architecture is fairly flexible and broad.

PHBs for expedited forwarding (EF) service, assured forwarding (AF) service, and best-effort service have been identified.[27,28] Expedited forwarding service is characterized by low loss, low latency, low jitter, and assured bandwidth. The quality of the EF service should resemble a point-to-point connection or a "virtual leased line". This performance implies that packets will encounter minimal queueing at each node. Without reserving bandwidth or connection admission control to block traffic, the EF service can be realized by ensuring that the aggregate arrival rate at any node never exceeds the minimum departure rate. Thus the ingress traffic must be restricted and policed at the network boundary, and traffic may have to be shaped (or groomed) between nodes. Also, this service class must be guaranteed some amount of minimum throughput by means of the highest static priority or round robin packet scheduling.

Assured forwarding service is a means to offer different levels of forwarding performance for IP packets. Four AF classes are defined, where each AF class is allocated a certain amount of buffer space and bandwidth in each node. IP packets that wish to use the AF services are assigned into one or more of these AF classes. The level of forwarding assurance of an IP packet depends on how much forwarding resources has been allocated to the AF class, and the current load of the AF class. For example, AF classes might correspond to gold, silver, and bronze service classes where gold service is likely to be better than silver, and silver better than bronze. Within each AF class, packets can be marked with one of three possible drop precedence values. In case of congestion, the drop precedence of a packet determines the relative importance of the packet within the AF class. Packets with a lower drop precedence value are protected by first discarding packets with a higher drop precedence value. Thus forwarding assurance also depends on a packet's drop precedence.

Naturally, the optical network is incapable of packet processing and queueing, and therefore unable to support PHBs. The optical network may provide optical channel trails with different characteristics (bit error rate, transmission delay and jitter, bandwidth, monitoring capabilities, reliability or level of protection), but traffic grooming and PHBs must be performed in the IP domain. Within the diffserv architecture, the optical network appears simply as links interconnecting LSRs.

Bandwidth brokers (BBs) are a new concept introduced into the diffserv architecture to centrally allocate and manage the resources of the network.[29] More specifically, a BB handles negotiation of quality of service between users and the network, negotiations between separate network domains, maintenance of resource utilization information for a network domain, and regulation of ingress traffic across the boundary of a network domain (essential for expedited forwarding service). The BB in diffserv performs admission control and resource allocation without the maintenance of per-flow state in each node as in intserv.

A common control plane in MPL(ambda)S implies a BB spanning both IP and optical domains. The BB maintains information about all resources and manages the diffserv trafffic flows on an end-to-end basis. Is there an advantage to a common BB for both network domains? Because the optical network has a relatively small role and contains different physical resources, separate bandwidth brokers in the IP and optical domains may make sense. One BB keeps track of diffserv traffic flows and resource utilization in the IP domain, and another BB manages optical channels and wavelengths, which are limited resources in the optical domain. As shown in Figure 2, the BB in the IP nework negotiates with the BB in the optical network to request optical channel trails with specified characteristics. In this case, the interface between the BBs may be simple because the optical network provides a simple service. Separate BBs have the advantage of distributing the resource management functions, but might incur more delay due to the negotiation process between the BBs.

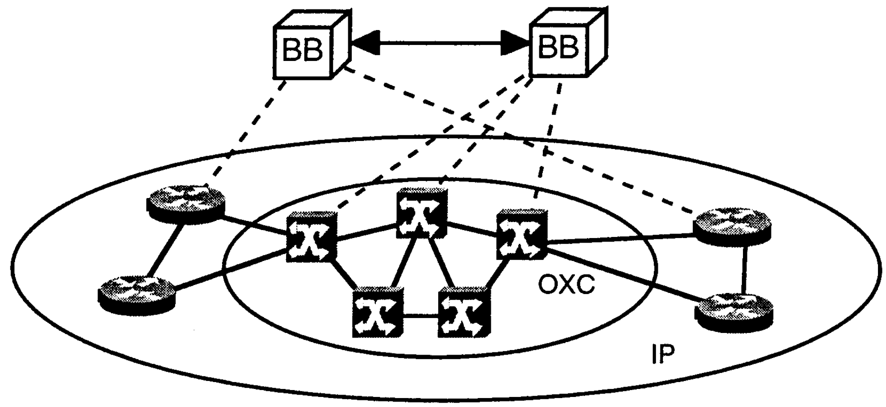

Figure 2. Bandwidth brokers in IP and optical domains.

5. FAULT PROTECTION AND RESTORATION

Minimizing service disruption in the possible event of various failures (fiber cuts, line card or switch failures, software failures) is paramount when terabits of data may be lost in every second. Because the IP and optical domains are physically different, fault management is usually approached in layers where a fault will be escalated to another layer if the first layer is unable to recover. While MPLS relies on LSP rerouting for fault recovery, the optical layer has more varied mechanisms. At the same time, recovery in the optical layer can be simpler, more robust, and efficient, but is more suited to fiber cuts and other physical media faults. Recovery in the electronic domain is still needed for line card or switch failures.

Survivability for optical networks encompasses both protection and network restoration capabilities.[30] Protection is hardware-based, pre-planned, and generally fast but limited to links and simple topologies like rings. Protection makes use of pre-assigned capacity between nodes. The simplest architecture has one working and one protection channel (1+1), while the most complex has n working and m protection channels (m:n). Optical channel trail protection is a dedicated end-to-end protection mechanism that can be used in any topology. A working trail is replaced by a protection trail if the working trail fails. Trail protection may also be 1+1 where the dedicated protection trail is only used for protection or m:n where extra traffic may be supported on the protection trails. Under 1+1 protection, a backup path is established for the protected primary path along a physically diverse route. A failure along the primary path results in an immediate switchover to the backup path. Under m:n shared protection, backup paths corresponding to physically diverse primary paths may share the same network resources. When a failure affects a primary path, it is assumed that the same failure will not effect the other primary paths whose backups share resources. Thus, the backup path for a primary path may be precomputed, but is activated only after failure of the primary path has been determined.

Optical network restoration is software-based, dynamic, and slower but applicable to general mesh topologies. It depends on OXCs to reroute optical channels around points of failure. Rerouting may be done locally or end-to-end. In local rerouting, the failed link is switched over to an alternate link between the adjacent OXCs when the link failure is detected. Link restoration does not effect the end-to-end route of the lightpath and may be the fastest recovery. When link restoration is not possible (e.g., in case of node failure), the affected lightpath may be rerouted to an alternate end-to-end path that completely avoids the OXCs or link segment where the failire occurred. For end-to-end restoration, alternate paths are typically pre-computed. Backup paths are usually physically diverse from the corresponding primary paths.

Because of the variety of mechanisms, optical channel trails may be provided with different grades of protection, for example: fully restorable with high priority; fully restorable but can be pre-empted by higher priority lightpaths if necessary; partially restorable on a best-effort basis; or unrestorable. These grades should be considered when LSPs are mapped to optical channel trails.

Automatic restoration of optical channel trails is commonly viewed as a service that should be offered by optical networks.[2] The IP layer will usually delay restoration actions to allow sufficient time for the optical network to recover. After the expiration of a holdoff time, it will be assumed that the optical network is unable to recover and the fault will be escalated up to the MPLS layer.

MPLS can reroute failed LSPs on the basis of a link, partial path, or end-to-end, as shown in Figure 3.[31] In link recovery, only individual LSP hops are rerouted. This can be done via a pre-established route between the two adjacent LSRs around the failed link. When a link failure is detected, the forwarding table of the upstream LSR is updated to use the label for the backup LSP, which causes traffic to flow around the failed link. Although fast, the result can be paths which are not optimal for traffic engineering. Hence, subsequent reconfiguration of LSPs may be done for optimization. In partial path rerouting, a new path segment is established from an unaffected upstream node to the egress node, thereby bypassing any downstream link or node failures. Finally, end-to-end rerouting in the MPLS layer is similar to end-to-end rerouting in the optical layer. In all recovery schemes, the backup route may be prespecified or dynamically computed. Prespecified backup routes will yield faster fault recovery by eliminating the need for a flooding procedure to find backup routes at time of failure, but primary and backup routes must be both found during connection setup. In all schemes, the rerouting may be accomplished by means of a signaling protocol or special control messages.

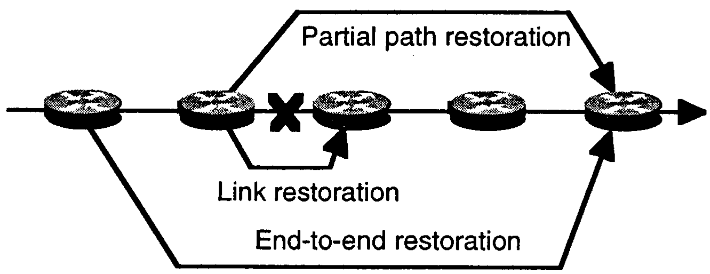

Figure 3. Rerouting schemes in MPLS.

A common control plane in MPL(ambda)S implies a unified fault recovery scheme contrasting with the usual layered approach. A unified approach might have two significant advantages. First, an important practical difficulty in layered protection is proper coordination between the recovery mechanisms in different layers. For example, the appropriate choices for holdoff timers is difficult to determine.[32] Lack of coordination between protection mechanisms at the optical and electronic layers might lead to inefficiency and duplication, or in the worst case, unpredictable race conditions or topology oscillations. A unified approach aware of all fault recovery mechanisms in the network promises to avoid the coordination or escalation problem.

Second, a unified scheme might be able to take advantage of knowledge of optical network characteristics to choose the best recovery action. A hypothetical example is shown in Figure 4 where the problem is to select the fastest backup route between LSR A and LSR B. One candidate backup route traverses the optical network, and another route does not. If it is known that the route through the optical network will take longer to establish or an optical channel is unavailable, the other backup route can be selected immediately.

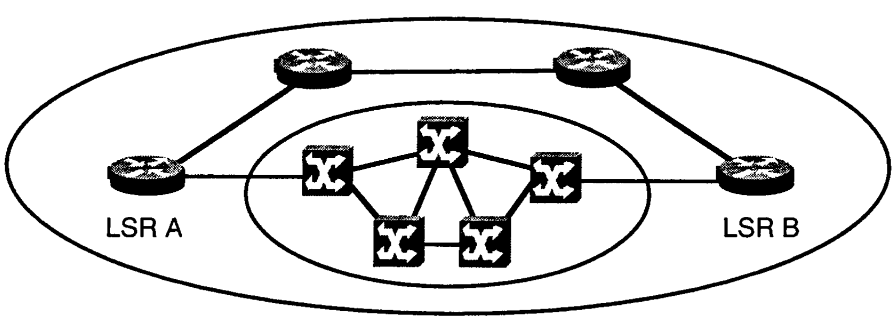

Figure 4. Example of choice between two backup routes.

6. CONCLUSIONS

This paper examined some practical issues related to the proposed MPL(ambda)S approach to extending the MPLS control plane to the optical transport network. In terms of signaling and routing protocols, it appears that extensions to existing MPLS protocols are possible to include the optical network. For support of diffserv, no significant advantage or disadvantage was found with MPL(ambda)S. Finally, we believe that the integrated control approach of MPL(ambda)S may offer potential advantages for fault recovery compared to the usual layered approach.

We did not address the larger question whether the MPL(ambda)S is the best or most practical approach to integrate an automatic switched optical network with Internet protocols, but it is clearly worthwhile to investigate as a practical approach to IP/WDM integration.

REFERENCES

1. H. Yoshimura, K-I. Sato, N. Takachio, "Future photonic transport networks based on WDM technologies," *IEEE Commun. Mag.* **37**, pp. 74-81, February 1999.
2. ITU-T Draft Rec. G.872, "Architecture of optical transport networks," Geneva, Feb. 1999.
3. ITU-T Draft Rec. G.ason, "Architecture for the automatic switched optical network (ASON)," Kyoto, Feb. 2000.
4. G. Bernstein, et al., "Optical domain service interconnect (ODSI) functional specification," ODSI Coalition, Aug. 2000.
5. B. Rajagopalan, et al., "IP over optical neworks: architectural aspects," *IEEE Commun. Mag.* **38**, pp. 94-102, Sept. 2000.
6. G. Bernstein, et al., "ODSI srevice discovery and address registration," ODSI Coalition, April 2000.
7. K. Arvind, et al., "Optical domain service interconnect (ODSI) signaling control specification," ODSI Coalition, Sept. 2000.
8. A. Viswanathan, N. Feldman, Z. Wang, R. Callon, "Evolution of multiprotocol label switching," *IEEE Commun. Mag.* **36**, pp. 165-173, May 1998.
9. E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol label switching architecture," Internet draft, work in progress.
10. R. Callon, et al., "A framework for multiprotocol label switching," Internet draft, work in progress.
11. D. Awduche, et al., "Multi-protocol lambda switching: combining MPLS traffic engineering control with optical crossconnects," Internet draft, work in progress.
12. J. Luciani, et al., "IP over optical networks – a framework," Internet draft, work in progress.
13. D. Awduche, et al., "Extension to RSVP for traffic engineering," Internet draft, work in progress.
14. D. Awduche, et al., "Extensions to RSVP for LSP tunnels," Internet draft, work in progress.
15. B. Jamoussi, ed., "Constraint-based LSP setup using LDP," Internet draft, work in progress.
16. L. Andersson, et al., "LDP specification," Internet draft, work in progress.
17. P. Ashwood-Smith, et al., "Generalized MPLS – signaling functional description," Internet draft, work in progress.
18. Y. Fan, et al., "Extensions to CR-LDP and RSVP-TE for optical path set-up," Internet draft, work in progress.
19. J. Lang, K. Mitra, J. Drake, "Extensions to RSVP for optical networking," Internet draft, work in progress.
20. H. Smit, T. Li, "IS-IS extensions for traffic engineering," Internet draft, work in progress.
21. D. Katz, D. Yeung, "Traffic engineering extensions to OSPF," Internet draft, work in progress.
22. K. Kompella, et al., "OSPF extensions in support of MPL(ambda)S," Internet draft, work in progress.
23. K. Kompella, et al., "IS-IS extensions in support of MPL(ambda)S," Internet draft, work in progress.
24. G. Wang, et al., "Extensions to OSPF/IS-IS for optical routing," Internet draft, work in progress.
25. E. Crawley, et al., "A framework for QoS-based routing in the Internet," Internet RFC 2386, Aug. 1998.
26. S. Blake, et al., "An architecture for differentiated services", Internet RFC 2475, Dec 1998.
27. J. Heinanen, et al., "Assured forwarding PHB group," Internet RFC 2597, June 1999.
28. V. Jacobson, K. Nichols, K. Poduri, "An expedited forwarding PHB," Internet RFC 2598, June 1999.
29. B. Teitelbaum, et al., "Internet2 QBone: building a testbed for differentiated services," *IEEE Network* **13**, pp. 8-16, Sept. 1999.
30. O.Gerstel, R. Ramaswami, "Optical layer survivability: a services perspective," *IEEE Commun. Mag.* **38**, pp. 104-113, March 2000.
31. T. Chen, T. Oh, "Reliable services in MPLS," *IEEE Commun. Mag.* **37**, pp. 58-62, Dec. 1999.
32. N. Ghani, S. Dixit, T-S. Wang, "On IP-over-WDM integration," *IEEE Commun. Mag.* **38**, pp. 72-84, March 2000.