Generalized Multiprotocol Label Switching: An Overview of Routing and Management Enhancements

Ayan Banerjee, John Drake, Jonathan P. Lang, and Brad Turner, Calient Networks

Kireeti Kompella, Juniper Networks

Yakov Rekhter, Cisco Systems

ABSTRACT

Generalized multiprotocol label switching, also referred to as multiprotocol lambda switching, supports not only devices that perform packet switching, but also those that perform switching in the time, wavelength, and space domains. The development of GMPLS requires modifications to current signaling and routing protocols. It has also triggered the development of new protocols such as the Link Management Protocol. In this article, we present the traffic engineering enhancements to the Open Shortest Path First Internet routing protocol [1] and IS-IS Intradomain Routing Protocol ([2, 3]), two popular routing protocols, to support GMPLS. We present the concepts of generalized interfaces, label-switched path hierarchy, and link bundling intended to improve GMPLS scalability. We also discuss the Link Management Protocol which can be used to make the underlying links more manageable.

INTRODUCTION

Today's synchronous optical network/synchronous digital hierarchy (SONET/SDH) transport network infrastructure provides a guaranteed level of performance and reliability for voice calls and leased lines, the predominant traffic types prior to 1995. Since 1995, however, there has been a dramatic increase in data traffic, driven primarily by the explosive growth of the Internet as well as the proliferation of virtual private networks (VPNs). Just before the turn of the millennium, the amount of data traffic worldwide surpassed voice traffic and will continue to outpace voice for years to come. At the same time, there is increasingly strong demand from customers to keep the cost of networking down. The need to carry more traffic, combined with the need to minimize the cost of carrying this traffic, results in a situation where service providers (SPs) need solutions that enable them to carry a large volume of traffic in the most cost-efficient manner.

Delivering solutions that enable SPs to carry a large volume of traffic in a cost-efficient manner has proven to be a challenge within the current data network architecture. Today's data networks typically have four layers: IP for carrying applications and services, asynchronous transfer mode (ATM) for traffic engineering, SONET/SDH for transport, and dense wavelength-division multiplexing (DWDM) for capacity (Fig. 1). This architecture has been slow to scale for very large volumes of traffic, and at the same time fairly cost-ineffective. Multilayer architectures typically suffer from the lowest common denominator effect where any one layer can limit the scalability of the entire network, as well as add to the cost of the entire network.

Effective transport should optimize the cost of data multiplexing as well as data switching over a wide range of traffic volumes. DWDM is a cost-efficient multiplexing technique that offers significant technical advantages. DWDM increases the bandwidth-carrying capacity of a single optical fiber by effectively creating multiple virtual fibers, each carrying multigigabits of traffic per second, on a single fiber. This provides multifold increase in bandwidth while leveraging the existing fiber infrastructure. Likewise, optical cross-connects (OXCs) are likely to emerge as the preferred option for switching multigigabit or even terabit data streams, since electronic per-packet processing is avoided.

It is widely expected that the predominant traffic carried over data networks will be IP-based, which suggests that the development of fast router technologies is essential for the aggregation of slower data streams into streams suitable for OXCs. Likewise, IP packet-based statistical multiplexing is likely to be the predominant multiplexing technology for data streams smaller than those suitable for DWDM.

0163-6804/01/$10.00 © 2001 IEEE

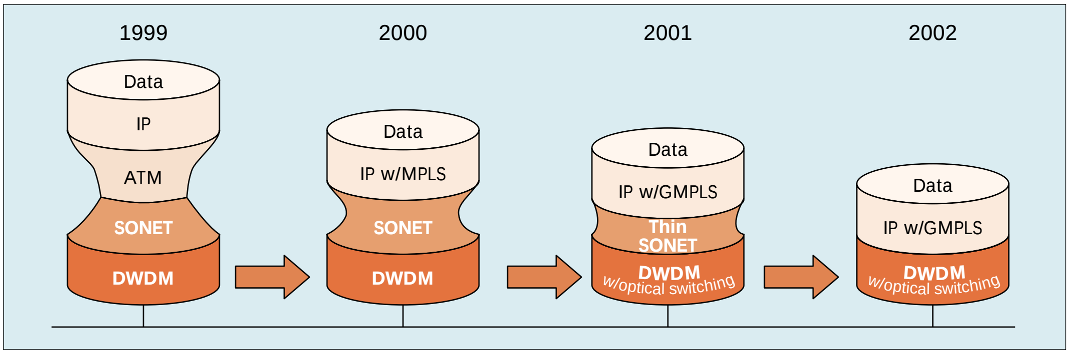

Figure 1. *The evolution toward photonic networking.*

Sustaining the advantages of flexible control models are the challenges of rapid provisioning, routing, monitoring, and efficient restoration. Reliability requirements make these challenges of paramount importance to photonic networks.

As the capabilities of both routers and OXCs grow rapidly, the high data rates of optical transport suggest the distinct possibility of bypassing the SONET/SDH and ATM layers. In order to bypass these layers their necessary functions must move directly to the routers, OXCs, and DWDMs. In the end, this results in a simpler, more cost-efficient network that will transport a wide range of data streams and very large volumes of traffic.

Operationally, this new optical architecture can be viewed from two vantage points, best described as an *overlay model* and a *peer model*. The overlay model (Fig. 2a) hides details of the internal network, resulting in two separate control planes with minimal interaction between them. One control plane operates within the core optical network, and the other between the core and the surrounding edge devices (called the *user-network interface*, UNI). The edge devices support lightpaths that are either dynamically signaled across the core optical network or statically provisioned without seeing inside the core's topology. This is very similar to current combined IP/ATM networks. The overlay model imposes administrative control boundaries between the core and edge by effectively hiding the contents of the core.

In the peer model (Fig. 2b), a single instance of the control plane spans an administrative domain consisting of the core optical network and the surrounding edge devices. This allows the SP edge devices to see the topology of the core network. Although an $O(N^2)$ mesh of point-to-point connections is still required if full connectivity between the edge devices is needed, it is used exclusively for the purpose of data forwarding. As far as the routing protocols are concerned, each edge device is adjacent to the photonic switch to which it is attached, rather than to the other edge devices. With $O(N)$ routing adjacencies a full mesh of $O(N^2)$ forwarding paths can be supported. This allows the routing protocols to scale to a much larger network.

Another approach is a hybrid model that combines both the peer and overlay models. Some edge devices serve as peers to the core network and share the same instance of a common control plane with the core network. Other edge devices could have their own control plane (or a separate instance of the control plane used by the core network), and interface with the core network through the UNI. This represents the highly desirable solution offering carriers and SPs substantial flexibility to deploy the most cost-effective model for their needs, be it peer, overlay, or some hybrid of these models.

Functionally, the peer model forms a superset of the overlay model. That is, the set of functions required to support the overlay model is a subset of the set of functions required to support the peer model. An overlay model can be derived from a peer model by administratively disabling topology sharing while preserving the connection signaling functions. This observation suggests that rather than having one set of protocols to support the overlay model and another to support the peer model, one suite of control plane protocols with enough flexibility to support both models would be the most efficient approach.

Sustaining the advantages of flexible control

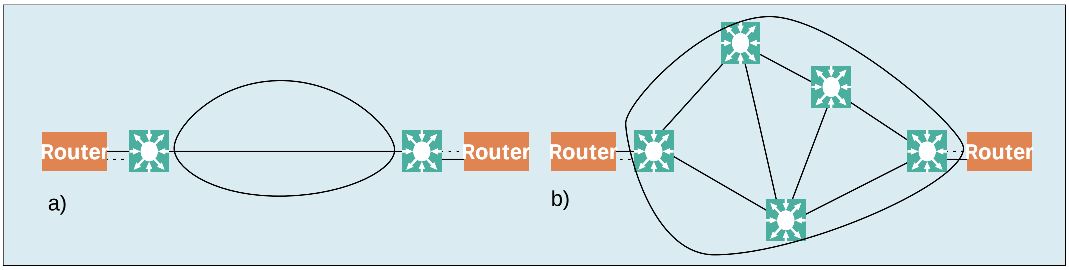

Figure 2. *a) Network level abstraction. The overlay model hides the internals of the optical network, essentially forming an optical cloud, and provides wavelength services to clients (e.g., routers, ADMs, and ATM switches) that reside at the edges of the network; b) link-level abstraction. The peer model opens the internals of the optical cloud, allowing the edge devices to participate in routing decisions and eliminating the artificial barriers between the transport and routing domains.*

Some modifications and additions are required to the MPLS routing and signaling protocols to adapt to the peculiarities of photonic switches. These are being standardized by the IETF under the umbrella of Generalized MPLS.

models are the challenges of rapid provisioning, routing, monitoring, and efficient restoration. Reliability requirements make these challenges of paramount importance to photonic networks. While several vendors are developing proprietary routing and signaling protocols to enable automatic provisioning, such implementations are unlikely to interoperate in multivendor deployments. Providers of "best-of-breed" optical solutions require the integration of photonic switches into a heterogeneous optical network, which will eventually combine next-generation equipment with legacy equipment. A common standardized control plane must be used to communicate between the various elements. Fortunately, such a standard is well known in the industry.

Over the last few years, IP routing has evolved to include new functionality under the umbrella of multiprotocol label switching (MPLS) [4], and recent work has been done on extending MPLS as a control plane that can be used not merely with routers, but also with legacy equipment (e.g., SONET, ADMs) and newer devices like OXCs [5]. These efforts offer the necessary standardized common control plane, an essential component in the evolution of open and interoperable optical networks. First, a common control plane simplifies operations and management, which reduces the cost of operations. Second, a common control plane provides a wide range of deployment scenarios ranging from overlay to peer, where the overlay model is realized by using just a subset of the functionality provided by the peer model. A common control plane allows the choice of peer or overlay (or a combination of both) to be driven by business considerations, rather than constrained by technology. At the same time, building the common control plane from proven signaling and routing avoids "reinventing the wheel" for protocol development, thereby minimizing risk while reducing time to market.

Some modifications and additions are required to the MPLS routing and signaling protocols to adapt to the peculiarities of photonic switches. These are being standardized by the Internet Engineering Task Force (IETF) under the umbrella of generalized MPLS, which can be summarized as follows:

- A new Link Management Protocol (LMP) designed to address issues related to link management in optical networks using photonic switches [6]
- Enhancements to the Open Shortest Path First/Intermediate System to Intermediate System (OSPF/IS-IS) routing protocols to advertise availability of optical resources in the network (e.g., generalized representation of various link types, bandwidth on wavelengths, link protection type, fiber identifiers) [7, 8]
- Enhancements to the Resource Reservation Protocol (RSVP)/Constraint-Based Routing Label-Distributed Protocol (CR-LDP) signaling protocols for traffic engineering purposes that allow a label-switched path (LSP) to be explicitly specified across the optical core [9]
- Scalability enhancements such as hierarchical LSP formation, link bundling, and unnumbered links

This article deals with enhancements made to the routing and link management protocols in support of generalized MPLS (GMPLS). The reader is referred to [9] for the enhancements to the signaling protocols.

MPLS BACKGROUND

MPLS is based on the following set of ideas:

- Forwarding information (label) separate from the content of IP header
- A single forwarding paradigm (label swapping), multiple routing paradigms
- Multiple link-specific realizations of the label swapping forwarding paradigm: "shim," virtual connection/path identifier (VCI/VPI), frequency slot (wavelength), time slot
- The flexibility to form forwarding equivalence classes (FECs)
- A forwarding hierarchy via label stacking

The separation of forwarding information from the content of the IP header allows MPLS to be used with devices such as OXCs, whose data plane cannot recognize the IP header. Label switch routers (LSRs) forward data using the label carried by the data. This label, combined with the port on which the data was received, is used to determine the output port and outgoing label for the data. The MPLS control plane operates in terms of the label swapping and forwarding paradigm abstraction. At the same time, the MPLS data plane allows multiple link-specific realizations of this abstraction. For example, a wavelength could be viewed as an implicit label. Finally, the concept of a forwarding hierarchy via label stacking enables interaction with devices that can support only a small label space. This property of MPLS is essential in the context of OXCs and DWDMs since the number of wavelengths (which act as labels) is not very large.

The MPLS framework includes significant applications such as constraint-based routing. Constraint-based routing is a combination of extensions to existing IP link-state routing protocols (e.g., OSPF and IS-IS) with RSVP or CR-LDP as the MPLS control plane, and the Constrained Shortest-Path-First (CSPF) heuristic. The extensions to OSPF and IS-IS allow nodes to exchange information about network topology, resource availability and even policy information. This information is used by the CSPF [10, sec. 7] heuristic to compute paths subject to specified resource and/or policy constraints. For example, either RSVP-TE [11] or CR-LDP [12] is used to establish the label forwarding state along the routes computed by a CSPF-based algorithm; this creates the LSP. The MPLS data plane is used to forward the data along the established LSPs.

Constraint-based routing is used today for two main purposes: traffic engineering and fast reroute. With suitable network design, the constraint-based routing of IP/MPLS can replace ATM as the mechanism for traffic engineering. Likewise, fast reroute offers an alternative to SONET as a mechanism for protection/restoration. Both traffic engineering and fast reroute are examples of how enhancements provided by MPLS to IP routing make it possible to bypass ATM and SONET/SDH by migrating functions provided by these technologies to the IP/MPLS control plane.

Paving a path for future evolution of MPLS technologies, as well as GMPLS enhancements, are several emerging synergies between LSRs and photonic switches, and between an LSP and an optical trail. An optical trail is an end-to-end path composed exclusively of photonic elements without optical-electronic conversions. Analogous to switching labels in an LSR, a photonic switch toggles wavelengths from an input to an output port. Establishing an LSP involves configuring each intermediate LSR to map a particular input label and port to an output label and port. Similarly, the process of establishing an optical trail involves configuring each intermediate photonic switch to map a particular input lambda and port to an output lambda and port. As in LSRs, photonic switches need routing protocols like OSPF or IS-IS to exchange link-state topology and other optical resource availability information for path computation. They also need signaling protocols like RSVP and LDP to automate the path establishment process. In the remainder of this article, we discuss the routing enhancements and link management.

ENHANCEMENTS TO ROUTING

In this section we discuss the enhancements made in support of GMPLS. These enhancements have been made to address some of the challenges of using MPLS to control optical and SONET/SDH time-division multiplexing (TDM) networks. These include:

1)The MPLS label space is comparatively large (one million per port), whereas there are a relatively limited number of lambdas and TDM channels (tens to hundreds per port today, scaling to thousands over the next few years).

2)MPLS LSPs can be allocated bandwidth from a continuous spectrum, whereas optical/TDM bandwidth allocation is from a small discrete set of values.

3)Today there are rarely more than 10 parallel links between a pair of nodes. To handle the growth of traffic providers will need to deploy hundreds of parallel fibers, each carrying hundreds of lambdas between a pair of network elements. This in turn raises three sub-issues:

   a. The overall number of links in an optical/TDM network can be several orders of magnitude larger than that of an MPLS network.

   b. Assigning IP addresses to each link in an MPLS network is not particularly onerous; assigning IP addresses to each fiber, lambda, and TDM channel is a serious concern, because of both the scarcity of IP addresses and the management burden.

   c. Identifying which port on a network element is connected to which port on a neighboring network element is also a major management burden and highly error-prone.

4)Fast fault detection and isolation, and fast failover to an alternate channel are needed.

5)The user data carried in the optical domain is transparently switched to increase the efficiency of the network. This necessitates transmitting control plane information decoupled from user data.

Note that all of the above are issues for MPLS networks as well; however, these issues are immediate and pressing for optical networks. The following sections describe how GMPLS addresses these issues.

To help understand the routing enhancements described below, we start with a brief description of link state protocols such as IS-IS and OSPF. Consider a network as a directed graph whose nodes are network elements (MPLS switches, cross-connects, etc.) and whose edges are links (fibers, cables, etc.). Each edge in the graph has associated attributes such as IP addresses, cost, and unreserved bandwidth. A link state protocol allows all the nodes to dynamically coordinate a coherent up-to-date picture of this graph, including the attributes of each edge. This picture of the graph is referred to as the *link state database*. Once the link state database is synchronized among all participating routers, each router uses the database to construct its own forwarding table. When a packet arrives at a router, the forwarding table is then consulted to determine how to forward the packet. Should the status of any link be changed, including adding or removing links, the link state database must be resynchronized, and all of the routers must recalculate their forwarding tables using the updated information in the link state database.

LSP HIERARCHY

LSP hierarchy is the notion that LSPs can be nested inside other LSPs, giving rise to a hierarchy of LSPs. This is achieved by considering an LSP as a link in the IS-IS or OSPF link state database. This simple notion offers a solution to issues 1 and 2 above.

MPLS LSPs that enter the optical transport domain at the same node and leave the domain at the same node may be *aggregated* and tunneled within a single optical LSP. This aggregation helps to conserve the number of lambdas used by the MPLS domain.

LSP hierarchy also helps deal with the discrete nature of optical bandwidth. When an optical LSP is set up, it gets a discrete bandwidth (say 2.488 Gb/s). However, when this optical LSP is treated as a link, that link's bandwidth need no longer be discrete. A 100 Mb/s MPLS LSP that crosses the optical transport domain can be tunneled through the optical LSP, leaving 2.388 Gb/s for other MPLS LSPs. Allocating an entire 2.488 Gb/s for every MPLS LSP that crosses the optical domain would be impractical.

A natural hierarchy exists that dictates the order in which LSPs can be nested. This hierarchy is based on the multiplexing capability of the LSP types. Note that LSPs always start and terminate on similar equipment (e.g., a lambda LSP originates and terminates on a device that supports lambdas). At the top of this hierarchy are nodes that have fiber-switch-capable (FSC) interfaces, followed by nodes that have lambda-

LSP hierarchy is the notion that LSPs can be nested inside other LSPs, giving rise to a hierarchy of LSPs. This is achieved by considering an LSP as a link in the IS-IS or OSPF link state database.

The reduction of management effort of configuring IP addresses, tracking allocated IP addresses, and dealing with the occasional duplicate address allocation is a significant savings, especially in the context of optical networks with their large number of links.

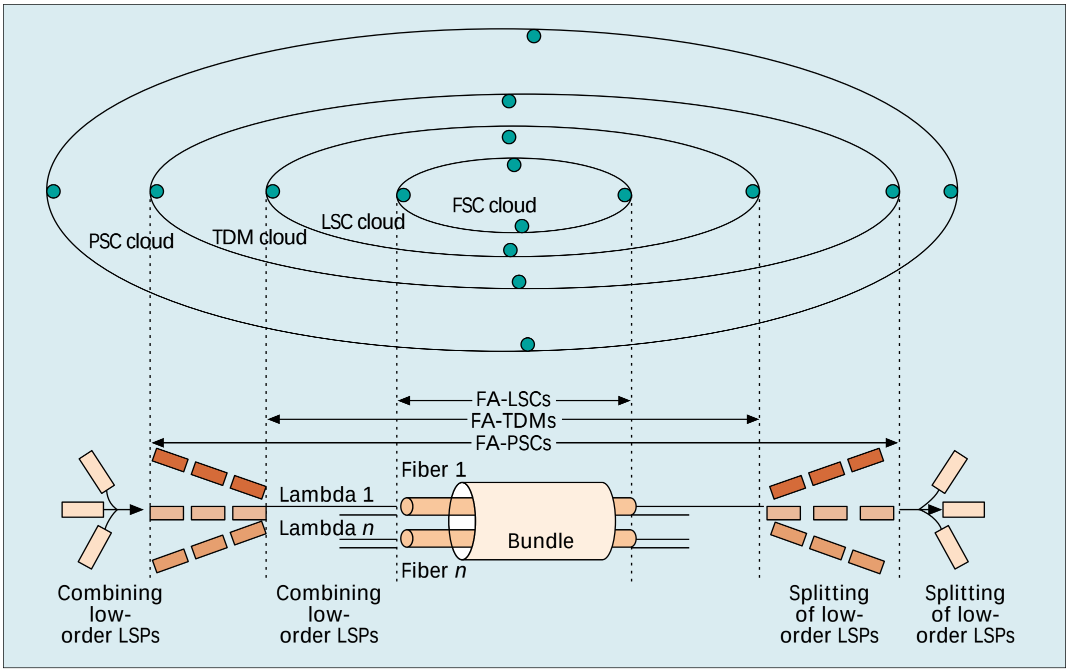

Figure 3. *The circles represent the interfaces on devices of similar nature. For example, the FSC circle consists of photonic cross-connect switches capable of switching entire fibers. The LSC circle consists of photonic or OXCs capable of switching wavelengths. The TDM circle consists of ATM or SONET cross-connects. Finally, the PSC circle consists of routers. Low-order LSPs are formed across the cloud of higher-order interfaces and announced into the IGP. This allows low-order LSPs to be grouped together and hierarchically tunneled through higher-order LSPs. Multiple PSC-LSPs are tunneled within a TDM-LSP, multiple TDM-LSPs are grouped and tunneled within an LSC-LSP, and so on. At the other end of the cloud they are split appropriately.*

switch-capable (LSC) interfaces, followed by nodes that have TDM-capable interfaces, followed by nodes with packet-switch-capable (PSC) interfaces. In a typical configuration (Fig. 3), the core cloud of FSC interfaces/nodes are connected to an outer cloud of LSC interfaces/nodes. These are connected to an outer cloud of TDM-capable nodes, which are finally connected to routers. Dissemination of this information is essential so that the paths within the clouds can be generated automatically with minimal manual configuration.

An LSP (Fig. 3) that starts and ends on a PSC interface can be nested (together with other LSPs) into an LSP of type TDM that starts and ends on a TDM interface (i.e., within nodes depicted as a TDM cloud in the picture). This TDM-LSP, in turn, can be nested (together with other TDM-LSPs) into an LSC-LSP that starts and ends on an LSC interface, which in turn can be nested (together with other LSC-LSPs) into an LSP that starts and ends on an FSC interface.

The LSPs appear as new link types in the IS-IS/OSPF routing database; the new link types are compatible with the existing flooding methods used for sharing conventional link information. Because of this flooding, each node has an identical link state database containing information about not just conventional links, but LSPs as well. A node, when performing path computation, is thus able to use not only conventional links, but also LSPs with appropriate constraints (e.g., the order of LSP tunneling has to be maintained). For more details, the reader is referred to [13]. This allows for *hierarchical* scaling of the link state database. Once a path is computed, the node uses RSVP/CR-LDP signaling mechanisms to establish label binding along the path. The details of the signaling mechanisms are beyond the scope of this article (see [9]).

LINK BUNDLING

As mentioned above, the link state database consists of all the nodes and links in a network, along with the attributes of each link. In light of issue 3a above, the link state database for an optical network can easily be several orders of magnitude bigger than that for an MPLS network.

To address this issue, we aggregate the link attributes of several parallel links of similar characteristics, and assign these aggregated attributes to a single "bundled" link. In so doing, the size of the link state database is reduced by a large factor, leading to vastly improved scaling of the link state protocol. The details of how links are bundled (i.e., how link attributes are aggregated) can be found in [14].

By summarizing the attributes of several links into one bundled link, some information is lost; for example, with a bundle of SONET links the switching capability of the link interfaces (OC-12, OC-48, OC-192) are flooded; however, the number of such interfaces and the exact time slots used are not announced. However, the benefit of improved scalability will significantly outweigh the value of the information lost. In addition, while the link state protocol carries a single bundled link, signaling requires that individual component links be identified. LMP [6] offers a means to accomplish this.

UNNUMBERED LINKS

All the links in an MPLS network are typically assigned IP addresses. When a path is computed through the network, the links that constitute the path are identified by their IP addresses; this information is conveyed to the signaling protocol, which then sets up the path. Thus, it would seem that every link must have an IP address. However, issue 3b describes the difficulty of doing this. Unnumbered links are used to resolve this problem; however, if an IP address is not used to identify a link, an alternative must be substituted.

What is required is a unique means of identifying links in a network; the task may be broken down into two steps. First, a mechanism is required to uniquely identify each node in the network; then each link emanating from that node is identified. Each node in the network is identified by a unique router ID; what remains is the latter problem of identifying the links emanating from a particular node. A solution to this problem, and the information that needs to be flooded by the routing protocols (OSPF and IS-IS) and that which needs to be communicated by the use of the signaling protocol, are described in [15].

Ultimately, each network node numbers its interfaces locally. The tuple [router ID, link number] serves as the identification for a link. The reduction of management effort in configuring IP addresses, tracking allocated IP addresses, and dealing with the occasional duplicate address allocation is a significant savings, especially in the context of optical networks with their large numbers of links.

LINK MANAGEMENT PROTOCOL

A consequence of generalizing MPLS to encompass non-PSC links is that a label is no longer an abstract identifier, but must now be able to map to time slots, wavelengths, and physical resources such as the ports of a switch. This requires that the association of these physical labels be created between adjacent nodes. For IGP scaling purposes, multiple links between nodes may be combined into a single bundled link as described above. LMP [6] runs between adjacent nodes and is used for both link provisioning and fault isolation. A key service provided by LMP is the associations between neighboring nodes for the component link IDs that may in turn be used as labels for physical resources. These associations do not have to be configured manually, a potentially error-prone process. A significant improvement in *manageability* accrues because the associations are created by the protocol itself. This addresses issue 3c.

Within a bundled link, the component links and associated control channel need not be transmitted over the same physical medium. LMP allows for decoupling of the control channel from the component links. For example, the control channel could be transmitted along a separate wavelength or fiber, or over a separate Ethernet link between the two nodes. This addresses issue 5. A consequence of allowing the control channel for a link to be physically diverse from the component links is that the health of a control channel of a link does not correlate to the health of the component links, and vice versa. Furthermore, due to the transparent nature of photonic switches, traditional methods can no longer be used to monitor and manage links.

LMP is designed to provide four basic functions for a node pair: control channel management, link connectivity verification, link property correlation, and fault isolation. Control channel management is used to establish and maintain connectivity between adjacent nodes, and consists of a lightweight keep-alive Hello protocol that is transmitted over the control channel. The link verification procedure is used to verify the physical connectivity of the component links, which is paramount due to the all too human-error-prone cabling process. The *LinkSummary* message of LMP provides the correlation function of link properties (e.g., link IDs, protection mechanisms, and priorities) between adjacent nodes. This is done when a link is first brought up and may be repeated any time a link is up and not in the verification procedure. Finally, LMP provides a mechanism to isolate link and channel failures in both opaque and transparent networks, independent of the data format, which enables issue 4.

Within a bundled link, the component links and associated control channel do not need be transmitted over the same physical medium. LMP allows for decoupling of the control channel from the component links.

CONCLUSION

GMPLS will be an integral part of deploying the next generation of data networks. It provides the necessary bridges between the IP and photonic layers to allow for interoperable and scalable parallel growth in the IP and photonic dimensions. With GMPLS dynamically bridging the gap between the traditional transport infrastructure and the IP layers, the path is being paved for rapid service deployment and operational efficiencies, as well as increased revenue opportunities. The necessary provisions have been put in place to support a smooth transition from a traditional segregated transport and service overlay model to a more unified peer model. The functionality afforded by GMPLS, its associated generalized notion of an LSP hierarchy, and bundling creates sufficient flexibility in support of either the segregation or unification of almost any operational paradigm desired by an operator. By streamlining support for multiplexing and switching in a hierarchical fashion and combining the flexible intelligence of MPLS traffic engineering, the business value of optical switching GMPLS will prove essential in any solution that aims to enable large volumes of traffic in a cost-efficient manner for service providers.

REFERENCES

[1] J. Moy, RFC 2328, “OSPF Version 2,” IETF.

[2] D. Oran, “OSI IS-IS Intra-Domain Routing Protocol,” IETF RFC 1142.

[3] R. Callon,, “Use of OSI IS-IS for Routing in TCP/IP and Dual Environments,” IETF RFC 1195.

[4] E. Rosen, A. Viswanathan, and R., Callon, “Multiprotocol Label Switching Architecture,” draft-ietf-mpls-arch-07.txt, July 2000, work in progress.

[5] D. Awduche et al., “Multiprotocol Lambda Switching: Combining MPLS Traffic Engineering with Optical Crossconnects,” Internet draft, draft-awduche-mpls-te-optical-02.txt, Mar. 2000, work in progress.

[6] J. P. Lang *et al.*, “Link Management Protocol,” Internet draft, draft-ietf-mpls-lmp-00.txt, Aug. 2000, work in progress.

GMPLS will be an integral part in the deployment of the next generation of data networks. It provides the necessary bridges between the IP and photonic layers to allow for interoperable and scalable parallel growth in the IP and the photonic dimensions.

[7] K. Kompella *et al.*, "IS-IS Extensions in Support of generalized MPLS," Internet Draft, draft-ietf-gmpls-extensions-00.txt, Sept. 2000, work in progress.

[8] K. Kompella *et al.*, "OSPF Extensions in Support of MPL(ambda)S," Internet Draft, draft-ompls-ospf-extensions-00.txt, July 2000, work in progress.

[9] P. Ashwood-Smith *et al.*, "Generalized MPLS — Signaling Functional Description," Internet Draft, draft-ietf-mpls-generalized-signaling-00.txt, Nov. 2000, work in progress.

[10] D. Awduche *et al.*, RFC 2702 "Requirements for Traffic Engineering Over MPLS," IETF.

[11] D. Awduche *et al.*, "RSVP-TE: Extensions to RSVP for LSP Tunnels," draft-ietf-mpls-rsvp-lsp-tunnel-07.txt, Aug. 2000, work in progress.

[12] L. Andersson *et al.*, "Label Distribution Protocol Specification," (draft-ietf-mpls-ldp-11), Aug. 2000, work in progress.

[13] K. Kompella and Y. Rekhter, "LSP Hierarchy with MPLS TE," Internet Draft, draft-ietf-mpls-lsp-hierarchy-00.txt, July 2000, work in progress.

[14] K. Kompella, Y. Rekhter, and L. Berger, "Link Bundling in MPLS Traffic Engineering," Internet Draft, draft-kompella-mpls-bundle-03.txt, Sept. 2000, work in progress.

[15] K. Kompella and Y. Rekhter, "Signalling Unnumbered Links in RSVP-TE," Internet Draft, draft-ietf-mpls-rsvp-unnum-02.txt, Sept. 2000, work in progress.

BIOGRAPHIES

AYAN BANERJEE (abanerjee@calient.net) is systems engineer at Calient Networks, where he assists in network and switch architecture design and analysis. Prior to joining Calient, he was a network scientist at BBN Technologies where he was involved in the design, analysis, and implementation of networking architectures and protocols. He received a B.Tech degree from the Indian Institute of Technology, Kharagpur, and an M.S. and Ph.D. from the University of California at Santa Barbara. He received the General Affiliates Dissertation Award for his work on the design and performance evaluation of dynamic communication algorithms on wireline and wireless networks and has also developed proprietary collision resolution algorithms for the multi-access cable modem network. His research interests include switch architectures, constraint-based routing, topology control, dynamic communication algorithms, and software radios.

JOHN DRAKE [SM] (jdrake@calient.net) is chief architect at Calient Networks, where he is responsible for the development of the company's routing and networking architectures, and technical management of vendor partnerships. Before joining Calient, he was distinguished engineer in the Office of CTO at FORE Systems, a member of the senior technical staff at IBM, and lead architect for IBM's Broadband Network Services. Some of his most significant accomplishments include the development of FORE's MPLS strategy, and leading the group that developed FORE's PNNI, MPOA, and LANE v2.0 implementations. He acquired 25 broadband networking patents during his tenure with FORE and IBM (e.g., distributed management communications network; quality of service management for source routing multimedia packet networks). He holds a B.S. in physics and mathematics from Carnegie-Mellon University.

JONATHAN P. LANG (jplang@calient.net) is systems engineer of network architecture at Calient Networks, where he focuses on the design of Calient's network and routing architecture. He is actively involved in the MPLS working group of the IETF, and has submitted numerous drafts extending MPLS to support intelligent optical networks. He has made several contributions to the OIF, and participates in the organization's architecture and signaling groups. Before joining Calient he worked for Whitetree, Inc., developing rate allocation algorithms for ATM. As a researcher and member of the Multidisplinary Optical Switching Technology (MOST) center at the University of California, Santa Barbara, he was responsible for developing the switch architecture and networking protocols for a new optical packet switch. He received an M.S. and a Ph.D. in electrical engineering from the University of California, Santa Barbara, and a B.S. in electrical engineering from the University of California, San Diego.

BRAD TURNER (bturner@calient.net) is a senior product manager for Calient Networks, where he works on routing and signaling software for Calient's photonic switch platform. Prior to joining Calient, Brad held product management positions at Cisco Systems and 3Com Corporation, where he specialized in layer 3 routing and layer 2 switching systems. He has a B.S. in computer science from Birmingham-Southern College.

KIREETI KOMPELLA (kireeti@juniper.net) is a senior protocols engineer at Juniper Networks. His primary interests are routing protocols, in particular, MPLS. His current focus is on traffic engineering, constraint-based routing, MPLS extensions for the control of non-packet networks, and other MPLS applications. In previous lives, his interests included file system design and number theory. He obtained his B.S. (EE) and M.S. (CS) degrees at IIT Kanpur, and his Ph.D. at the University of Southern California.

YAKOV REKHTER (yakov@cisco.com) works at Cisco Systems, Inc., where he is a Cisco Fellow. He is one of the leading designers of tag switching, BGP/MPLS VPNs, and MPLS traffic engineering. He is also one of the leading designers of Border Gateway Protocol (BGP). Among his most recent interests is multiprotocol lambda switching (a.k.a. GMPLS). He is the author of over 40 RFCs, as well as numerous presentations and articles on TCP/IP and the Internet.