Framework for MPLS-Based Control of Optical SDH/SONET Networks

Greg Bernstein, Ciena Corporation
Eric Mannie, Ebone
Vishal Sharma, Metanoia, Inc.

Abstract

The suite of protocols that defines Multi-Protocol Label Switching is in the process of enhancement to generalize its applicability to the control of optical networks. One area of prime consideration is to use these generalized MPLS protocols in upgrading the control plane of optical transport networks. This article describes those extensions to MPLS directed toward controlling SDH/SONET networks. SDH/SONET networks are ideal candidates for this process since they possess a rich multiplex structure, and a variety of protection/restoration options are well defined and widely deployed. We discuss the extensions to MPLS routing protocols to disseminate information needed for transport path computation and network operations, and the extensions to MPLS label distribution protocols needed for provisioning of transport circuits.

SDH [1]/SONET [2] is the predominant optical networking technology used in today's wide and metropolitan area networks and forms the underlying infrastructure for most data and voice services. SDH/SONET signals are transported optically and may be switched either optically (as a whole signal) or electronically (as a whole or in parts). While SDH/SONET and DWDM technology has advanced to meet the raw capacity requirements of today's fast growing data service transport, the control plane for such optical networks remains relatively primitive, and has not advanced as quickly [3].

For example, a lack of multi-vendor interoperability in provisioning new end-to-end SONET/SDH circuits leads to long provisioning times. Also, a lack of multi-vendor interoperability in automated topology discovery and resource status can lead to inefficient utilization of network resources.

A major advantage of the MPLS architecture for use as a general network control plane is its clear separation between the data forwarding plane, the signaling (connection control) plane, and the routing (topology discovery/resource status) plane. This article describes how the MPLS protocol suite is being extended to Generalized MPLS (GMPLS) [1] and specialized [5] to dynamically establish, maintain, and tear down time division multiplexed (TDM) optical circuits. The overall structure and required functionality of an IP-centric control plane can be found in [3].

This article first reviews MPLS and optical TDM technology (SONET/SDH). The extensions to the GMPLS routing protocols to support SONET/SDH are discussed next, followed by a discussion of the signaling attributes that need to be conveyed by a GMPLS label distribution protocol. The article concludes by examining some outstanding issues.

MPLS Overview

Packet-based MPLS uses labels to make forwarding decisions at the network nodes, in contrast to traditional destination-based hop-by-hop forwarding used in IP networks. An MPLS [6] network consists of MPLS nodes called label switch routers (LSR) connected by circuits called label switched paths (LSP). Border LSRs in an MPLS cloud act either as ingress or egress LSRs depending on the direction of the traffic being forwarded. MPLS allows the establishment of LSPs between ingress and egress LSRs. Each LSP is associated with a forwarding equivalence class (FEC), which may be thought of as a set of packets that receive identical forwarding treatment at an LSR (e.g., the set of destination addresses lying in a given address range). To establish an LSP, a label distribution protocol, i.e., a signaling protocol such as LDP/CR-LDP or RSVP-TE, is required. Between two adjacent LSRs a short, fixed-length identifier called a label (significant only between the two LSRs) locally identifies an LSP. The signaling protocol is responsible for the inter-node communication that assigns and maintains these labels.

When a packet enters an MPLS-based packet network, it is classified according to its FEC and, possibly, additional rules that together determine the LSP along which the packet is sent. For that purpose, the ingress LSR attaches an appropriate label to the packet and forwards the packet to the next hop. The label may be attached to a packet either in the form of a header encapsulating the packet (the "shim" header) or it may be written in the VPI/VCI field (or DLCI field) of the layer 2 encapsulation of the packet. In SDH/SONET networks a label is simply associated with a segment of a circuit, and is used in the signaling plane to identify this segment (e.g., a time-slot) between two adjacent nodes.

When a packet reaches a core packet LSR, this LSR uses the label as an index into a forwarding table to determine the next hop and the corresponding outgoing label, writes the new label into the packet, and forwards the packet to the next hop. When the packet reaches the egress LSR, the label is removed and the packet is forwarded using adequate forwarding, such as normal IP lookups. We will see that for a SDH/SONET network these operations do not occur in quite the same way.

0890-8044/01/$10.00 © 2001 IEEE

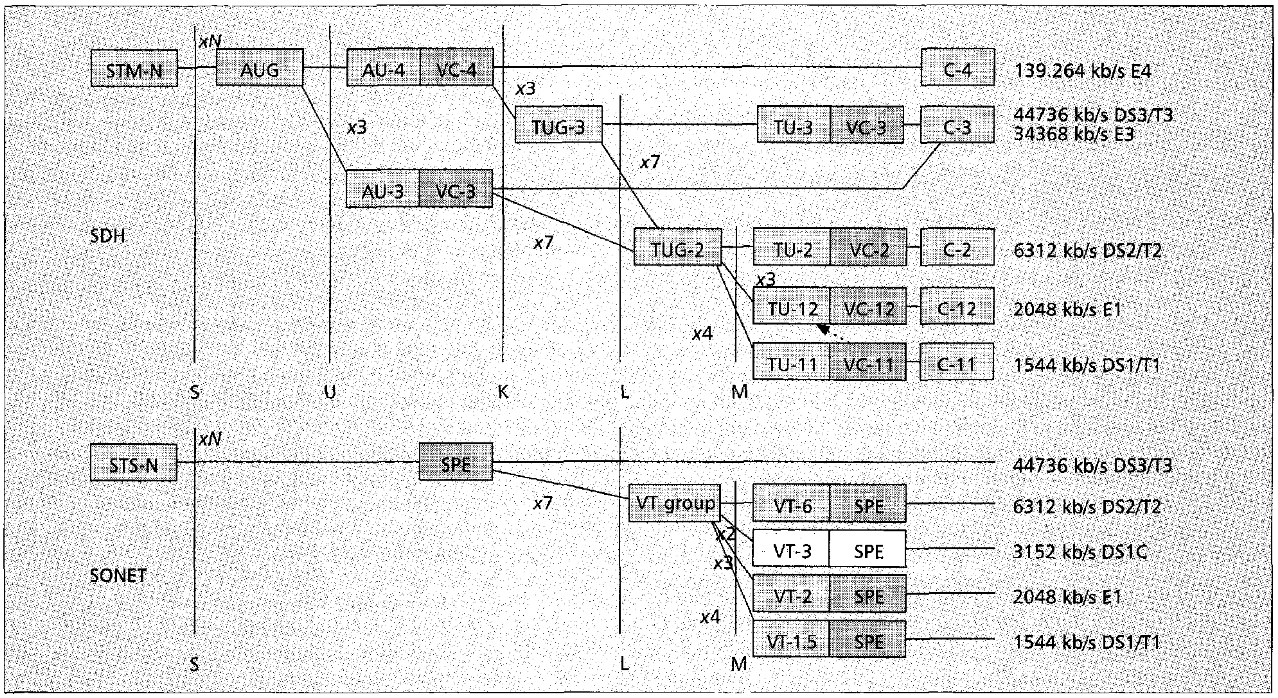

■ **Figure 1.** *SDH and SONET multiplexing structure and typical PDH payload signals.*

SDH / SONET Overview

There are currently two different multiplexing technologies in use in optical networks: wavelength division multiplexing (WDM) and time division multiplexing (TDM). This article focuses on TDM technology. SDH [1] and SONET [2] are two TDM standards widely used by operators to transport and multiplex different tributary signals over optical links, thus creating a multiplexing structure that we call the *SDH/SONET multiplex*. These two standards have several similarities, and to some extent SONET can be viewed as a subset of SDH. Internetworking between the two is possible using gateways.

The fundamental signal in SDH is the STM-1, which operates at a rate of about 155 Mb/s, while the fundamental signal in SONET is the STS-1, which operates at a rate of about 51 Mb/s. These two signals are made of contiguous frames that consist of a *transport overhead* (header) and a *payload*. To solve synchronization issues, the actual data is not directly transported in the payload but rather in another internal frame that floats over two successive SDH/SONET payloads, and is named a virtual container (VC) in SDH and a synchronous payload envelope (SPE) in SONET.

The SDH/SONET architecture identifies three different layers, each of which corresponds to one level of communication between SDH/SONET equipment. These are, starting with the lowest, the regenerator section/section layer, the multiplex section/line layer, and the path layer. Each of these layers has its own overhead (header). The transport overhead of a SDH/SONET frame is mainly subdivided into two parts that contain the regenerator section/section overhead and the multiplex section/line overhead. In addition, a pointer (in the form of the H1, H2, and H3 bytes) indicates the beginning of the VC/SPE in the payload.

The VC/SPE itself is made up of a header (the path overhead) and a payload. This payload can itself be subdivided into sub-elements (signals) in a fairly complex way. In SDH an STM-1 frame itself may contain either one VC-4 or three multiplexed VC-3s. The structure of the SDH/SONET multiplex is shown in Fig. 1.

The leaves of these multiplex structures are time slots (positions) of different sizes that can contain tributary signals (e.g., E1, E3, etc.) which are mapped into the leaves using standardized mapping rules. What is important for MPLS-based control of SONET/SDH is to identify the elements that can be switched from an input multiplex on one interface to an output multiplex on another interface. These elements are either those that can be re-aligned via a pointer, i.e., a VC-x in the case of SDH and a SPE in the case of SONET, or those, such as TUG-2, TUG-3, or STM-*N*, for which proper pointer processing can be done for each constituent of the group.

An STM-*N*/STS-*N* signal is formed from *N* x STM-1/STS-1 signals via byte interleaving. The VCs/SPEs in the *N* interleaved frames are independent and float according to their own clocking. To transport tributary signals in excess of the basic STM-1/STS-1 signal, the VCs/SPEs can be concatenated, i.e., glued together. In this case their relationship with respect to each other is fixed in time, which relieves an end system of any inverse multiplexing bonding processes. Different types of concatenations are defined, with specific rules.

For instance, the standard SONET concatenation allows the concatenation of M x STS-1 signals within an STS-N signal with $M \leq N$ and $M = 3, 12, 48, 192$. The SPEs of these M x STS-1s can be concatenated to form an STS-Mc, which is shorthand for describing an STS-M signal whose SPEs have been concatenated.

MPLS Routing for SDH/SONET

Modern transport networks based on SDH/SONET excel at interoperability in the performance monitoring (PM) and fault management (FM) areas. However, they do not interoperate in the areas of topology discovery or resource status. Although link state routing protocols, such as IS-IS and OSPF, have been used for some time in the IP world to compute destination-based next hops for IP datagrams (without routing loops), it is in providing timely topology and network status information in a distributed manner, at any network node, that they have immense value. If resource utilization information is disseminated along with the link status (as was done in

Signal type	SDH	SONET
Lower order	VC-11, VC-12, VC-2	VT-1.5 SPE, VT-2 SPE, VT-3 SPE, VT-6 SPE
Higher order	VC-3, VC-4	STS-1 SPE

■ Table 1. *SDH/SONET switched signal groupings.*

ATM's PNNI routing protocol) then a very complete picture of network status is available to a network operator for use in planning, provisioning, and operations.

Therefore, information needed to compute the path of a connection through the network must be distributed via the routing protocol. In TDM this information includes, but is not limited to, the available capacity of the network links, the switching and termination capabilities of the nodes and interfaces, and the protection properties of the link. Hence the information discussed in the following sections must be appropriately advertised via the routing protocol. Due to this increase in information transferred in the routing protocol it is important to separate the relatively static parameters concerning a link from the dynamic parameters that may be subject to frequent changes.

Switching Capabilities

The main switching capabilities that characterize a SDH/SONET end system, and thus get advertised into the link state routing protocol, are: the switching granularity, supported forms of concatenation, and the level of transparency.

Switching Granularity — The signals that can be switched within the SDH/SONET hierarchies are those that are either directly referenced via a pointer, i.e., the VCs in SDH and the SPEs in SONET, or those for which proper pointer processing may be done for their constituents, such as TUG-2 and TUG3 in SDH. These signals are subdivided into lower-order signals and higher-order signals, as shown in Table 1.

The equipment in use today has diverse capabilities. Some equipment only switches signals starting at VC-4 for SDH or STS-1 for SONET, while some equipment only allows the switching of aggregates (concatenated or not) of signals such as 16 VC-4s, i.e., a complete STM-16, and nothing below. Other equipment switches subsets of the lower-order signals as well.

Signal Concatenation Capabilities — Different types of concatenations that have been defined or proposed for SDH/SONET equipment are: standard contiguous concatenation [2], arbitrary contiguous concatenation [5], and virtual concatenation [1].Each of these has different rules governing its size, placement, and binding.

The size of SONET standard concatenated signals (STS-Mc signals), for example, must be a multiple of 3. Similarly, their starting location and interleaving are also constrained [2]. This helps with SDH compatibility (since there is no STS-1 equivalent signal in SDH) and reduces the number of connection types.

The major disadvantages of these restrictions are a restricted bandwidth assignment (which inhibits finer grained traffic engineering), and the lack of flexibility in starting time slots and their interleaving for STS-Mc (i.e., where the rest of the signal gets put in terms of STS-1 slot numbers), which requires re-grooming due to bandwidth fragmentation.

Thus, some framer manufacturers now support "flexible" or arbitrary concatenation where there are no restrictions on the size of an STS-Mc (as long as $M \leq N$) and no constraints on the STS-1 timeslots used to convey it, i.e., the signals can use any combination of available time slots.

Standard and arbitrary concatenations are network services, while virtual concatenation is a SDH/SONET end-system service, recently approved by ITU-T [1] and the committee T1 of ANSI. This service allows SDH/SONET end systems to "glue" together the VCs or SPEs of separate signals rather than having the signals carried through the network as a single unit.

SDH/SONET Transparency — The purpose of SDH/SONET is to carry its payload signals in a transparent manner. This can include some of the layers of SDH/SONET itself, i.e., the path overhead can never be touched since it actually belongs to the client. It may also be useful to transport, multiplex, and/or switch lower layers of the SDH/SONET signal transparently.

The SDH/SONET overhead is broken into three layers: the regenerator section/section layer, the multiplex section/line layer, and the path layer. All these layers are concerned with fault and performance monitoring. The regenerator section/section overhead is primarily concerned with framing, while the multiplex section/line overhead is primarily concerned with multiplexing and protection. To perform multiplexing, an SDH/SONET network element should be multiplex section/line terminating. However, not all SDH/SONET multiplexers/switches perform SDH/SONET pointer adjustments on all the STM-1s/STS-1s passing through them or, if they perform the pointer adjustments, they do not terminate the multiplex section/line overhead.

For example, a multiplexer may take four SONET STS-48 signals and multiplex them onto an STS-192 signal without performing standard line pointer adjustments on the individual STS-1s. This can be looked at as a carrier's carrier service since it may be desirable to pass SONET signals, like an STS-12 or STS-48, with some level of transparency through a network and still take advantage of TDM. Table 2 summarizes the levels of SDH/SONET transparency. In addition, a new type of flexible transparency service between the path and line layers has recently been defined [5].

Protection

SDH and SONET networks offer a variety of protection options at both the multiplex section/line level and path level. Standardized SONET/SDH line level protection techniques include linear 1+1 and linear 1:N automatic protection switching (APS) and both two-fiber and four-fiber bi-directional line-switched rings (BLSRs) [7]. At the path layer, SONET offers unidirectional path-switched ring protection. Both ring and 1:N line protection also allow for "extra traffic" to be carried over the protection line when that line is not being used, i.e., when it is not carrying traffic for a failed working line. These protection methods, which are summa-

Transparency type	Comments
Path layer (multiplex section/ line terminating)	Standard higher-order SDH/SONET path switching. Multiplex section/line overhead is terminated or modified.
Multiplex section or line level (regenerator section/ section terminating)	Preserves multiplex section/line overhead and switches the entire (line) multiplex as a whole. Regenerator section/section overhead is terminated or modified.
Regenerator section/section layer	Preserves section overhead. Does not touch any of the SDH/SONET bits.

■ Table 2. *SDH/SONET transparency types and their properties.*

Protection type	Extra traffic optionally supported	Comments
1+1 unidirectional	No	Requires no coordination between the two ends of the circuit. Dedicated protection line.
1+1 bi-directional	No	Coordination via K byte protocol. Lines must be consistently configured. Dedicated protection line.
1:1	Yes	Dedicated protection.
1:N	Yes	One Protection line shared by *N* working lines. $N \leq 14$
4F-BLSR (4-fiber bi-directional line switched ring)	Yes	Dedicated protection, with alternative ring path and span protection
2F-BLSR (2-fiber bi-directional line switched ring)	Yes	Dedicated protection, with alternative ring path
UPSR (uni-directional path switched ring)	No	Dedicated protection via alternative ring path. Typically used in access networks.

■ Table 3. *Common SDH/SONET protection mechanisms.*

rized in Table 3, are completely separate from any MPLS layer protection or restoration mechanisms.

It may be desirable to route some connections over lines that support protection of a given type, while others may be routed over unprotected lines, or as "extra data" over protection lines. Also, to assist in the configuration of these various protection methods, it can be extremely valuable to advertise the link protection attributes in the routing protocol. For example to configure a 1:*N* protection group via two nodes the lines must be "numbered the same" with respect to both ends of the connection, or else the APS (K1/K2 byte) protocol will not operate correctly.

The amount of information to disseminate concerning protection is an open issue. Table 4 presents one extreme, while the other extreme can be represented by a simple enumerated list of: extra-traffic/protection line, unprotected, shared (1:*N*)/working line, dedicated (1:1, 1+1)/working line, enhanced (ring)/working line.

Available Capacity Advertisement

For each SDH/SONET LSR interface, the LSR must maintain an internal table indicating each signal allocated in the multiplex structure. This is the most complete and accurate view of the link usage and available capacity. This information needs to be advertised in some way to all other SDH/SONET LSRs in the same domain for use in path computation. This requires a trade-off between:

- The amount of detail in the available capacity information to be reported via a link state routing protocol.
- The frequency of updates and the conditions under which this information is updated.
- The percentage of connection establishments that are unsuccessful on their first attempt.
- The extent to which network resources can be optimized.

Different levels of summarization are being considered today for the available capacity. At one extreme, all signals allocated on an interface could be advertised, while at the other extreme, a single aggregated value of the available bandwidth could be advertised.

Consider first the relatively simple structure of SONET and its most common current and planned usage. DS1s and DS3s are the signals most often carried within a SONET STS-1. Either a single DS3 occupies the STS-1 or up to 28 DS1s (4 each within the 7 VT groups) occupy the STS-1. With a reasonable VT1.5 placement algorithm within each node, it may be possible to just report on aggregate bandwidth usage in terms of the number of STS-1s currently allocated, and the number of VT-1.5s currently allocated in the STS-1 currently being filled with DS1 signals. In this manner a network optimization program could try to determine the optimal placement of DS3s and DS1s to minimize wasted bandwidth due to half-empty STS-1s on various links within the transport network.

Similarly, consider the set of super rate SONET signals (STS-*N*c). If the links between the two switches support arbitrary concatenation, then the reporting is straightforward since any of the STS-1s within an STS-M can be used for the STS-*N*c being transported over that link. However, if contiguous concatenation is also supported, then reporting becomes trickier, since there are now constraints on where the STS-1s can be placed. SDH has still more options and constraints. Therefore the best way to advertise bandwidth resource availability/usage in SDH/SONET is not clear yet.

LSP Provisioning for SDH/SONET

Traditionally, end-to-end circuit connections in SDH/SONET networks have been set up via network management systems (NMSs), which issue commands (usually under the control of a human operator) to the various network elements involved in the circuit, via an equipment vendor's element management system (EMS). Very little multi-vendor interoperability has been achieved via management systems. A common signaling protocol, such as RSVP with TE extensions [8] or CR-LDP [9], appropriately extended for circuit switching applications, could therefore help to relieve these interoperability problems.

Protection related link information	Comments
Protection type	Indicates which of the protection types delineated in Table 5.
Protection group id	Indicates which of several protection groups (linear or ring) that a node belongs to. Must be unique for all groups that a node participates in.
Working line number	Important in 1:*N* case and to differentiate between working and protection lines.
Protection line number	Used to indicate if the line is a protection line.
Extra traffic supported	Yes or no.
Layer	If this protection parameter is specific to SONET then this parameter is not needed, otherwise it would indicate the signal layer that the protection is applied.

■ Table 4. *Parameters defining protection mechanisms..*

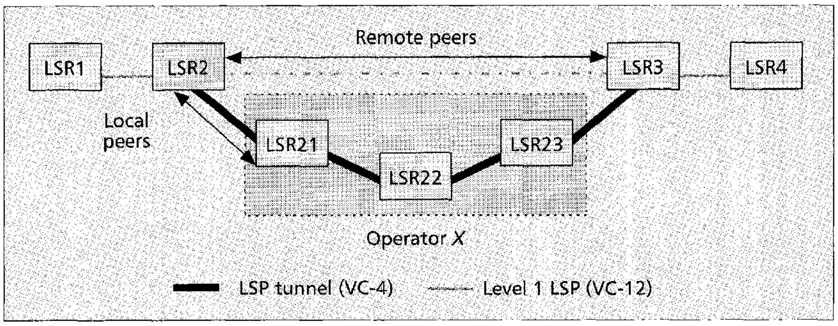

■ Figure 2. *Example of LSP tunnel within the SDH hierarchy.*

Label Structure in SDH/SONET

The signaling protocol used to establish an SDH/SONET LSP must have specific information elements in it to map a label to the particular signal type that it represents and to the position of that signal in the SDH/SONET multiplex. With a carefully chosen label structure, the label itself can be made to function as this information element. In particular, the well-defined and finite structure of the SDH/SONET multiplexing tree leads to a signal numbering scheme.

For SDH the possible leaves of the multiplex tree are VC-4, VC-3, VC-2, VC-12, or VC-11. *N* x STM-1 signals may be interleaved together to form an STM-*N* signal. Therefore, we must identify the STM-1 that is itself composed of sub-signals. To do this, the standard (K, L, M) numbering scheme, defined in G.707 [5], can be extended to:

(S, U, K, L, M) or S.U.K.L.M (in dot notation), where

S:1 → *N* indicates a specific STM-1/STS-1 inside an STM-*N*/STS-*N* multiplex;

U:0 → 4 is an index indicating a VC-4 or a particular VC-3;

K:0 → 4 is an index indicating the content of a VC-4;

L:0 → 8 is an index indicating the content of a TUG-3, VC-3; or STS-1 SPE; and

M:0 → 10 is an index indicating the content of a TUG-2 or VT group.

Each letter indicates a possible branch number starting at the parent node in the naming tree. Branches are numbered in increasing order, starting from the top of the naming tree. The numbering starts at 1, and zeros are used to indicate a non-significant field.

- S is the index of a particular STM-1/STS-1. S = 1 → *N* indicates a specific STM-1/STS-1 inside an STM-*N*/STS-*N* multiplex. For example, S = 1 indicates the first STM-1/STS-1,
- U is only significant for SDH and must be ignored for SONET. It indicates a specific VC inside a given STM-1. U = 1 indicates a single VC-4, while U = 2 → 4 indicates a specific VC-3 inside the given STM-1.
- K is only significant for SDH and must be ignored for SONET. It indicates a specific branch of a VC-4. K = 1 indicates that the VC-4 is not further subdivided and contains a C-4. K = 2 → 4 indicates a specific TUG-3 inside the VC-4. K is not significant when the STM-1 is divided into VC-3s.
- L indicates a specific branch of a TUG-3, VC-3, or STS-1 SPE. It is not significant for an unstructured VC-4. L = 1 indicates that the TUG-3/VC-3/STS-1 SPE is not further subdivided and contains a VC-3/C-3 in SDH or the equivalent in SONET. L = 2 → 8 indicates a specific TUG-2/VT group inside the corresponding higher-order signal.
- M indicates a specific branch of a TUG-2/VT group. It is not significant for an unstructured VC-4, TUG-3, VC-3, or STS-1 SPE. M=1 indicates that the TUG-2/VT group is not further subdivided and contains a VC-2/VT-6. M = 2->3 indicates a specific VT-3 inside the corresponding VT group; these values *must not* be used for SDH since there is no equivalent of VT-3 with SDH. M = 4 → 6 indicates a specific VC-12/VT-2 inside the corresponding TUG-2/VT group. M = 7 → 10 indicates a specific VC-11/VT-1.5 inside the corresponding TUG-2/VT group [5, 6]. Note that M = 0 denotes an unstructured VC-4, VC-3, or STS-1 SPE.

Example 1: S > 0, U = 1, K = 1, L = 0, M = 0, denotes the unstructured VC-4 of the *S*th STM-1.

Example 2: S > 0, U = 1, K > 1, L = 1, M = 0, denotes the unstructured VC-3 of the *K*th-1 TUG-3 of the VC-4 of the *S*th STM-1.

Example 3: S > 0, U = 0, K = 0, L = 1, M = 0, denotes the unstructured STS-1 SPE of the *S*th STS-1.

Example 4: S>0, U = 0, K = 0, L > 1, M = 9, denotes the 3rd VT-1.5 in the *L*th-1 VT group in the *S*th STS-1.

Multiplex Hierarchies and LSP Tunnels

The LSP tunneling feature in MPLS turns out to be a necessity for the control of switching networks that operate on multiple levels of a multiplex hierarchy, i.e., both TDM and WDM hierarchies. For example, in Fig. 2 we show an SDH connection consisting of seven LSRs (switches). Three of these LSRs operate strictly on higher-order signals, i.e., LSR21, LSR22, and LSR23. Two of these LSRs operate strictly on lower-order signals, i.e., LSR1 and LSR4. LSR2 and LSR3 have an unstructured VC-4 established between them that acts as an LSP tunnel for the lower-order LSP (VC-12) between LSRs 1 and 4. Note that LSR21–L23 do not need to know anything about the contents of the VC-4, and if possible, should not be sent information about it.

M	SDH	SONET
0	Unstructured VC-4/VC-3	Unstructured STS-1 SPE
1	VC-2	VT-6
2		1st VT-3
3		2nd VT-3
4	1st VC-12	1st VT-2
5	2nd VC-12	2nd VT-2
6	3rd VC-12	3rd VT-2
7	1st VC-11	1st VT-1.5
8	2nd VC-11	2nd VT-1.5
9	3rd VC-11	3rd VT-1.5
10	4th VC-11	4th VT-1.5

■ Table 5. *The encoding of the M field in the SDH/SONET multiplex entry.*

Signaling Elements

Information that is required to completely specify the signal and its characteristics for the desired connection must be transferred via the label distribution protocol. This information can be partitioned into three groups. The first group specifies the nature/type of the LSP in terms of the particular signal (or collection of signals) within the SDH/SONET multiplex that the LSP represents, and is used by all the nodes along the path of the LSP. The second group specifies the payload carried by the LSP in terms of the termination and adaptation functions required at the endpoints, and is used by the source and destination nodes of the LSP. The third group specifies certain link selection constraints, which control, at each hop, the selection of the underlying link that is used to transport this LSP.

LSP Encoding Type, Signal Type, and Connection Grouping — The nature of the SDH/SONET signal that will be switched and delivered end-to-end is specified collectively by the *LSP encoding type and signal type* fields carried in an LSP request. Another element specifying the nature of the desired LSP is the extent, if any, of connection grouping, which is specified by a combination of two fields, the *requested grouping type* and the *requested number of components*.

LSP encoding type	Signal type	
SDH	VC-11	STM-4 MS
	VC-12	STM-4 RS
	VC-2	STM-16
	TUG-2	STM-16 MS
	VC-3	STM-16 RS
	TUG-3	STM-64
	VC-4	STM-64 MS
	STM-1	STM-64 RS
	STM-1 MS	STM-256
	STM-1 RS	STM-256 MS
	STM-4	STM-256 RS

■ Table 6. *Permitted LSP encoding types and their corresponding signal types for SDH.*

LSP Encoding Type and Signal Type — The *LSP encoding type* indicates the technology of the LSP being requested, i.e., either SDH or SONET. The *signal type* field indicates the specific type of signal being requested, and is interpreted in the context of the *LSP encoding type*. Thus, the signal type provides transit switches with information required to determine the type of time-slots/labels that can support this LSP. As an example, the permitted LSP encoding types with their permitted signal types for SDH are shown in Table 6. A detailed discussion of the encoding types appears in [5].

Connection Grouping — A number of non-concatenated signals may be routed together as a group. A classical routing constraint for such a grouping is that all these signals must be within the same higher-order hierarchy in order to experience essentially the same propagation delay.

The *requested grouping type* indicates the type of grouping being requested. The values for SDH/SONET are "no grouping," "virtual concatenation," "contiguous arbitrary concatenation" (or flexible concatenation), and "contiguous standard concatenation." For contiguous standard concatenation, there must be a standard number of components (3, 12, 48, etc.), while for contiguous arbitrary concatenation, the number of components is arbitrary (2, 3, 4, ...). In all cases, the components must be routed in the same higher-order container.

The *requested number of components* (RNC) indicates the number of identical individual signal types that are requested to be grouped into an LSP, as specified in the RGT field. All components here are assumed to have identical characteristics. The field is set to zero when no grouping is requested.

Payload Type — An LSP request must also carry an identifier of the payload that will be carried by the LSP. The payload identifier is interpreted in the context of the LSP encoding type, and is used by the endpoints of the LSP. As an example, Table 7 depicts suggested payload type values for SDH/SONET.

Link Protection Type — The *link protection type* field carried in an LSP request indicates the level of protection that an LSP desires on the links at each hop along its path. In other words, the link protection is local to the interface between two adjacent nodes, and controls how the underlying link at a particular hop is protected. It is, therefore, distinct from MPLS-level protection [10], which involves protection of the actual LSP (which may be done either end-to-end or locally).

The *link protection type* may be represented as a vector of flags, where one or more protection levels may be turned on simultaneously. A value of 0 implies that this connection does not care about which, if any, link protection is used. More than one bit may be set when multiple protection types are acceptable. When multiple bits are set and multiple protection types are available, the choice of protection type is a local (policy) decision. The proposed flags are:

- Extra traffic: Indicates that links that are reserved for automatic recovery in case of a fault elsewhere in the network may be used for this LSP. Thus, the LSP can be dropped even if there is no fault on the links along which this LSP is routed.
- Unprotected: Indicates that unprotected links may be used for this LSP. Thus, the LSP does not lose service as long as the link is up, but loses service when this link goes down, since the link itself is not protected by a backup link.
- Shared: Indicates that protected (working) links whose protection resources are shared with some number, say N, of other working links may be used by this LSP. This means that if there is a fault along this particular link, the LSP will lose service on this hop only if the backup link is already in use by traffic from one of the remaining N–1 working links (due to an earlier fault on one of those links).
- Dedicated: Indicates that links with dedicated protection, e.g., 1:1 or 1+1 protection, may be used by this LSP. This means that a protection link is reserved for the working link over which this LSP is routed, so that this LSP is always protected against any fault on its working link
- Enhanced: Indicates that links that are multiply protected, such as via ring and span switching in a 4-fiber MS-SPRING/BLSR.

Thus, the link protection represents both a property of a link (which needs to be appropriately advertised in routing), as well

LSP encoding type	Payload/client type
SDH	Unknown
	Asynchronous mapping of E4
	Asynchronous mapping of DS3
	Asynchronous mapping of E3
	Bit synchronous mapping of E3
	Byte synchronous mapping of E3
	Asynchronous mapping of DS2
	Bit synchronous mapping of DS2
	Byte synchronous mapping of DS2
	Asynchronous mapping of E1
	Byte synchronous mapping of E1
	Byte synchronous mapping of 31 * DS0
	Asynchronous mapping of DS1
	Bit synchronous mapping of DS1
	Byte synchronous mapping of DS1
	ATM mapping
SONET	Unknown
	DS1 SF asynchronous
	DS1 ESF asynchronous
	DS3 M23 asynchronous
	DS3 C-bit parity asynchronous
	ATM
	POS

■ Table 7. *The payload type indicator in the context of the LSP encoding type for SDH/SONET..*

as a constraint on which links may be used for a given path (which is signaled during connection setup as specified above).

Conclusions

In this article we presented a detailed account of the issues involved in applying MPLS-based control to TDM networks by focusing on the two main areas of application of these methods to TDM networks, namely, routing and signaling. As the industry moves toward this mode of operation, however, several issues require further consideration. One such area is routing protocol scalability, which requires a careful trade-off between the granularity of information advertised in routing and the volume of routing traffic generated. Similarly, the number of underlying layers whose information is compressed into a single instance of the routing protocol is also an open question, and depends on the type of network model (overlay or peer) that a service provider intends to deploy. Finally, the bandwidth requirements of the control channel and the implementation of the control channel itself are important issues that impact the realization of the distributed automated network control as proposed in this article.

References

[1] G.707-2000, Network Node Interface for the Synchronous Digital Hierarchy (SDH), International Telecommunication Union.

[2] ANSI T1.105-1995, "Synchronous Optical Network (SONET) — Basic Description Including Multiplex Structure, Rates and Formats."

[3] G. Bernstein, J. Yates, and D. Saha, "IP-Centric Control and Management of Optical Transport Networks," *IEEE Commun. Mag.*, Oct. 2000.

[4] L. Berger *et al.*, "Generalized MPLS — Signaling Functional Description," Internet Draft: <draft-ietf-mpls-generalized-signaling-04.txt>, May 2001.

[5] E. Mannie *et al.*, "GMPLS Extensions for SONET and SDH Control," Internet Draft: <draft-ietf-ccamp-gmpls-sonet-sdh-00.txt>, May 2001.

[6] B. Davie, Y. Rekhter, and M. Kaufman, *MPLS Technology and Applications*, San Mateo, California, Jan. 2000.

[7] J. Manchester, P. Bonenfant, and C. Newton, "The Evolution of Transport Network Survivability," *IEEE Commun . Mag.*, Aug. 1999.

[8] L. Berger *et al.*, "Generalized MPLS Signaling — RSVP-TE Extensions," Internet Draft:<draft-ietf-mpls-generalized-rsvp-te-03.txt>, May 2001.

[9] L. Berger *et al.*, "Generalized MPLS Signaling — CR-LDP Extensions," Internet Draft: <draft-ietf-mpls-generalized-cr-ldp-03.txt>, May 2001.

[10]V. Sharma and F. Hellstrand, eds., "Framework for MPLS-based Recovery," Internet Draft: <draft-ietf-mpls-recovery-frmwrk-02.txt>, Mar. 2001.

Biographies

GREG BERNSTEIN (GregB@ciena.com) is senior scientist at Ciena Corporation specializing in optical network control and management planes. He was senior director for software engineering development at Lightera Networks (acquired by Ciena Corporation) where he led the development of intelligent optical switches. He was with Northern Telecom from 1994 until the founding of Lightera Networks in early 1998. At Northern Telecom, he was a senior manager responsible for new product development in the areas of ATM and Internet Protocol. He has filed or been granted seven US/European patents in the areas of packetized voice, congestion control, and optical networking. Dr. Bernstein earned his Ph.D., M.S., and B.S. in electrical engineering and computer science from the University of California at Berkeley

ERIC MANNIE is technology and standards strategy manager at Ebone (GTS), where he is responsible for investigating and analyzing new technologies. His current scope of activities is concerned with the integration of IP and optical technologies. In that context he is a co-author of several GMPLS drafts and editor of some of them. Before joining Ebone (GTS) he worked as an IP telephony consultant for DigiCable, a consultancy company. Before that he worked for eight years as a researcher at the University of Brussels (ULB) where he was also a developer and a consultant for manufacturers, operators, and national and European governmental agencies. His activities were focused on IP technologies, IP and ATM integration, MPLS, videoconferencing and IP telephony, cable modems, and network management.

VISHAL SHARMA (v.sharma@ieee.org) is principal consultant at Metanoia, Inc., where he specializes in technical consulting for telecommunications vendors and service providers. His focus is on system/network architecture, algorithm and protocol design, and system analysis. He earned a B. Tech. degree from the Indian Institute of Technology, Kanpur, and M.S. and Ph.D. degrees from the University of California, Santa Barbara. He has worked and consulted at companies such as Motorola, Tellabs, Jasmine Networks, ACT Networks, and Digital Instruments.