Multiprotocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Crossconnects

Daniel Awduche, Movaz Networks
Yakov Rekhter, Juniper Networks

ABSTRACT

This article describes an approach to the design of control planes for optical crossconnects which leverages existing control plane techniques developed for MPLS traffic engineering. The proposed approach combines recent advances in MPLS traffic engineering control plane constructs with OXC technology to provide a framework for real-time provisioning of optical channels, foster development and deployment of a new class of OXCs, and allow the use of uniform semantics for network management and operations control in hybrid networks consisting of OXCs and label switching routers. The proposed approach is particularly advantageous for OXCs intended for data-centric optical internetworking systems.

INTRODUCTION

This article describes an approach to the design of control planes for optical crossconnects (OXCs), which is based on the multiprotocol label switching (MPLS) traffic engineering control plane. The main idea is to leverage recent advances in control plane technology developed for MPLS traffic engineering [1–6]. This approach is driven by pragmatic considerations, since it exploits an existing technology base to foster rapid development and deployment of a new class of OXCs that address the optical transport needs of the Internet. This approach will assist in optical channel layer bandwidth management, dynamic provisioning of optical channels, and network survivability through enhanced protection and restoration capabilities.

As used in this article, an OXC is a path switching element in an optical transport network that establishes routed paths for optical channels by locally connecting an optical channel from an input port (fiber) to an output port (fiber) on the switch element. Additional characteristics of OXCs, as used in this article, are discussed later.

The proposed OXC control plane uses the Internal Gateway Protocol (IGP) extensions for MPLS traffic engineering (with additional enhancements) to distribute relevant optical transport network state information, including topology state information. This state information is subsequently used by a constraint-based routing system to compute paths for point-to-point optical channels. The proposed OXC control plane also uses an MPLS signaling protocol [3,4] to instantiate point-to-point optical channels.

This article does not specify the details of the extensions and domain-specific adaptations required to map the MPLS traffic engineering control plane onto the optical domain. However, we provide a high-level overview of the architectural issues involved in making such adaptations in a later section.

ADVANTAGES

The advantages of the proposed approach are numerous:

- It offers a framework for optical bandwidth management and the real-time provisioning of optical channels.
- It exploits recent advances in MPLS control plane technology and also leverages accumulated operational experience with IP distributed routing control.
- It obviates the need to reinvent a new class of control protocols for optical transport networks and allows reuse of software artifacts originally developed for the MPLS traffic engineering application. Consequently, it fosters the rapid development and deployment of a new class of OXCs.
- It facilitates the introduction of control coordination concepts between data network elements and optical network elements.
- It simplifies network administration in facilities-based service provider networks by providing uniform semantics for network management and control in both the data and optical domains.
- It paves the way for the eventual introduction of DWDM multiplexing capabilities on IP routers.

0163-6804/01/$10.00 © 2001 IEEE

The requirements of the Internet are mandating IP-centric networks to be cost effective, survivable, and scalable, and to provide control capabilities that facilitate network performance optimization.

- Lastly, it establishes a preliminary framework for the notion of an optical Internet.

BACKGROUND

The growth, performance, and survivability requirements of the Internet are mandating IP-centric networks to be cost effective, survivable, and scalable, and to provide control capabilities that facilitate network performance optimization. Some of these requirements are being addressed by the multiprotocol label switching (MPLS) traffic engineering capabilities under development by the Internet Engineering Task Force (IETF) [1–4].

The underlying optical transport network also needs to be versatile, reconfigurable, and cost effective, and to support a variety of protection and restoration capabilities due to the rapidly changing requirements of the Internet.

A result of these trends, therefore, is the evolution of optical transport networks from simple linear and ring topologies to (partial) mesh topologies.

Underscoring the importance of versatile networking capabilities in the optical domain, a number of standardization organizations have initiated work items to study the requirements and architectures for reconfigurable optical networks. For example, International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Recommendation G.872 [7] speaks of an optical transport network (OTN) as “a transport network bounded by optical channel access points.” The ITU-T G.872 OTN architecture is based on a layered structure, which includes:

- An optical channel (OCh) layer network
- An optical multiplex section layer network
- An optical transmission section layer network

The OCh layer is the most relevant to the discussions in this article. The OCh layer network supports end-to-end networking of optical channel trails between access points. The OCh layer network provides the following functions: routing, monitoring, grooming, and protection and restoration of optical channels. In this situation, programmable OXCs will be critical to the realization of the OCh layer functions, especially in mesh optical networks.

Other standards organizations and interoperability forums actively pursuing projects related to dynamically reconfigurable optical networks include the American National Standards Institute (ANSI) T1X1.5 committee, the Optical Internetworking Forum (OIF), and the IETF.

In all these cases, the successful realization of the vision of versatile optical networking depends very much on the synthesis of appropriate control plane technologies with programmable and reconfigurable OXCs.

OXCS, LSRS, OPTICAL TRAILS, AND EXPLICIT LSPS

Consider a hybrid IP-centric optical internetworking environment consisting of both label switching routers (LSRs) and OXCs, where the OXCs are programmable and support wavelength conversion/translation.

At a level of abstraction, an LSR and an OXC exhibit a number of isomorphic relations. Enumerating these relations exposes the reusable software artifacts from the MPLS traffic engineering control plane model. Architecturally, both LSRs and OXCs emphasize problem decomposition by decoupling the control plane from the data plane.

The data plane of an LSR uses the label swapping paradigm to transfer a labeled packet from an input port to an output port [8]. The data plane of an OXC uses a switching matrix to connect an OCh trail from an input port to an output port.

An LSR performs label switching by first establishing a relation between an <input port, input label> tuple and an <output port, output label> tuple. Likewise, an OXC provisions OCh trails by first establishing a relation between an <input port, input optical channel> tuple and an <output port, output optical channel> tuple. These relations are determined by the control plane of the respective network elements, and are locally instantiated on the device through a switch controller.

The functions of the control plane (for both LSRs and OXCs) include resource discovery, distributed routing control, and connection management. In particular, the control plane of the LSR is used to discover, distribute, and maintain relevant state information associated with the MPLS network, and to instantiate and maintain label switched paths (LSPs) under various MPLS traffic engineering rules and policies. An LSP is the path through one or more LSRs followed by a specific forwarding equivalence class [8]. An explicit LSP is one whose route is defined at its origination node.

The control plane of the OXC, on the other hand, is used to discover, distribute, and maintain relevant state information associated with the OTN, and to establish and maintain OCh trails under various optical traffic engineering rules and policies. An OCh trail provides a point-to-point optical connection between two access points. At each intermediate OXC along the route of an OCh trail, the OXC switch fabric connects the trail from an input port to an output port.

A distinction between OXCs and LSRs is that the former do not perform packet-level processing in the data plane, while the latter are datagram devices which may perform certain packet-level operations in the data plane. A significant conceptual difference is that with LSRs the forwarding information is carried explicitly as part of the labels appended to data packets, while with OXCs the switching information is implied from the wavelength or optical channel.

In this article we use the generic term OXC to refer to all categories of programmable and reconfigurable crossconnects for OTNs, irrespective of the technologies that underlie them.

The OXC control plane design approach described in this article is independent of the underlying OXC switch technologies. It is also independent of specific OXC implementation details. Local adaptation mechanisms can be used to tailor the control plane onto various OXC implementations with different hardware capabilities. As an example, a local adaptation function can map a channel/port input/output relation into specialized low-level instructions to actuate a rearrangement of the crossconnect switch fabric such that the required input/output relation is realized.

Explicit LSPs and Optical Channel Trails

At a conceptual level, explicit LSPs and optical channel trails exhibit certain commonalities. Essentially, they are both fundamentally unidirectional point-to-point path connection abstractions. An explicit LSP provides a parameterized packet forwarding path (traffic trunk) between an ingress LSR and an egress LSR. Correspondingly, an OCh trail provides a (possibly parameterized) OCh between two endpoints for the transport of client digital signals.

The payload carried by both LSPs and optical trails are transparent to intermediate nodes along their respective paths. Both LSPs and optical trails can be parameterized to stipulate their performance, behavioral, and survivability requirements from the network.

A constraint-based routing scheme can be used to select appropriate paths for both LSPs and optical trails. Generally, such paths may satisfy some demands and policy requirements subject to some constraints imposed by the operational environment.

Generic Requirements for the OXC Control Plane

This section contains the requirements for the OXC control plane, with emphasis on the routing components of these requirements. There are three key aspects to these requirements:

- The capability to establish OCh trails expeditiously (in seconds or even milliseconds rather than days or months).
- The capability to support traffic engineering functions — the introduction of dense wavelength-division multiplexing (DWDM) and automatically switched optical networks is unlikely to eliminate the need for traffic engineering. Instead, it will simply mandate OXCs to also support some traffic engineering capabilities.
- The capability to support various protection and restoration schemes.

Historically, the "control plane" of OTNs has been implemented via network management. This approach has the following drawbacks:

- It leads to relatively slow convergence following failure events (typical restoration times are measured in minutes, or even days and weeks, especially in systems that require explicit manual intervention). The only way to expedite service recovery in such environments is to preprovision dedicated protection channels.
- It complicates the task of interworking equipment from different manufacturers, especially at the management level (generally, a custom umbrella network management system, NMS, or operations support system, OSS, is required to integrate otherwise incompatible element management systems from different vendors).
- It precludes the use of distributed dynamic routing control capabilities in such environments.
- It complicates the task of internetwork provisioning (due to the lack of EDI between operator NMSs).

Thus, another important motivation for the approach described in this article is to improve the responsiveness of the optical transport network, and to increase the level of interoperability within and between service provider networks.

MPLS Traffic Engineering as a Generic Control Plane for OXCs

The requirements for the OXC control plane described in the previous section have already been addressed by the MPLS traffic engineering control plane, under development by the IETF [1–6].

An Overview of the MPLS Traffic Engineering Control Plane

The MPLS traffic engineering control plane is a synthesis of new concepts in IP traffic engineering (enabled by MPLS) and the conventional IP network layer control plane. The high-level requirements for traffic engineering over MPLS were articulated in [1]. It is the combination of the notions defined in [1] with the conventional IP control plane constructs that effectively establishes a framework for the MPLS traffic engineering control plane model [1; see also 2]).

The MPLS traffic engineering control plane includes the following components:

- Resource discovery.
- State information dissemination, which is used to distribute relevant information concerning the state of the network, including topology and resource availability information. In the MPLS context, this is accomplished by extending conventional IP link state routing protocols to carry additional information in their link state advertisements [5–6].
- Path selection, which is used to select an appropriate route through the MPLS network for explicit routing. It is implemented by introducing the concept of constraint-based routing, which is used to compute paths that satisfy certain specifications subject to certain constraints, including constraints imposed by the operational environment [1].
- Path management, which includes label distribution, path placement, path maintenance, and path revocation. These are used to establish, maintain, and tear down LSPs in the MPLS context. The label distribution, path placement, and path revocation functions are implemented through a signaling protocol, such as the Resource Reservation Protocol (RSVP) extensions [3] or constraint routed label distribution protocol (CR-LDP) [4].

These components of the MPLS traffic engi-

A constraint-based routing scheme can be used to select appropriate paths for both LSPs and optical trails. Generally such paths may satisfy some demands and policy requirements subject to some constraints imposed by the operational environment.

The MPLS traffic engineering control plane (with some minor extensions) would be very suitable as the control plane for OXCs. An OXC that uses the MPLS traffic engineering control plane would effectively become an IP addressable device.

neering control plane are separable and independent of each other. This is a very attractive feature because it allows an MPLS control plane to be implemented using a composition or synthesis of best-of breed-modules.

In [1] several new MPLS control plane capabilities were proposed that allow various traffic engineering policies to be actualized in MPLS networks. Many of these capabilities are also relevant and applicable to OTNs with reconfigurable OXCs.

We summarize some of these capabilities below, focusing on the set of attributes associated with traffic trunks. A traffic trunk is an aggregation of traffic belonging to the same class that is forwarded through a common path. In general, the traffic trunk concept is a technology-independent abstraction. In [1] it was used within the context of MPLS and allowed certain attributes of the traffic transported through LSPs to be parameterized. The traffic trunk concept can also be extended, in an obvious way, to the OTN.

As stipulated in [1], the attributes that can be associated with traffic-trunks include:

- Traffic parameters which indicate the bandwidth requirements of the traffic trunk
- Adaptivity attributes, which specify the sensitivity of the traffic trunk to changes in the state of the network and, in particular, indicates whether the traffic trunk can be rerouted when “better” paths become available.
- Priority attributes, which impose a partial order on the set of traffic trunks, and allow path selection and path placement operations to be prioritized
- Preemption attributes which indicate whether a traffic trunk can preempt an existing traffic trunk in its path
- Resilience attributes, which stipulate the survivability requirements of the traffic trunk and, in particular, the response of the system to faults that impact the path of the traffic trunk
- Resource class affinity attributes, which further restrict route selection to specific subsets of resources and, in particular, allow generalized inclusion and exclusion policies to be implemented

It should be clear that a subset of these capabilities can be mapped onto an OTN by substituting the term *optical channel trail* for the term traffic trunk.

The MPLS control plane also supports the notion of an abstract node, which is essentially a set of nodes (e.g., a subnet, an autonomous system) whose internal topology is opaque to the origination node of an explicit LSP. So, in the most general manner, the route of an explicit LSP (or traffic trunk) can be specified as a sequence of single hops and/or abstract nodes.

The MPLS control plane is very general and also oblivious to the specifics of the data plane technology. In this regard, the MPLS control plane can be used in conjunction with a data plane that does not necessarily process IP packet headers and does not know about IP packet boundaries. For an existence proof, note that the MPLS control plane has been implemented on IP-LSRs (LSRs implemented on routers), ATM-LSRs (LSRs implemented on asynchronous transfer mode switches), and frame relay-LSRs (LSRs implemented on frame relay switches).

The MPLS control plane may also be implemented on OXCs, as discussed in this article.

SYNTHESIZING THE MPLS TRAFFIC ENGINEERING CONTROL PLANE WITH OXCS

Given that both OXCs and LSRs require control planes, one option would be to have two separate, independent, and incompatible control planes, one for OXCs and another for LSRs. To understand the drawbacks of this approach, especially in IP-centric optical internetworking systems, one need look no further than the experience with IP over ATM, where IP has its own control plane — Border Gateway Protocol (BGP), Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF) — and ATM its own control plane (private network–network interface, PNNI). For some of the drawbacks see [1, 2].

Given that the control planes for both OXCs and LSRs have relatively similar requirements, an alternative approach is to develop a coherent control plane technology that can be used for LSRs and OXCs. Such a uniform control plane will eliminate the administrative complexity of managing hybrid optical internetworking systems with separate, dissimilar control and operational semantics. Specializations may be introduced in the control plane, as necessary, to account for inherent peculiarities of the underlying technologies and networking contexts.

All of the above observations suggest, therefore, that the MPLS traffic engineering control plane (with some minor extensions) would be very suitable as the control plane for OXCs. An OXC that uses the MPLS traffic engineering control plane would effectively become an IP addressable device. Thus, this proposition also solves the problem of addressing for OXCs. The distribution of topology state information, establishment of OCh trails, OTN traffic engineering functions, and protection and restoration capabilities would be facilitated by the MPLS traffic engineering control plane.

An out-of-band IP communications system can be used to carry and distribute control traffic between the control planes of OXCs, perhaps through dedicated supervisory channels (e.g., using dedicated wavelengths or channels, or an independent out-of-band IP network). In this environment, Simple Network Management Protocol (SNMP) or some other network management technology could be used for element management. From the perspective of control semantics, an OXC with an MPLS traffic engineering control plane would resemble an LSR.

The physical fiber between a pair of OXCs would represent a single link in the OTN network topology. Individual wavelengths or channels would be analogous to labels. If there are multiple fibers between a pair of OXCs, then, as an option, these multiple fibers could be logically grouped together through a process called bundling and represented as a single link in the OTN network topology.

IS-IS or OSPF, with extensions for traffic engineering [5, 6] and possibly additional optical network specific extensions would be used to distribute information about the OTN topology as well as the information about available bandwidth and available channels per fiber. This information will then be used to compute explicit routes for OCh trails. An MPLS signaling protocol, such as RSVP extensions [3], will be used to instantiate the optical channel trails. Using the RSVP extensions, for example, the wavelength information or OCh information (as the case may be) will be carried in the LABEL object, which will be used to control and reconfigure the OXCs.

To bootstrap the system, OXCs must be able to exchange control information. One way to support this is to preconfigure a dedicated control wavelength between each pair of adjacent OXCs, or between an OXC and a router, and to use this wavelength as a supervisory channel for exchange of control traffic. Another possibility is to construct a dedicated out-of-band IP network for the distribution of control traffic.

Even though an OXC equipped with an MPLS traffic engineering control plane would (from a control perspective) resemble an LSR, there are some important distinctions and limitations. One distinction is that there are no analogs of label merging in the optical domain. This implies that an OXC cannot merge several wavelengths into one wavelength. Another distinction is that an OXC cannot perform the equivalent of label push and pop operations in the optical domain. This is because the analog of a label in the OXC is a wavelength or an OCh, and the concept of pushing and popping wavelengths is infeasible with contemporary commercial optical technologies.

In the proposed approach, an OXC will maintain a wavelength forwarding information base (WFIB) per interface (or per fiber). This is because lambdas and/or channels (labels) are specific to a particular interface (fiber), and the same lambda and/or channel (label) could be used concurrently on multiple interfaces (fibers).

The MPLS traffic engineering control plane is already being implemented on data plane technologies that exhibit some of the aforementioned distinctions. For example, an ATM-LSR supports only a subset of the MPLS functionality. In particular, most ATM-LSRs are incapable of merging LSPs, and may not be able to perform label push/pop operations as well. Also, similar to the approach proposed here for OXCs, ATM-LSRs have a per-interface label forwarding information base (LFIB).

Yet another important distinction concerns the granularity of resource allocation. An MPLS LSR which operates in the electrical domain can potentially support an arbitrary number of LSPs with arbitrary bandwidth reservation granularities (bounded by the maximum reservable bandwidth per interface, label space, and amount of required control overhead). In contrast, an OXC can only support a relatively small number of OCh trails (this may change as the technology evolves), each of which will have discrete bandwidth granularities (e.g.,OC-12, OC-48, OC-192, and OC-768). A special degenerate case occurs when the control plane is used to establish OCh trails that all have a fixed bandwidth (e.g., OC-48).

If the bandwidth associated with an LSP is small relative to the capacity of an OCh trail, very inefficient utilization of network resources could result if only one LSP is mapped onto a given OCh trail. To improve utilization of resources, it is necessary to be able to map several low-bandwidth LSPs onto a relatively high-capacity OCh trail. For this purpose, a generalized notion of *nested LSPs* may be used. Since an OXC cannot perform label push/pop operations, the start/end of a nested LSP has to be on a router (since nesting requires label push/pop). In this nesting situation, it is the wavelength of the *container* OCh trail itself that effectively constitutes the outermost label.

The transparency and multiprotocol properties of the MPLS control plane approach would allow an OXC to route OCh trails carrying various types of digital payloads (including IP, ATM, sycnchronous optical network, etc.) in a coherent and uniform way.

Since an OXC cannot perform label push/pop operations, the start/end of a nested LSP has to be on a router. In this nesting situation, it is the wavelength of the "container" optical channel trail itself that effectively constitutes the outermost label.

ARCHITECTURAL CONSIDERATIONS FOR DEPLOYMENT IN OPERATIONAL NETWORKS

This section provides a high-level overview of the architectural considerations for deployment of the proposed control plane in operational networks consisting of LSRs and OXCs. These architectural issues have implications on the degree of control isolation, control coupling, and control cohesion between LSRs and OXCs.

Essentially, there are two extremal architectural options for deployment of the proposed control plane in an operational context consisting of LSRs and OXCs.

- Overlay option: One option is to use different instances of the control plane in the OTN (OXC) and IP (LSR) domains. In this situation, each instance of the control plane will operate independent of the other. Interworking (including control coordination) between the two domains can be established through static configuration or through some other procedures that are outside the scope of this article. This partitioned and explicitly decoupled deployment option allows maximal control isolation between the OTN and IP domains. This scheme is conceptually similar to the model in use today, whereby the OTN simply provides point-to-point channels between IP network elements with very minimal control interaction between the two domains.
- Peer option: Another option is to use a single instance of the control plane that subsumes and spans LSRs and OXCs. In such an environment an LSP could traverse an intermix of routers and OXCs, or span just routers or just OXCs.

One of the advantages of the control plane design approach described in this article is that it potentially allows network administrators the leeway to make these deployment architectural decisions based on their specific objectives, network contexts, and service models.

Other architectural options are also possible which allow various degrees of control isolation and control integration between the OXCs and LSRs. To improve scalability the control plane may use a routing hierarchy (e.g., routing areas). A hierarchy may be applied in either of the situations mentioned above. Furthermore, in the overlay option with different control plane instances for OXCs and LSRs, a hierarchy could be enabled for each control plane instance independent of the other.

In the deployment option with a single instance of the control plane, each routing area may maintain a link state database that contains:

- Physical LSPs (fiber links)
- Optical LSPs (OCh trails)
- Logical LSPs (conventional LSPs)

As a general rule, all of these path-oriented connection entities could simply be considered LSPs with different characteristics. The origination LSR (the head-end) of each LSP entity may locally decide whether to advertise the LSP (with appropriate attributes) so that other LSRs could use it as a link for subsequent path computations.

There are significant trade-offs to the above deployment options, including aspects related to scalability and fault isolation. These trade-offs are outside the scope of this article.

One of the advantages of the control plane design approach described in this article is that it potentially allows network administrators the leeway to make these deployment architectural decisions based on their specific objectives, network contexts, and service models.

SUMMARY

This article outlines how the MPLS traffic engineering control plane could be adapted and reused as the control plane for OXCs. Such a control plane would be used to distribute optical transport network topology state information and to set up optical channel trails. Such a control plane would support various traffic engineering functions in the optical domain, and enable a variety of protection and restoration capabilities. Furthermore, such a control plane technology would expedite the development and deployment of a new class of data-centric OXCs. Additionally, the proposed approach would simplify integration of OXCs and LSRs. Finally, the proposed approach would provide coherent semantics for network management and operations control in hybrid optical internetworking systems consisting of LSRs and OXCs.

ACKNOWLEDGMENTS

We would like to acknowledge John Drake for his contributions to this article.

REFERENCES

[1] D. Awduche *et al.*, "Requirements for Traffic Engineering over MPLS," RFC 2702, Sept. 1999.
[2] D. Awduche, "MPLS and Traffic Engineering in IP Networks," *IEEE Commun. Mag.*, Dec., 1999.
[3] D. Awduche *et al.*, "Extensions to RSVP for LSP Tunnels," Internet draft, 1999, work in progress.
[4] B. Jamoussi *et al.*, "Constraint-Based LSP Setup Using LDP," Internet draft, 1999, work in progress.
[5] H. Smit and T. Li, "IS-IS Extensions for Traffic Engineering," Internet draft, 1999, work in progress.
[6] D. Katz and D. Yeung, "Traffic Engineering Extensions to OSPF," Internet draft, 1999, work in progress.
[7] ITU-T G.872, "Architecture for Optical Transport Networks," 1999.
[8] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," Internet draft, 1999, work in progress.

BIOGRAPHIES

DANIEL AWDUCHE (awduche@movaz.com) is vice president of Network Architecture at Movaz Networks, an optical networking equipment company. Prior to Movaz, he served as Distinguished Technical Member and acting director of Global Network Architecture at UUNET, a WorldCom company. He is active in the IETF where he contributed to the development of MPLS and helped pioneer the concept of multiprotocol lambda switching (MPLambdaS). A respected authority on traffic engineering, he has been very prominent in articulating some of the modern concepts of Internet traffic engineering.

YAKOV REKHTER (yakov@juniper.net) is a distinguished engineer at Juniper Networks. Prior to Juniper he was a Cisco Fellow at Cisco Systems. He was one of the leading architects and a major software developer of the NSFNET Backbone Phase II. He co-designed the Border Gateway Protocol (BGP). He was also one of the lead designers of Tag Switching, BGP/MPLS VPNs, and MPLS traffic engineering. Among his most recent activities is the work on MPLambdaS. He is the author or co-author of over 40 IETF RFCs, and numerous papers and articles on TCP/IP and the Internet.