ELSEVIER

Computer Networks 40 (2002) 111–129

COMPUTER NETWORKS

www.elsevier.com/locate/comnet

Invited Paper

Internet traffic engineering using multi-protocol label switching (MPLS)

Daniel O. Awduche [a], Bijan Jabbari [b,*]

[a] *Movaz Networks, One Technology Parkway South, Norcross, GA 30092, USA*
[b] *ECE Department, George Mason University, M/S 1G5, Fairfax, VA 22030-4444, USA*

Abstract

With the rising popularity of the Internet there have arisen corresponding requirements for network reliability, efficiency, and service quality. Internet service providers are responding to these developments by critically examining every aspect of their operational environment, looking for opportunities to scale their networks and optimize performance. In this context, traffic engineering has emerged as a major consideration in the design and operation of large public Internet backbone networks. However, the classical Internet interior gateway routing protocols hinder the practical realization of sophisticated traffic engineering policies in legacy IP networks. The advent of multi-protocol label switching (MPLS) offers the prospect to address some of the shortcomings associated with traffic engineering in IP networks. This paper discusses the techniques and practices of traffic engineering in contemporary IP networks, emphasizing the role of MPLS in performance optimization of the public Internet. We also examine the impact of generalized MPLS (GMPLS) on traffic engineering in IP-over-optical networks as the underlying technologies continue to mature.

Keywords: MPLS; GMPLS; MPλS; Traffic engineering; Routing; Internet performance optimization

1. Introduction

Multi-protocol label switching (MPLS) has come a long way from its early beginnings. Initially conceived as a means to expedite packet forwarding in legacy routers with software-based forwarding engines, MPLS has resulted in fundamental advancements in IP control plane technology, Internet traffic engineering, virtual private networks (VPNs), connection management in optical networks, and IP-over-optical inter-networking architectures. When MPLS is combined with differentiated services and constraint-based routing, various types of QoS capabilities can be implemented in IP networks. These QoS capabilities provide a pathway to transition the Internet from a best effort environment to a true multi-service infrastructure.

The main architectural concept underlying MPLS is the clear separation of the control plane from the data plane in network switching elements. The data plane consists of forwarding

---

[*] Corresponding author. Tel.: +1-703-993-1618; fax: +1-703-993-1601.

E-mail addresses: awduche@movaz.com (D.O. Awduche), bjabbari@gmu.edu (B. Jabbari).

1389-1286/02/$ - see front matter © 2002 Published by Elsevier Science B.V.
PII: S1389-1286(02)00269-4

components that perform simple label switching operations according to the classical label swapping paradigm. The control plane, on the other hand, is concerned with network level coordination functions, such as routing and signaling, to facilitate the movement of traffic across the entire network.

As the Internet continues to grow rapidly, Internet service providers are confronted with a number of fundamental challenges. One challenge concerns the need to scale and expand the network infrastructure to accommodate an increasing demand for bandwidth. Another challenge concerns the need to manage the installed capacity effectively and efficiently to enhance the end-user perception of network service quality, while minimizing costs. Related to the above considerations is yet another important challenge imposed by the need for enhanced survivability. Traffic engineering will play a pivotal role in the effective and efficient management of the installed capacity in public IP networks. It will also perform important functions in the introduction of sophisticated resilience capabilities (fault recovery and restoration) into the infrastructure.

One of the first major applications of MPLS in operational IP networks was traffic engineering [1,2,4,12,14]. The impetus for MPLS-based traffic engineering originated not from the research community, but instead from major Internet service providers who grapple with the growth and performance challenges of the rapidly evolving Internet. The requirements for traffic engineering over MPLS were articulated by Awduche et al. in [1]. The essence of Internet traffic engineering (TE) is the performance optimization of IP networks. In this context, the performance measures of interest to end users of the network infrastructure emphasize QoS objectives such as low delay, low delay variation, high throughput, low packet loss, and predictable service. On the other hand, the performance measures of interest to service providers also emphasize minimizing costs through efficient utilization of network assets to enhance business outcomes in the commercial and highly competitive Internet environment. Thus, the essential goals of Internet traffic engineering are to optimize traffic oriented performance characteristics while simultaneously minimizing network costs through efficient utilization of network resources.

The applicability of MPLS to Internet traffic engineering arose from the limitations of conventional shortest path interior gateway routing protocols, such as IS–IS and OSPF, which employ simple, distributed, and unconstrained shortest path algorithms to establish forwarding paths through the network. The main issue with these protocols is that they do not take capacity constraints and traffic characteristics into account in making routing decisions. The outcome, therefore, is that some segments of the network become congested while other segments along alternative routes remain under-utilized. The way in which MPLS addresses the IP traffic engineering problem will be discussed in detail in Section 4 of this paper.

Another important application of MPLS presently under consideration concerns QoS management in IP networks. MPLS by itself does not provide QoS capabilities. However, when MPLS is combined with constraint-based routing and differentiated services, together they allow sophisticated QoS capabilities to be introduced in the Internet [1,13,18].

Still another important application of MPLS relates to VPNs. Essentially, a VPN is a network inter-connecting multiple sites belonging to one organization (intranet) or belonging to a group of related organizations (extranet), which is provisioned over a shared public network. Typically, VPNs employ tunneling techniques to isolate traffic belonging to a VPN from other traffic within the network. Tunneling also allows a VPN to use private addressing schemes and to carry different types of traffic. For this application, MPLS can be viewed as a tunneling technology that supports the implementation of VPN services.

Lastly, the MPLS traffic engineering control plane has been extended and generalized to serve as the control plane for different types of switched transport networks, ranging from packet-switched networks and time division multiplexing capable interface (TDM) technologies, to automatically switched optical transport networks [3,20,21]. This generic MPLS-based control plane technology is presently being standardized by the IETF within the concept of generalized MPLS (GMPLS).

The advent of MPLS along with the standardization issues surrounding it has provoked a significant organizational realignment within the IETF, culminating in the formation of a new IETF Directorate termed the "SUB-IP area".

This paper provides a state of the art review of MPLS technology, focusing primarily on the traffic engineering application. The remainder of this paper is structured as follows: Section 2 describes the MPLS components. Section 3 provides a discussion of traffic engineering and its process model. Section 4 presents an overview of the general considerations surrounding traffic engineering in IP networks. Operational considerations for MPLS traffic engineering are also discussed in this section. This is followed in Section 5 by a short description of the key concepts of differentiated services (Diffserv)-aware MPLS traffic engineering. Section 6 considers briefly the analytical models for traffic engineering that have been proposed in the literature. Section 7 is devoted to GMPLS, which is an adaptation and generalization of the MPLS traffic engineering control plane to support different types of transport networks. Section 8 discusses some of the future directions in MPLS traffic, covering concepts such as policy-based MPLS network management, service level agreement management, and customer network management. Finally, Section 9 contains our concluding remarks.

2. MPLS fundamentals

The basic premise behind MPLS is quite simple: The main idea is to attach a short fixed-length label to packets at the ingress to an MPLS domain [17]. Throughout the interior nodes of the MPLS domain, the labels rather than the original packet headers, are used to make forwarding decisions. The assignment of labels to packets is based on the concept of forwarding equivalence class (FEC) [11]. According to this concept, packets belonging to the same FEC are assigned the same label at an ingress node to an MPLS domain and generally traverse through the same path (or multi-path) across the MPLS network.

The definition of forwarding equivalence class can be quite general. A FEC may consist of packets entering a network through the same ingress node and exiting the network through the same egress node. A FEC may also consist of packets belonging to the same service class, entering and exiting the network through the same ingress and egress nodes, and requiring similar QoS or packet treatment across the MPLS domain. A FEC may even consist of packets belonging to the same flow. Generally, the association of FECs to packets can be based on information contained in the packets, or on information extraneous to the packet (such as the ingress port through which the packet entered the node), or a combination of both. In essence, MPLS enables the allocation and binding of labels to various granularities of flows in a packet-switched network.

The path traversed by a "forwarding equivalence class" is called a label switched path (LSP). A signaling protocol is used to establish and teardown LSPs. The signaling protocol is involved in label allocation, label distribution, and label binding. An explicit LSP is one whose route is determined at its originating node. Within the context of explicit routing for traffic engineering and quality of service applications, the signaling protocol may also convey various types of attributes associated with explicit LSPs.

One of the characteristics that distinguishes MPLS from earlier label swapping technologies (such as frame relay and ATM) is the concept of 'label stacking,' which is an ordered set of labels affixed to a packet. Label stacking allows multiple labels to be assigned to the same packets at one or more nodes in the network, in a hierarchical arrangement. Routers which can forward both MPLS labeled packets and conventional IP packets are called label switching routers (LSRs).

From a topological perspective, the LSRs at the edge of an MPLS network that assign labels to packets are generally referred to as label edge routers (LERs). Fig. 1, depicts an MPLS network containing LERs at the boundary to the network and conventional LSRs within the core. It should be noted that LERs are simply roles played by LSRs with respect to FEC assignment at an

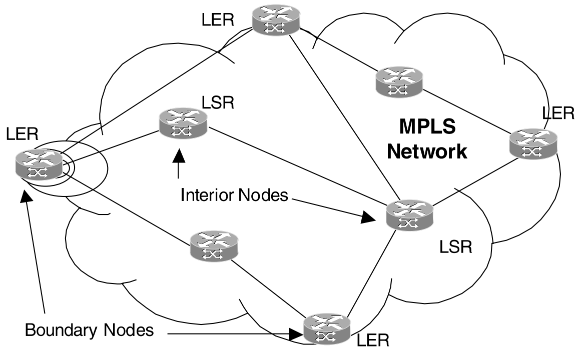

Fig. 1. Interior and boundary nodes in an MPLS network.

ingress node (or removal of labels at an egress node) in an MPLS network.

MPLS consists of a forwarding (or transport) plane and a control plane. The two are decoupled and independent of one another. Fig. 2, depicts a conceptual view of the MPLS control plane and forwarding plane. The MPLS control plane is a collection of protocols that collectively establish network level functionality in MPLS networks. The protocols themselves are implemented as software processes that communicate with each other across node boundaries using message passing. The protocol specifications detail the message formats, syntax, semantics, and transaction sequence for the message exchange. One of the main functions performed by the MPLS control plane is to facilitate the establishment of label switched paths in MPLS networks. The establishment of LSPs may be subject to various types of preferences and constraints. This means the control plane needs to distribute and manage network topology and resource availability information using a routing protocol, and perform signaling functions to establish and tear-down LSPs.

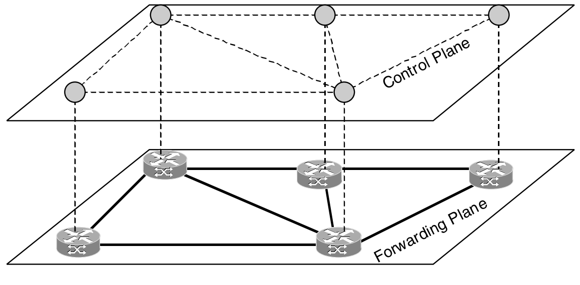

Fig. 2. Conceptual view of MPLS control plane and forwarding plane.

In practice, signaling is one of the most fundamental aspects of the MPLS control plane. Indeed, much of the work of the IETF MPLS working group has centered around developing signaling protocols for label distribution and LSP management.

For MPLS traffic engineering applications, the control plane consists of the legacy IP routing and signaling protocols along with the extensions that have been incorporated into them to support the new requirements imposed by traffic engineering (ISIS-TE, OSPF-TE, RSVP-TE, CR-LDP, BGP). The two main subsystems of the MPLS-TE control plane are (1) the signaling protocol with all pertinent extensions, e.g., RSVP-TE or CR-LDP [5,15]; and (2) the routing protocol with applicable extensions, e.g., OSPF-TE. As an example, the signaling protocol RSVP-TE consists of extensions to the IETF's RSVP protocol to support the establishment of parameterized explicit label switched paths in MPLS networks. We will have more to say about the MPLS traffic engineering control plane in Section 4.

The MPLS forwarding plane consists of the datapath within a network element through which user traffic traverses. The forwarding plane performs label swapping operations using lookup tables and miscellaneous packet treatment functions such as scheduling, queue management, rate shaping, policing, and others. The forwarding plane is generally implemented in hardware to support high speed operations. Fig. 3, depicts a functional view of the MPLS control and forwarding planes.

3. Traffic engineering process model

Internet traffic engineering deals with the performance optimization of operational IP networks. Optimization in this context refers to the transport of IP packets in the most efficient, reliable, and expeditious manner possible through a given network [1,4]. Traffic engineering can also be applied for both congestion avoidance and congestion re-

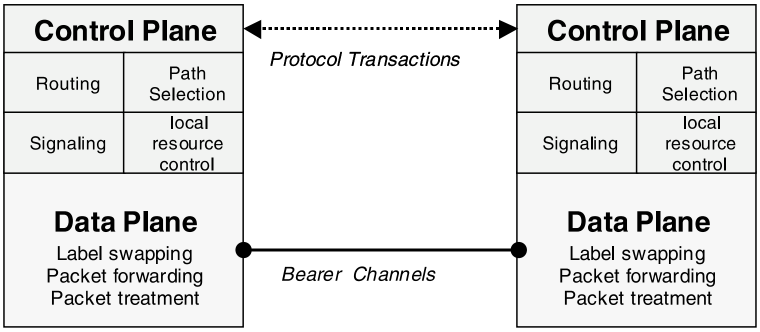

Fig. 3. Functional view of control and forwarding planes.

covery in backbone IP networks, especially congestion problems caused by poor resource allocation. For example, in a particular network scenario, some resources might be over-utilized and congested, while resources along alternative viable paths remain under-utilized. This type of congestion problems, caused by inefficient resource allocation, is one of the major issues that traffic engineering aims to obviate [1]. To accomplish these goals, Internet traffic engineering applies technology and scientific principles to the measurement, modeling, characterization, and control of Internet traffic. Roughly speaking, it is often asserted by practitioners in the field that *traffic engineering* in large scale IP networks essentially boils down to the ability to place traffic where the capacity exists to accommodate it; whereas *network engineering*, on the other hand, boils down to the ability to install capacity where the traffic exists.

Although the scope of TE is broad—encompassing the application of technology and scientific principles to the measurement, modeling, characterization, and control of Internet traffic—the paramount aspect of traffic engineering in service provider networks is the transport of IP traffic through a given network in the most efficient, economical, reliable, and expeditious manner possible [4].

In [2], the context for Internet traffic engineering was described. This includes a network context, a problem context, a solution context, and an operational and implementation context. The network context relates to network structure, network policies, network characteristics, network constraints, network quality attributes, and network optimization criteria [2]. The problem context concerns identification, abstraction, representation, formulation, and specification of the desirable features of acceptable solutions [2]. The solution context involves analysis, evaluation of alternatives, prescription, and resolution. Finally, the operational and implementation context involves planning, organization, and execution [2].

The main difficulty with Internet traffic engineering has been the limited capabilities of IP technologies concerning traffic control, resource control, and measurement. The simplicity and distributed nature of Internet link-state routing protocols has been viewed as one of the advantages of IP networks. Such protocols are also called distributed database protocols because each node within a routing area maintains an identical copy of the area link-state database, which is updated and synchronized periodically using a reliable flooding mechanism that disseminates link-state advertisements. Route computation is based on shortest path algorithms (typically Dijkstra's) using administratively specified link metrics. Each node performs route computation independently. The basic problem that arise with these protocols is that resource availability and traffic characteristics are not taken into consideration in making routing decisions, which can result in congestion in some network segments, even in networks with a preponderance of under-utilized links. This phenomenon is sometimes referred to as super-aggregation of traffic, especially when the shortest paths of multiple traffic streams converge on specific router interfaces resulting in congestion problems.

In the absence of effective control over traffic routing, any aspiration towards network performance optimization and QoS provisioning is likely to remain elusive. The reason for this is that routing has a substantial influence on many of the key performance measures of operational networks, such as congestion, throughput, delay, and resource utilization. Therefore, the Internet will remain a best effort environment without the introduction of more sophisticated routing control capabilities, other than simple unconstrained shortest path algorithms. The ability to enable constraint-based routing in IP networks is one of the achievements of MPLS that makes it particularly useful for traffic engineering.

Another issue with conventional IP routers is that it is not feasible to estimate the network traffic matrix from interface statistics on the routers. Still another issue relating to measurement is that when congestion occurs in the core of the network, it is very difficult to determine which source–destination pairs contribute to the congestion and the proportion of traffic contributed by each pair. Recently, MPLS and Diffserv have emerged as two complementary technologies that can facilitate the traffic engineering function in IP networks. Before we delve into the applications of MPLS to traffic engineering in IP networks, it is worthwhile to review the traffic engineering process model (see also [4]).

We illustrate the basic concepts of Internet traffic engineering by describing the traffic engineering process model. The process model represents the different phases in the lifecycle of traffic engineering in an operational context. The process model is iterative and cyclic. There are four main phases to this process model: (1) policy formulation phase, (2) data acquisition phase, (3) analysis and characterization phase, and (4) performance optimization phase. The interaction between the phases is characterized by major and minor workflow cycles as shown in Fig. 4.

3.1. Policy formulation phase

Effective traffic engineering requires first of all the formulation of an appropriate control policy. This activity is performed during the policy formulation phase in the traffic engineering process model. Generally, the policies will depend on the network context, the business model, the cost structure, prevailing policies, and the optimization criteria. The policy formulation phase may apply a conceptual business model, a performance model, and a revenue or utility model to aid in the creation of appropriate policies. The policies formulated during this phase provide guiding principles governing the management, control, and operation of the network.

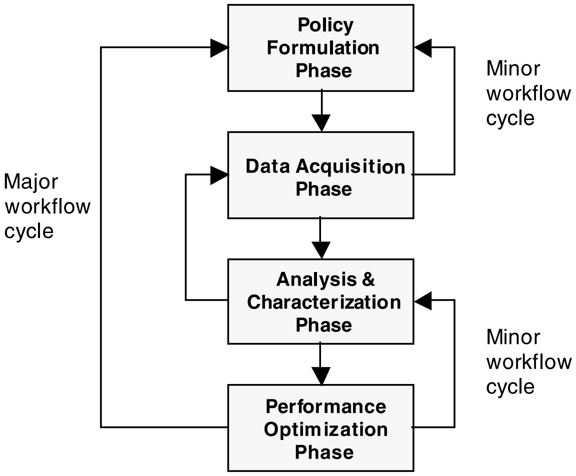

Fig. 4. Traffic engineering process model.

Within the context of MPLS-based traffic engineering, the policy formulation phase involves several considerations, such as deciding whether to conduct strategic or tactical traffic engineering in the network, determining the measurement philosophy and methodology for the network, and determining the update policy for LSPs in the network. Strategic traffic engineering in the MPLS context involves careful planning of the LSP virtual topology and adherence to a systematic methodology for reconfiguration (including modification) of existing LSPs. It may also involve careful consideration of forecasted traffic patterns in the future to come up with an evolutionary plan that accounts for existing and future traffic demands. Strategic traffic engineering also involves careful attention to how, where, and when new LSPs are activated to address performance issues in the network. Tactical MPLS traffic engineering is a more ad hoc approach to optimizing network performance by establishing and managing explicit LSPs purposely to address very specific network performance problems. For example, new LSPs may be created (in an ad hoc manner) to deliberately divert traffic away from congested network resources onto under-utilized alternatives. Tactical traffic engineering is sometimes referred to as the "hybrid approach" because it involves using LSPs to control traffic paths in some segments of the network and using interior gateway routing protocol metrics to control paths in other segments of the network.

3.2. Data acquisition phase

During the data acquisition phase, empirical statistics are collected from the operational network through a measurement system. These statistics should be carefully chosen to capture relevant operational characteristics, such as traffic patterns, link utilization, traffic trends, and packet drop statistics.

Sometimes, it may not be feasible to obtain empirical statistics from an operational network for many reasons. The network may not exist, for example, during network planning and network design. The measurement system may not cover the whole network. Finally, the empirical measures of interest may not be directly observable. In such instances, mathematical models may be used when all required empirical data is unavailable. Mathematical models may also be used to supplement and complement empirical statistics. The data acquisition phase is essentially the feedback component of the traffic engineering process model.

In the MPLS context, the data acquisition phase may involve monitoring, measuring, and storing various performance and fault statistics associated with LSPs and the underlying network infrastructure. Data acquisition may entail measurement of traffic performance characteristics, measurement of resource utilization, measurement of routes traversed by specific traffic streams, and measurement of traffic statistics between specified nodes in the network.

3.3. Analysis and characterization phase

The analysis and characterization phase involves analysis and characterization of the traffic workload derived from the measurement phase. This is essentially the performance evaluation aspects of traffic engineering. Performance evaluation can be qualitative or quantitative, and may be proactive or reactive [2]. In general, various techniques and methodologies can be applied during this phase, such as simulation and analytical techniques based on mathematical models.

One of the objectives of the analysis and characterization phase is to understand the underlying phenomenon occurring within the network, and particularly to understand the root cause of anomalous network behavior. Another objective of this phase is to determine the performance of the network under various scenarios using different types of performance measures. Structural bottlenecks such as hot-spots, and various time-series characteristics of the network such as peak rates, busy hour, and seasonality may also be identified during this phase.

Traffic engineering in large networks is a complex endeavor. Therefore, there is a need for offline analysis and simulation tools to support the traffic engineering function, especially the analysis, characterization, and optimization aspects of this activity. The tools may include various mathematical models and optimization techniques, resource models, traffic models, queuing models, time-series models, routing analysis models, models for resource dimensioning, and many others. Analysis tools are particularly useful in MPLS networks because of the potential operational complexity that may be associated with managing a large number of LSPs.

3.4. Performance optimization phase

The performance optimization phase is the fourth phase in the traffic engineering process model and involves applying an appropriate decision process to select the best course of action to enhance performance of the network. Optimization in the traffic engineering sense is not a one-time process, but rather involves a continual and iterative process of network performance improvement.

Optimization may involve regulating the inflow of traffic into the network, controlling the mapping of traffic onto available network resources using routing control capabilities, expanding the network topology by adding additional links, increasing the capacity of existing links, or controlling local packet treatment policies (queueing, scheduling, dropping policy, etc.) at individual network elements. The optimization phase may resort to the use of traffic control mechanisms to accomplish the objective. Traffic controls are those mechanisms that are used to regulate the flow of traffic through a network and to guide the routing of traffic through the network. Performance optimization may also initiate a planning process to increase the capacity of the network.

In the MPLS context, performance optimization may involve: (1) creating new LSPs and carefully controlling their routes using an appropriate path selection mechanism; (2) rerouting existing LSPs to alleviate congestion problems or to circumvent a network anomaly, or to establish a more balanced traffic distribution; (3) deactivating and tearing down an existing LSP; (4) modifying the parameters of existing LSPs to modulate their behavioral characteristics; (5) modifying the attributes associated with network resources that influence the placement of LSPs over them; (6) adding additional capacity to the network; (7) creating multiple LSPs with common endpoints and partitioning and allocating the traffic between the endpoints across the parallel LSPs; (8) starting a network planning process to expand the network topology and capacity; and (9) modifying nodal traffic management parameters. This phase may also involve additional activities in the MPLS context such as modifying the parameters of routing and signaling protocols.

3.5. Taxonomy of traffic engineering systems

We now provide an overview of taxonomy of traffic engineering systems following the discussion in [2]. The taxonomy provides a classification system for different types of traffic engineering methodologies. The taxonomy derives from “traffic engineering styles” which are abstractions of important traffic engineering methodologies. The classification system is dominated by a set of fundamental dichotomies. Fig. 5 illustrates some of the basic dichotomies in the classification of traffic engineering systems.

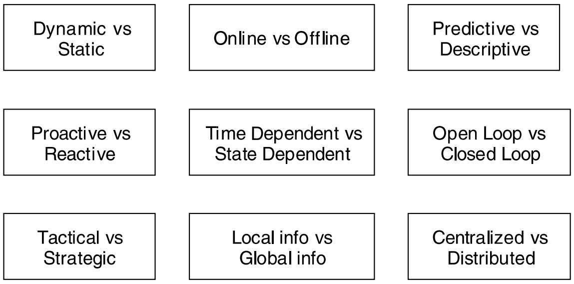

Fig. 5. Dichotomies in the classification of traffic engineering systems.

As shown in Fig. 5, the dichotomies that underpin the traffic engineering taxonomy include:

- Dynamic versus static traffic engineering.
- Offline versus online traffic engineering.
- Predictive versus descriptive traffic engineering.
- Proactive versus reactive traffic engineering.
- Time-dependent versus state-dependent traffic engineering.
- Open loop versus closed loop traffic engineering.
- Tactical versus strategic traffic engineering.
- Traffic engineering methodologies based on local information versus methodologies based on global information.
- Centralized versus distributed traffic engineering.

4. MPLS-based traffic engineering in IP networks

The fundamental requirements for traffic engineering over MPLS were laid out in RFC-2702 [1]. The motivation for MPLS-based traffic engineering can be traced back to the limitations of classical IP routing protocols which are based on shortest path concepts using a single additive metric, without consideration of network constraints, resource availability, and traffic characteristics. Another limitation of legacy IP systems has to do with measurement, particularly the lack of ability to determine the traffic matrix of an IP network. A traffic matrix is one of the most important input parameters for traffic engineering. It is very difficult to optimize the performance of a network in the absence of reliable data concerning the traffic matrix and other operational statistics. It turns out that MPLS can assist in both the enhancement of routing control functions and in the estimation of traffic matrices in IP networks.

The effect of the limitations associated with legacy IP systems, as noted earlier, is that traffic in a network can be localized to a subset of network resources, causing congestion in those segments of the network, even though excess capacity may exist elsewhere within the same network along alternative feasible paths. It is quite difficult to address this situation by manipulating the link metrics associated with interior gateway routing protocols.

We shall shortly discuss some of the capabilities making MPLS attractive for traffic engineering in IP networks. In the next paragraph, however, we first take a look at the classical overlay model based on IP over ATM and IP over frame relay, which was one of the techniques employed by a number of large backbone Internet service providers to circumvent some of the issues surrounding traffic engineering with legacy IP interior gateway routing protocols.

4.1. Overlay traffic engineering: IP over ATM and IP over frame relay networks

Prior to the advent of MPLS, large Internet service providers with dense core network topologies discovered that virtual connection-based abstractions with originating connection control capabilities are a good way to compensate for some of the limitations of legacy routing protocols, and to control and modulate the placement of traffic onto available network resources. The way they have gone about this, however, has been quite expensive. The general methodology is to introduce a secondary technology with traffic management and virtual circuit switching capabilities (such as ATM or frame relay) into the IP infrastructure in an overlay configuration. Elements of the secondary technology are placed at the core of the IP infrastructure and are surrounded by an epidermis of IP routers [4]. The virtual circuits of the secondary technology serve as point to point connections between IP routers over which routing protocols establish adjacencies so that routers connected directly by virtual circuits appear to each other as neighbors in the IP routing layer. Thus, in essence, the virtual circuits of the secondary technology appear as physical links to IP routing protocols.

One of the key characteristics of the overlay network is that the control plane of the interior virtual-circuit-based network is completely decoupled and independent of the control plane of the client overlay IP network. In the case of IP over ATM, the interior ATM network uses PNNI as the control plane to establish and deactivate virtual circuits. The client IP network uses conventional IETF IP control plane protocols (e.g., OSPF, BGP, etc.) Fig. 6 illustrates the topological configuration of the classical IP over ATM overlay network.

In the overlay configuration, the traffic management and constraint-based routing capabilities of the secondary technology (e.g., ATM or frame relay) can be exploited to implement traffic engineering objectives. For example, virtual circuits can be rerouted to move traffic away from congested resources onto under-utilized alternatives. These types of configurations also allow the service provider to derive an estimate of the traffic matrix by measuring and characterizing the traffic flow

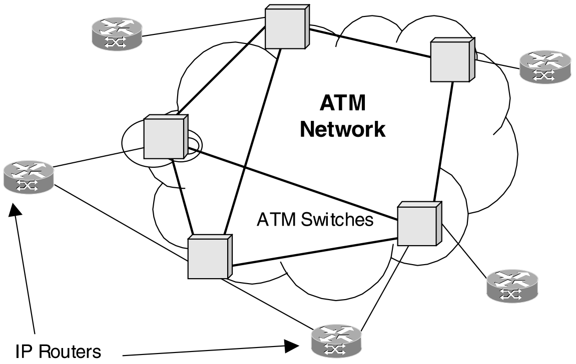

Fig. 6. Classical IP over ATM overlay network.

across the virtual circuits that inter-connect the routers.

There are many disadvantages, however, associated with the IP over ATM and IP over frame relay overlay techniques for traffic engineering. The most substantial limitations have to do with the added cost of building and managing two independent networks with dissimilar technologies and different operational semantics. Additionally, the number of virtual circuits in the overlay approach, hence the number of adjacencies between routers, generally grows as a function of the square of the number of routers in the network. This is the so-called $O(N^2)$ scaling problem with the overlay model for traffic engineering in IP networks.

4.2. Attractiveness of MPLS for traffic engineering

The attractiveness of MPLS for traffic engineering arose from the fact that it can provide equivalent (and sometimes superior) capabilities to the overlay model in an integrated fashion on a single network element. Some of the advantages that MPLS offers relative to the overlay model are (1) fewer network elements, (2) lower operating costs, (3) greater reliability because fewer network elements exist along the routed path, (4) potentially less latency, and (5) simplified network architectures. It should be noted that MPLS also supports the overlay model, giving service providers the option to deploy overlay or integrated solutions using a common MPLS technology in IP networks. The IETF RFC-2702 [1] outlined a set of capabilities which when added to MPLS allows it to serve as an effective means to implement various traffic engineering policies in IP networks. These requirements have resulted in the extension of IP signaling protocols (e.g., RSVP-TE [5]) and routing protocols (e.g., IS–IS [10] and OSPF) to support the new traffic engineering capabilities. The IP signaling and routing protocols along with the extensions mandated by MPLS traffic engineering represent what is generally termed the "MPLS traffic engineering control plane". This was reviewed briefly in Section 2 of this paper. Subsequent proposals within the IETF have expanded these requirements to encompass additional capabilities to support Diffserv-aware traffic engineering in networks deploying both MPLS and Diffserv [18]. Furthermore, these capabilities have been extended to provide control plane capabilities for other transport network technologies under the banner of MPλS [3] and GMPLS [24].

4.3. Fundamental problems of traffic engineering over MPLS

As noted in [1], there are three fundamental problems surrounding traffic engineering over MPLS networks. The first problem concerns mapping ingress traffic into FECs. The second problem involves mapping FECs onto LSPs. The third and last main problem involves mapping LSPs onto the physical network topology.

4.4. Protocol extensions to support MPLS traffic engineering

One of the main objectives of the MPLS traffic engineering requirements is to introduce various capabilities to allow constraint-based routing to be implemented cost-effectively in IP networks. The requirements for traffic engineering over MPLS propose several attributes that can be associated with "traffic trunks" to specify their behavioral characteristics and performance requirements, and various attributes that can be associated with network resources to specify various resource attributes and constraints, and to modulate the routing of traffic trunks over them. A traffic trunk essentially consists of traffic belonging to the same class that are routed through a common path or multi-path. In contemporary MPLS terminology, the term LSP-tunnel is generally used to refer to both the traffic trunk and the explicit LSP through which it traverses.

The traffic engineering extensions to MPLS support the assignment of various types of attributes to LSP-tunnels, such as bandwidth characteristics, resource affinities, resilience attributes, priority attributes, preemptive capabilities, and many others. The bandwidth characteristics indicate the capacity requirements of the LSP from the network. The resource affinities is a powerful means to indicate general classes of resources to include or exclude from the path of an LSP-tunnel. The resilience attributes indicate the survivability requirements of an LSP-tunnel. The priority attributes impose a partial order between different LSP-tunnels. The preemptive capabilities stipulate the conditions under which one LSP-tunnel can preempt another when they contend for the same resources. The MPLS traffic engineering extensions also support the ability to associate various attributes with network resources, such as capacity constraints, over-subscription factors, resource class attributes, and others. The attributes associated with resources are disseminated in link state advertisements by the interior gateway routing protocols which have been extended to support this new capability.

With the advent of MPLS traffic engineering, conventional IP routing protocols such as IS–IS and OSPF have been extended to advertise new types of capabilities and constraints associated with links. Some of the new enhancements to IP routing protocols include the assignment of traffic engineering metrics to links, assignment of resource class attributes to links, the advertisement of maximum link bandwidth, and the advertisement of maximum reservable link bandwidth. The value of the maximum reservable link bandwidth can be manipulated by a network operator to over-subscribe or under-subscribe a link.

The extensions to IP signaling protocols (RSVP and LDP) have been much more fundamental than the corresponding extensions to routing. In the case of RSVP, several new objects have been added to support the establishment and teardown of explicit LSPs with various types of behavioral attributes. The new version of RSVP with the traffic engineering extensions is called RSVP-TE and is documented in [5] as a standards track Internet RFC. Furthermore, while the original RSVP specification was intended to be used by hosts to request and reserve network resources for microflows, the traffic extensions permit RSVP-TE to be used by network elements (e.g., label switching routers) to establish parameterized explicit LSPs and assign network resources to them. Some of the new objects introduced in the RSVP-TE specification include: a LABEL-REQUEST object, RECORD-ROUTE object, LABEL object, EXPLICIT-ROUTE object, and new SESSION objects. New RSVP error messages have also been added to provide notification of anomalous conditions [5,6]. It should be noted that even though two signaling protocols are currently supported for MPLS traffic engineering, the RSVP-TE specification [5,6] has emerged as the dominant protocol utilized in operational networks and implemented by most network equipment manufacturers.

With the traffic engineering capabilities for MPLS, the operational aspects associated with establishment of LSP-tunnels are substantially simplified. An LSP-tunnel can be established by configuring its characteristics (endpoints plus the desired performance and behavioral attributes) at an originating LSR. The originating LSR will then employ an appropriate constraint-based path computation algorithm to compute a path through the network satisfying the LSP-tunnel specifications subject to various constraints that exist within the network. Once the path is successfully computed, the originating LSR will subsequently use an appropriate signaling protocol (e.g., RSVP-TE) to establish the LSP-tunnel.

4.5. Operational considerations for MPLS traffic engineering

We now turn our attention to some pragmatic aspects by considering some of the operational issues in the deployment of MPLS traffic engineering solutions.

The way traffic engineering is actually conducted in some operational IP networks is that the network operator will configure LSR interfaces and assign routing and traffic engineering attributes to them. This information is subsequently flooded throughout the routing area by the interior gateway routing protocol with traffic engineering extensions (e.g., OSPF-TE). Once the protocol specified aspects of LSR interface configuration management are concluded, the operator then commences to configure parameters relating to LSPs. The attributes of LSPs that are configured include the destination endpoint, the miscellaneous parameters of the LSP such as bandwidth, priorities, affinities, and resilience properties. The bandwidth assigned to an LSP can be based on some notion of “effective bandwidth” which is derived during the analysis and characterization phase of the traffic engineering process model. Once these attributes of an LSP are configured, a constraint-based routing mechanism at the originating LSR will then compute an appropriate path through the network satisfying the LSP attributes subject to prevailing network constraints. Once a path is selected, the signaling protocol (e.g., RSVP-TE) will then be invoked to dynamically establish the LSP.

LSP topology: One aspect that needs to be given adequate thought from an operational perspective is the layout of the LSP virtual topology. By this, we mean determining which nodes will function as endpoints of LSP-tunnels. Experience suggests that networks with regular and well structured virtual topologies are easier to manage, but this regularity may come at the cost of loss in efficiency. In any case, special consideration should be given to whether to deploy a large number of LSP-tunnels or a smaller number. A large number of LSPs allows optimizing the network more effectively, but may however result in significant operational complexity. On the other hand, fewer LSPs are easier to manage, but may result in avoidable network inefficiencies.

Load balancing across multiple parallel LSPs with common endpoints is an important practical traffic engineering problem in operational networks. Load balancing across multiple LSPs is imperative in circumstances where the traffic demand between the common endpoints exceeds the capacity of a link or router interface along the maximum bandwidth path between the endpoints. The basic concept behind load balancing across multiple parallel LSPs is to partition arriving traffic (according to some principle of partitioning) and assign the partitioned traffic onto the parallel LSPs (according to some principle of allocation) to achieve a network performance objective. The assignment and allocation of arriving traffic to parallel LSPs can be based on dynamic or static considerations, and may be open loop or closed loop. The open loop scenario occurs when the partitioning and assignment of traffic onto the parallel LSPs does not utilize dynamic feedback information from the network to modulate the decision process. Closed loop load balancing refers to the scenario in which the partitioning and assignment of traffic onto parallel LSPs is influenced by dynamic feedback information from the network. Load balancing across multiple parallel LSPs may also consider local policies involving CoS considerations, especially in Diffserv-aware MPLS networks.

4.6. Network survivability

Reliable network operation is an important aspect of Internet traffic engineering. In particular, adequate consideration must be given to the issue of survivability of LSPs when network faults occur in an operational context. The ability to offer enhanced survivability capabilities on a per LSP basis is one of the many benefits that MPLS offers in IP networks. Different types of protection, restoration, and local repair schemes are feasible with MPLS. The reader is referred to [25–27] for a framework of MPLS recovery techniques and a discussion of signaling enhancements. The Internet draft [28] discusses the extensions to RSVP-TE for establishment of backup LSP-tunnels for local repair (i.e., recovery at intermediate segments of LSP-tunnels when failures occur) of explicit LSP routes.

4.7. Measurement considerations

Another important operational consideration in MPLS traffic engineering is the measurement system. Generally, it is desirable to have a view into the route traversed by each LSP in the network, to obtain traffic statistics within an LSP, to monitor bandwidth requirements of each LSP, and to monitor the dynamics of LSPs in the network. In a differentiated services environment, it may also be desirable to measure the delay along an LSP under different conditions. Deriving a traffic matrix from measured statistics is one of the fundamental issues in traffic engineering.

5. Diffserv-aware traffic engineering

We now turn our attention to the MPLS extensions to support Diffserv-aware traffic engineering. MPLS and Diffserv are two important components of resource allocation in future IP networks. The traffic engineering capabilities provided by MPLS facilitate effective routing and global resource allocation within a given domain, integrating advanced constraint-based routing with bandwidth resource allocation. Diffserv deals with local resource allocation (so-called per hop behaviors—PHB). In particular, it deals with the allocation of buffer and link resources to packets based on the Diffserv code point (DSCP) in the packet headers.

There are actually two aspects to MPLS support of Diffserv. The first has to do with basic support for Diffserv within MPLS itself. The second concerns the actual traffic engineering considerations in MPLS–Diffserv networks (that is, networks that concurrently implement MPLS and Diffserv). The MPLS support for basic Diffserv is specified in [13], and essentially stipulates how Diffserv behavior aggregates can be mapped onto LSPs. Two types of LSPs are defined to support this capability. The first type of LSPs are those that can carry different types of ordered behavior aggregates within the same LSP. These are called EXP-inferred-LSPs or E-LSPs because the behavior aggregate of each packet is inferred from the EXP bits (experimental bits) in the MPLS label associated with the packet. The second type of LSPs are those that carry only one type of behavior aggregate. These are called Label-inferred-LSPs or L-LSPs, because the behavior aggregate for packets is inferred from the label assigned to each packet.

The motivation for the MPLS requirements to support Diffserv aware traffic engineering derives from the fact that the original MPLS traffic engineering proposals focused on the optimization of aggregated traffic trunks, without adequate consideration to the issue of preferential treatment to different types of traffic in a Diffserv environment. There are two issues in the original MPLS-TE requirements that have direct bearing on Diffserv-aware traffic engineering, namely the concepts of priorities and preemption. The detailed requirements for Diffserv-aware traffic engineering are contained in [18]. These requirements allow traffic engineering to be applied at a finer granularity, on a per class basis, in MPLS–Diffserv networks, so that the service and performance requirements of each class can be accommodated. Diffserv-aware traffic engineering is particularly relevant in network scenarios where capacity is scarce, and where traffic belonging to different behavior aggregates contend for network resources.

With Diffserv-aware traffic engineering, different bandwidth constraints can be specified on network elements for different classes of traffic—so that different classes have different views of network resource availability. Essentially, this allows the service provider to carve the network into different capacitated virtual networks for different traffic classes which co-exist within the same common underlying infrastructure. One of the requirements for Diffserv-aware traffic engineering is that the class specific virtual networks have to be work conserving, which means that if a high priority traffic class does not use up all its allotted bandwidth, the remaining bandwidth can be utilized by other service classes.

6. Analytical approaches to MPLS traffic engineering

We now turn our attention to analytical modeling and mathematical formulation of the MPLS traffic engineering problem. These activities are mostly concerned with the third and fourth phases of the traffic engineering process model, described in Section 3. There are many sub-problems involved in the performance optimization of operational MPLS networks. Three of the most significant problems include: (1) constraint-based routing, (2) traffic partitioning and assignment, and (3) restoration. It should be noted that even though these problems are well known in other application domains, they are still in a state of infancy concerning MPLS, and much remains to be done. Accordingly the literature addressing these areas is somewhat limited, at this time.

The problem of constraint-based routing deals, in general, with the computation of paths for LSPs subject to various types of constraints. The constraints themselves may be inherent to the network (e.g., available bandwidth) or they can be administratively specified (e.g., affinities and resource class attributes, and diversity requirements for protection and restoration). The computational aspects of constraint-based routing can be performed online or offline. Generally, these problems are NP-complete. This means that simple heuristics that produce good "engineering" solutions in reasonable time must be employed for online path computation. On the other hand, offline path computation can employ more sophisticated heuristics. Among the early work in the area of offline MPLS constraint-based routing is the contribution by Fahim [19], which utilized a centralized global optimization algorithm.

A heuristic online path selection method, referred to as minimum interference routing algorithm (MIRA) is presented in [20]. This algorithm, which requires a priori knowledge of ingress–egress pairs, attempts to defer loading of certain so called "critical links". The critical link is defined as those links whose congestion will cause blocking of future LSP setup requests between more than one ingress–egress pair. There are, however, computational complexities associated with this algorithm.

The second important MPLS traffic engineering optimization problem deals with the optimal partitioning and assignment of traffic to parallel LSPs between pairs of MPLS ingress and egress nodes. One aspect of this problem deals with the dynamic control of the partitioning of traffic and the assignment of the partitions to parallel LSPs to optimize network performance. Mathematical formulations of this problem are provided in [8,9,16]. In the next paragraph, we briefly review the approach described in [8,9].

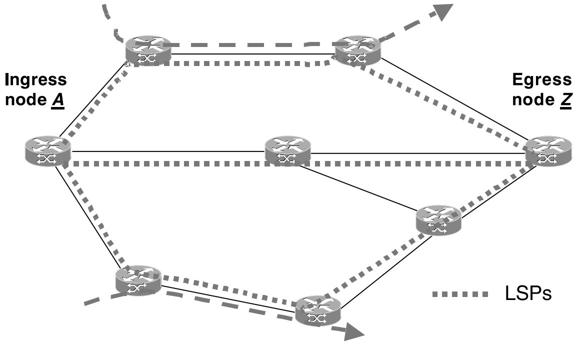

Fig. 7. Illustrative network with three parallel LSPs between two nodes.

Fig. 7 depicts three paths between ingress node ***A*** and egress node ***Z***. The question addressed in [8,9] is how to map the input traffic (which arrives according to a stochastic process) dynamically and efficiently onto the parallel paths. The approach is based on developing an analytical model to obtain the optimal partitioning of ingress traffic and the subsequent mapping of the traffic onto the parallel LSPs, taking into account the current state of the network. Each LSP was modeled by a sequence of queues and to simplify the problem, so that each node along the LSP was represented by a queue and each queue was characterized as an M/M/1/K system. The model takes into account the aggregate traffic arriving at each nodal queue, some of which is contributed by traffic from the target LSP traversing the node, while the remaining traffic is contributed by all other LSPs that traverse the node. An iterative methodology was then applied to solve the resulting problem.

The third important MPLS traffic engineering problem, path restoration in MPLS networks, has not been studied extensively in the literature. Several proposals dealing with MPLS protection, restoration, and local recovery have been submitted to the IETF. There are many approaches that have been proposed for restoration in ATM networks which might be applicable to MPLS with some modifications. The method of mapping traffic onto parallel LSPs can also be used to implement graceful performance degradation under failure scenarios. In this approach, when a failure impacts one of the LSPs in the parallel configuration, the traffic originally assigned to it is reassigned to the remaining LSPs using an appropriate partitioning and assignment methodology.

7. Generalized MPLS

Perhaps the most significant advancement in the evolution of the MPLS is the extension and generalization of the MPLS traffic engineering control plane to serve as the control plane for other types of transport networks, including TDM networks (e.g., SONET/SDH) and optical transport networks. This effort has been embarked upon by the IETF under the acronym Generalized MPLS (GMPLS).

GMPLS is a suite of control plane protocols that provides consistent and uniform semantics for signaling, routing, and link management in different types of transport networks [24]. GMPLS strongly advocates and promotes explicit separation of the control plane from the underlying dataplane or transport infrastructure. GMPLS allows products from different vendors to inter-operate at a control level in different types of switched transport networks. GMPLS also allows new and innovative ways to inter-connect various technologies and different layers, without restricting the way individual layers interwork with each other. Ultimately, GMPLS will simplify the design, deployment, and operations management of heterogeneous networks consisting of an assortment of packet switched and circuit switched equipment from different manufacturers.

7.1. Origins of GMPLS

Fig. 8 depicts the evolution of GMPLS. The origins of GMPLS can be traced back directly to the multi-protocol lambda switching (MPλS) concept originally proposed by Awduche and Rekhter, which was submitted to the IETF as an Internet Draft in 1999 (see [3]). The main idea underlying MPλS is the adaptation, specialization, and reuse of control plane concepts originally developed for MPLS traffic engineering in optical networks. The traffic engineering control plane itself (traffic engineering extensions to IP routing and signaling protocols) arose in response to the requirements stipulated in [1].

The applicability of the MPLS traffic engineering control plane to the optical domain, as propounded in the MPλS proposition, depends very much on the conceptual commonalities that exist between label switching routers and optical cross-connects, coupled with the commonalities that exist between explicit LSPs and optical channel trails. These commonalities were highlighted in the MPλS proposal, along with various interesting architectural possibilities brought about by the new approach [3]. However, there are several unique features of optical transport networks that must be taken into consideration in the design of control plane technologies for this domain. These include the need to convey additional topology state information to capture some of the peculiar characteristics of optical networks, and the need for more complex constraint-based path selection algorithms. Furthermore, failures occurring in the control plane of optical networks should not impact established optical connections carrying user traffic.

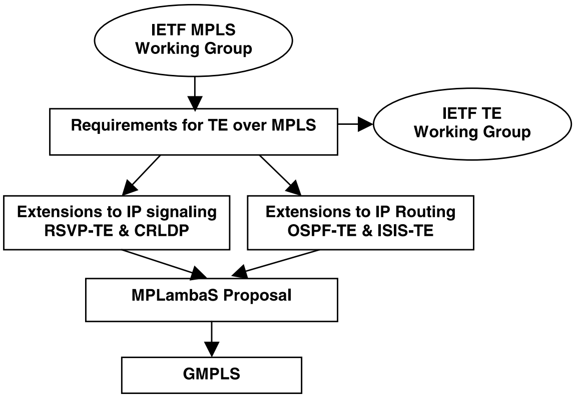

Fig. 8. Evolution of GMPLS.

GMPLS has generalized the MPλS concept, so that the same control plane concepts can be used in other switched transport technologies, such as TDM and optical networks for example, as well as traditional packet and cell switched networks. In particular, GMPLS extends the concept of a “label”, so that: (1) in a packet-switched network, a label represents a short tag attached to a packet; (2) in a TDM network, a label represents a time slot; (3) in a wavelength-switched network, a label represents a wavelength; and (4) in a fiber-switched network, a label represents a fiber. To support these extensions, a broad range of interface types have been defined over which a GMPLS control plane can exercise control. These include: packet switch capable interfaces (PSC), TDM, lambda switch capable interface (LSC), and fiber switch capable interface (FSC).

GMPLS consists of three main aspects: routing, signaling, and link management. GMPLS also explicitly decouples the control channel (over which control information flow) from the transport or bearer channels (over which user traffic traverse). This decoupling has important implications on the fault handling characteristics of the control plane, because failure in the control plane does not necessarily imply failure of the transport plane, unlike in conventional IP routing for example. GMPLS also requires bi-directional control channels between two adjacent nodes, even if the neighboring nodes are inter connected by unidirectional links at the transport level.

The routing component of GMPLS essentially consists of new extensions to conventional IP routing protocols (IS–IS and OSPF), on top of the previous extensions for MPLS traffic engineering (see e.g., [22]). The main issues for routing in general transport networks center around selecting the best path (or set of paths) for the transport of traffic across the network. This activity requires neighbor discovery, network resource discovery, topology state information acquisition and dissemination, topology state information management, and path selection. The last issue mentioned, namely path selection, is clearly a critical consideration in the design of control planes for switched transport networks, but it is not directly covered by the GMPLS specifications, because the algorithmic aspects of explicit path computation do not require direct inter-operability, except to provide topology state information and associated constraints which serve as input in the path selection process. The bulk of the GMPLS extensions to conventional IP routing protocols deal with the ability to acquire, represent, disseminate, and manage new types of link information. The concept of link is generalized in GMPLS to admit a variety of constructs with different properties that support topological adjacency between two nodes. Additional mechanisms, such as link bundling, have been introduced to enhance the scalability of the routing component in transport networks in which multiple links can exist between two nodes (e.g., DWDM systems). The concept of forwarding adjacency, which allows a node to advertise a link which was previously established by its own control plane, has also been made an integral aspect of GMPLS routing. In the case of OSPF, for example, the opaque LSA has been augmented with new TLVs to support additional traffic engineering characteristics of transport networks. Some of the new link characteristics include: incoming and outgoing interface identifiers, link protection type, shared risk link groups, and interface switching descriptor.

In the case of signaling, GMPLS has introduced many enhancements to the MPLS traffic engineering signaling protocols [21]. The concept of ‘label’ has been generalized, as noted earlier, to support the reconfiguration of various types of switching elements in transport networks. Accordingly a ‘generalized label object,’ has been added to the signaling protocols (e.g., RSVP-TE). Another new signaling extension with GMPLS is support for bi-directional LSPs (the original MPLS specifications supported only unidirectional LSPs). The concept of ‘suggested label’ has also been included to allow an upstream node to suggest a label to a downstream node (the original MPLS signaling protocols support downstream-on-demand label distributed where labels are exclusively assigned by downstream nodes). The intent of the suggested label is to reduce setup latency, by allowing the upstream not to reconfigure its switching matrix before it receives an explicit label binding from the downstream node, but it can be applied in optical networks with limited wavelength conversion capability to perform wavelength assignment by upstream nodes. Another new signaling capability, driven by optical networks, is the ‘label set’ concept, which allows an upstream node to restrict the range of labels that a downstream node can allocate. Again, this capability can be used for wavelength assignment. The label set restriction can be imposed on a single hop or along the entire LSP path. A failure indication mechanism has been added to the signaling protocol to allow a downstream node to notify upstream of failures. This feature is imperative for fault recovery in optical networks where loss of light and other traffic impairments are usually detected by downstream nodes, but not upstream nodes. Therefore, upstream nodes need to be notified of faults detected by downstream nodes to effect appropriate recovery policies. Another GMPLS signaling enhancement is the ability to include technology specific characteristics in the signaling protocol. Lastly, a mechanism is provided to support the splicing together of two LSP segments. In short, almost every aspect of the original MPLS traffic engineering signaling protocol extensions have been further refined and enhanced to support the GMPLS concept.

The third main component of GMPLS is the link management protocol (LMP), which is a new protocol that has been specifically developed with the advent of GMPLS [23]. The motivation for LMP arose from the observation that the control channel between two adjacent nodes in general transport networks (e.g., optical networks) is generally decoupled from the bearer channels. This means that it is not possible to make useful inferences about the condition of the bearer channels from the condition of the control channels. Furthermore, the bearer channels between two adjacent nodes may consist of thousands of links. LMP was designed to deal with these types of scenarios. LMP runs between two adjacent nodes and is used for both link property correlation and control channel management. LMP implements a 'hello' protocol to detect control channel failure between adjacent nodes. The link property correlation feature is used to advertise various properties of component links in the underlying bearer network between two adjacent nodes. LMP can also be used for link property correlation and fault management between adjacent nodes.

8. Future directions

This section explores some possible future directions in the evolution of the operational aspects of MPLS-based Internet traffic engineering. We will highlight aspects relating to policy-based MPLS network management, customer network management (CNM), and the related issue of advanced service level agreement (SLA) management. Constraint-based routing will remain an area of focal activity into the foreseeable future concerning MPLS traffic engineering. The areas of IP over optical architectures and inter-connection models will be an area of significant research and development activities, especially when both the IP and optical domains utilize GMPLS control plane technologies. Inter-domain traffic engineering is yet another important research problem worth mentioning.

8.1. Policy-based MPLS network management

The ultimate goal of policy-based network management is to provide the capability to manage heterogeneous networks in a uniform fashion, preferably from business directives without fixation on the underlying technologies. Policy-based management is the next critical phase in the evolution of MPLS traffic engineering. Policy-based network management involves establishing a level of abstraction in the network control and management software systems that allows masking the technological characteristics of the network. We distinguish between two levels of policy-based network management: (1) high order policy-based management and (2) low order policy-based management. High order policy-based management is concerned with creating an abstraction layer between business logic and network logic. Low order policy-based management is implemented within the network itself and involves resolving low level policy issues within the network. Examples of low order policy-based management within the network are the application of control policies in the selection of paths for LSPs, in the assignment of bandwidth and other resources to LSPs, in the reservation of resources for LSPs, in the mapping of traffic onto LSPs, in establishing criteria for service policies within a network element (queueing, scheduling, rate shaping, policing), and in the recovery and restoration of traffic under network fault conditions.

As shown in Fig. 9, a high order policy-based management infrastructure contains the following basic components: (1) a policy management interface, (2) a policy decision point, (3) a policy repository, and (4) policy enforcement points. The policy repository is an interface to the policy management system. It allows users to specify and submit policy statements derived from business and engineering directives in the form of policy schemas. The policy repository stores persistent policy information. The policy decision point

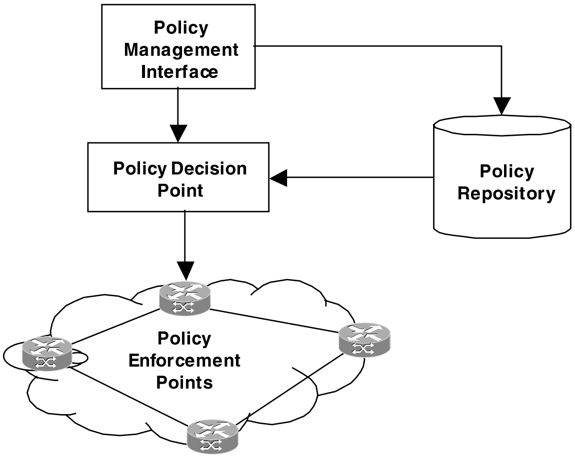

Fig. 9. Generic functional components for MPLS policy-based management.

converts high order policy logic into network management and control logic. The policy enforcement points (which will typically reside within the network) are responsible for executing the final decisions by implementing or activating network control functions. Within the context of policy-based management of MPLS networks, the MPLS traffic engineering control plane can be viewed as both a policy decision point and a policy enforcement point. The research and development challenge is to explore and understand the various facets relating to policy-based management of MPLS networks, and in particular heterogeneous multi-technology GMPLS networks as the underlying infrastructure continues to evolve.

8.2. Customer network management

CNM will be another important future phase in the evolution of MPLS-based networks. CNM allows a customer to modify and monitor the services he or she enjoys from the network by interacting with a CNM portal situated within the service provider network. The business drivers for CNM center around operational cost reduction and enhancing the economies of scale of the network, so that the cost of customer service will not increase proportionately with the number of customers subscribed to the network. This means particularly that the cost of operating the network will not grow proportionately with the size of the network, and the subscribed customer base. Closely related to the concept of CNM is the notion of end-to-end flow through provisioning which enables provisioning of network services in a completely automated fashion, without human intervention at intermediate points in the network. Another related concept is the idea of *advanced SLA* management which can exploit the MPLS traffic engineering capabilities in large-scale IP networks.

8.3. IP over optical inter-connection architectures and models

Inter-connection models for IP over optical network architectures will be an area of considerable interest. The inter-connection models that are under active consideration today include the overlay model, the peer model, and the augmented model. The peer model, in particular, presents many interesting conceptual and practical challenges relating to security, scalability, fault containment, performance optimization, routing control, signaling control, link management, resource allocation, etc. These issues are largely unexplored because the peer model is the consequence of the recent introduction of MPLS and GMPLS.

8.4. Inter-domain traffic engineering

Most of the industrial activities relating to Internet traffic engineering have centered around intra-domain traffic engineering, that is traffic engineering within a given autonomous system in the Internet. The issue of inter-domain traffic engineering, that is traffic engineering across autonomous systems, is an important topic in need of more rigorous studies [7].

9. Conclusion

This paper described the basic concepts of MPLS and its applications to Internet traffic engineering. The process model for traffic engineering was also discussed, along with traffic engineering considerations in combined MPLS and Diffserv networks. Continuing advances in technology will result in changes in the way traffic engineering is performed in the Internet. For example, the emergence of intelligent optical internetworking systems in the future, with sophisticated bandwidth provisioning capabilities and dynamic wavelength routing based on GMPLS will have a significant impact on traffic engineering in core IP networks. Coupled with these are fundamental research and development issues that remain unexplored in constraint-based routing, policy-based management of MPLS networks, CNM, and IP over optical architectures and interconnection models utilizing GMPLS.

References

[1] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, J. McManus, Requirements for traffic engineering over MPLS, RFC 2702, September 1999.

[2] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja, X. Xiao, Overview and principles of Internet traffic engineering, IETF Internet Draft, Work in Progress, November 2001.

[3] D. Awduche, Y. Rekhter, Multiprocotol lambda switching: combining MPLS traffic engineering control with optical crossconnects, IEEE Communications Magazine, March 2001.

[4] D. Awduche, MPLS and traffic engineering in IP networks, IEEE Communications Magazine, December 1999.

[5] D. Awduche, L. Berger, D. Gan, T. Li, G. Swallow, V. Srinivasan, RSVP-TE: extensions to RSVP for LSP Tunnels, RFC-3209, December 2001.

[6] D. Awduche, A. Hannan, X. Xiao, Applicability statement for extensions to RSVP for LSP-tunnels, RFC-3210, December 2001.

[7] D. Awduche et al., An approach to optimal peering between autonomous systems in the Internet, in: International Conference on Computer Communications and Networks (ICCCN'98), October 1998.

[8] E. Dinan, D. Awduche, B. Jabbari, Analytical framework for dynamic traffic partitioning in MPLS networks, in: IEEE International Conference on Communications, ICC 2000, New Orleans, Louisiana, 18–22 June, 2000.

[9] E. Dinan, D. Awduche, B. Jabbari, Optimal traffic partitioning in MPLS, in: Networking 2000, 14–19 May, 2000, Paris, France.

[10] H. Smit, T. Li, IS–IS extensions for traffic engineering, in: IETF Internet Draft, Work in Progress, February 2001.

[11] E. Rosen, A. Viswanathan, R. Callon, Multiprotocol label switching architecture, RFC 3031, January 2001.

[12] X. Xiao, A. Hannan, B. Bailey, L. Ni, Traffic engineering with MPLS in the Internet, IEEE Network Magazine, March 2000.

[13] F. Le Faucheur et al., MPLS Support of Differentiated Services, IETF Internet Draft, Work in Progress, February 2001.

[14] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, NetScope: traffic engineering for IP networks, IEEE Network Magazine (2000).

[15] B. Jamoussi et al., Constraint-Based LSP Setup using LDP, IETF Internet Draft, Work in Progress, February 2001.

[16] A. Elwalid, C. Jin, S. Low, I. Widjaja, MATE: MPLS adaptive traffic engineering, in: Proceedings of INFOCOM'01, April 2001.

[17] E. Rosen, D. Tappan, G. Fedorkow, Y. Rekhter, D. Farinacci, T. Li, A. Conta, MPLS label stack encoding, RFC-3032, January 2001.

[18] Le Faucheur et al., Requirements for support of Diff-Serv-aware MPLS traffic engineering, IETF Internet Draft, Work in Progress, 2001.

[19] F. Fahim, Constraints based routing algorithms and applications, in: Proceedings of MPLS'98, McLean, Virginia, October 1998.

[20] K. Kar, M. Kodialam, T.V. Lakshman, Minimum interference routing of bandwidth guaranteed tunnels with MPLS traffic engineering applications, IEEE Journal on Selected Areas in Communications 18 (12) (2000).

[21] P. Ashwood-Smith et al., Generalized MPLS: Signaling Functional Description, IETF Internet Draft, Work in Progress, 2001.

[22] K. Kompella et al., OSPF Extensions in Support of Generalized MPLS, IETF Internet Draft, Work in Progress, 2001.

[23] J. Lang et al., Link Management Protocol (LMP), IETF Internet Draft, Work in Progress, 2001.

[24] E. Mannie et al., Generalized MPLS (GMPLS) Architecture, IETF Internet Draft, Work in Progress, 2001.

[25] V. Sharma et al., Framework for MPLS-based Recovery, IETF Internet Draft, Work in Progress, July 2001.

[26] A. Banerjee, Generalized MPLS: An overview of signaling enhancements and recovery techniques, IEEE Communications Magazine (July 2001) 144–151.

[27] C. Huang, V. Sharma, K. Owens, S. Makam, Building reliable MPLS networks using a path protection mechanism, IEEE Communications Magazine (March 2002) 156–162.

[28] P. Pan et al., Fast Reroute Extensions to RSVP-TE for LSP Tunnels, IETF Internet Draft, Work in Progress, January 2002.