Virtual Network Resource Management for Next-Generation Networks

Alberto Leon-Garcia, University of Toronto
Lorne G. Mason, McGill University

ABSTRACT

The multilayer structure of transport networks presents a key challenge in network resource management. In this article we present a virtual network approach to the management of transport network resources. We describe the approach as developed for the management of large-scale ATM networks and discuss the application to the next generation of multilayer transport networks.

INTRODUCTION

Resource management in transport networks continues to present a challenge to network operators. Internet Protocol (IP), asynchronous transfer mode (ATM), synchronous optical network (SONET), and optical networks are all managed using separate systems. The ideal of a single resource management system within an integrated operations environment remains elusive because of the complexity inherent in supporting a diverse set of legacy and emerging transport technologies and the lack of suitable standards. In this article we present an integrated approach for managing network resources using the concept of virtual networks. This approach was developed for the management of ATM transport resources, but we believe provides scalability and capabilities that make it relevant to today's multilayer networks.

The evolution of network resources management and network scale are intertwined. Network resource management is concerned with the fair allocation of transport resources to competing users. Each increase in network scale implies an increase in complexity that must be accompanied by an advance in network management. Typically, the advance incorporates a new level of abstraction or virtualization that strikes the right balance between granularity, locality, and manageability. An example of this is the introduction in the early 1990s of the SONET architecture, which provided a major advance in network management capability. Today the SONET transport architecture is challenged by the explosion in Internet traffic and optical transmission capacity. The next generation of management systems must manage packet, time-division multiplexed (TDM) as well as optical flows that range in scale six orders of magnitude from megabits to terabits per second. Clearly we are ready for the next advance in network management.

The article is organized as follows. In the first part, we present the virtual network approach that was developed for ATM network resources management. In doing so, we also discuss the advances from ATM traffic management that laid the groundwork of the management of packet traffic. In the latter part, we apply the virtual network concept to today's multilayer networks.

MANAGING PACKET FLOWS: FROM ATM TO MPLS

During the 1990s there were two competing visions for service convergence, one using the ATM-based broadband integrated services digital network (B-ISDN) and the other evolving from the Internet, by way of the integrated services IP (IntServ) and differentiated services IP (DiffServ) architectures [1]. IntServ provided guaranteed service and controlled load service in addition to best effort. The Resource Reservation Protocol (RSVP) was developed to establish soft-state reservations for individual microflows, providing end-to-end quality of service (QoS) control through call admission and queue management procedures. The large number of soft states in the core network required to keep track of individual flows led to problems of scalability for the IntServ approach.

0163-6804/03/$17.00 © 2003 IEEE

The DiffServ architecture was advanced to achieve scalability in the core network. Differentiated class of service (CoS) was provided to flow aggregates rather than individual flows at the expense of relinquishing end-to-end QoS control for individual microflows. DiffServ offers a range of differentiated per-hop behaviors (PHBs) to aggregated flows, which are determined by packet classification and traffic conditioning at edge routers. DiffServ does not require explicit signaling as the PHB is encoded in each packet header and the standard Internet routing protocols such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) are employed. As a result DiffServ offers no capability for source routing. The bandwidth broker concept described in [1] does provide some traffic engineering capability for DiffServ.

By contrast ATM employs signaling for capacity reservation and call admission control (CAC), and the PNNI routing protocol. As a result ATM was not able to leverage the embedded investment in the Internet protocols such as OSPF, BGP and RSVP. On the other hand, ATM had the advantage of strong QoS support and was a proven high speed label switching technology. Routers employed in the Internet had to read the variable length IP header and consult the routing table to determine how the packet should be routed. As a consequence it was more difficult to perform this operation at high speed than the simple label swapping technique used with the fixed-length ATM cells.

The feverish period of competition led to the IP switching concept advanced by Ipsilon and Fujitsu, which attempted to combine the strengths of routing and switching by establishing cut-through paths for persistent flows while handling short flows by standard IP routing. Ipsilon proposed using the ATM switching hardware, but replaced the ATM control plane with protocols better suited to IP networks. Related research on Tag Switching at Cisco, ARIS at IBM, and CSR at Fujitsu were attempts to more efficiently handle flows of IP frames. These researchers ultimately collaborated under the aegis of the Internet Engineering Task Force (IETF) to produce the multiprotocol label switching (MPLS) standard [2].

MPLS can readily be installed in existing IP networks to add functionality and efficiency. Using RSVP or other IP-based signaling protocols, label switched paths (LSPs) can be set up on demand, by either users who are requesting connections to provide a given level of service quality or the network itself if such an action results in better performance and/or efficiency. At the same time, best effort traffic can be handled more efficiently with MPLS than by means of conventional routers. This layer 2.5 protocol can run over existing ATM switches where the virtual channel/virtual path identifier (VCI/VPI) header is treated as a label or over Ethernets and other layer two protocols by a shim label. This power and flexibility as well as the ability of MPLS to interwork with IntServ and DiffServ as well as best effort architectures has made it the protocol of choice for the envisaged next-generation Internet. With this advance, most observers believe that MPLS will take the place of ATM as the vehicle for convergence of services.

Initially MPLS is being introduced within some administrative domains, which typically do not span end-to-end path flows. Bilateral service level agreements (SLAs) are established between neighboring MPLS domains, which control the volume of traffic exchanged and its treatment in the downstream domain. In the initial deployment of MPLS, there is no end-to-end reservation, but rather only across a single MPLS domain. This facilitates the introduction of MPLS along with traffic engineering within a domain [3]. Here traffic engineering refers to the capability to distribute the flow from an origin node to a destination node across a set of LSPs linking the origin and destination nodes. This permits more efficient use of network resources and improved transport performance. The traffic engineering extension of OSPF, OSPF-TE, represents an effort to leverage the existing Internet routing protocols and provide the traffic engineering functionality. On the other hand, the absence of end-to-end control means that unless additional reservation and CAC procedures are put in place MPLS cannot provide end-to-end QoS support.

While an end-to-end ATM network provides strong QoS support, the large installed base of IP applications in end systems, as well as the low-cost availability of Gigabit Ethernet, preempted the extension of ATM from islands in the core to the local area and desktop. Nevertheless, the installed ATM base in the core is effectively reused as a particular implementation of MPLS. Moreover, by adding additional functionality to MPLS (e.g., admission control and bandwidth reservation), QoS support approaching that of the original ATM dream can be envisioned.

The installed ATM base in the core is effectively reused as a particular implementation of MPLS. Moreover, by adding additional functionality to MPLS, such as admission control and bandwidth reservation, QoS support approaching that of the original ATM dream can be envisioned.

Managing Layer Networks

In CITR-sponsored research, a framework was developed for large-scale ATM network resource management (NRM). An NRM system is responsible for fair allocation of transport network resources to competing users. The scope of the NRM system includes both traffic management (ATM cell-level controls and connection setup/release) and network management (performance monitoring and configuration management). Telecommunications Information Networking Architecture Consortium (TINA-C) provided a framework for developing these two traditionally separate NRM controls on a common object-oriented software platform in a distributed processing environment, allowing the sharing of common network resource information and management functionality. A key element of the NRM system was the extension of the notion of layer networks to include virtual networks [4, 5].

The layer network concept from transport networks can be applied to ATM to view the connection topology and resources at the VC, VP, and transmission path (TP) layers. The

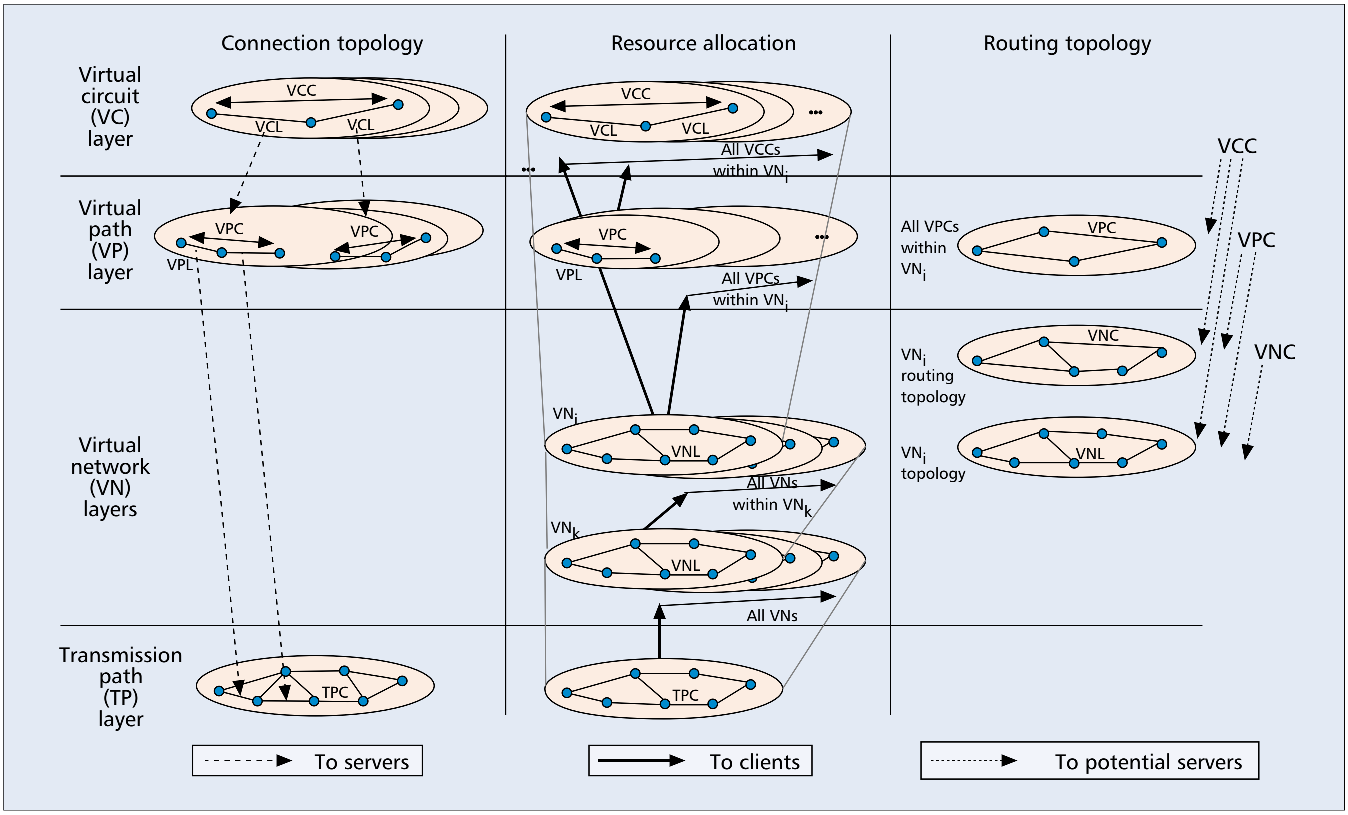

■ **Figure 1.** *Layer network extension to include the virtual network layer.*

VP layer is considered as a *server* to the VC layer, and the TP layer can be considered as a server for both the VC and VP layers. In the NRM context, a server layer allocates resources to its *client* layers. A client layer can also define the restricted topology over which routing procedures may be carried out by the client layer connections. Thus, in the NRM system, each layer provides a simplified logical view of a specific portion of network resources within which management functions can be carried out. With reference to Fig. 1, the virtual circuit and transmission path layers are represented by the first and third rows, respectively. The CITR work extended the layers to include a *virtual network* as shown by the middle row in Fig. 1.

The role of the NRM is to allocate resources among competing users and to monitor resource usage. The competition of resources occurs at several levels: At the cell level VCs and VPs compete for buffer resources and transmission time; at the connection level, VCs and VPs compete for end-to-end connectivity that provides the appropriate level of QoS. The notion of *equivalent bandwidth* [6] plays a key role in bridging these two levels. The equivalent bandwidth of a connection takes the QoS requirements at the cell level and quantifies the bandwidth resource requirements of a connection. This provides a means to assess the ability of a network element to accommodate a connection. Equivalent bandwidth also provides a means to quantify the resources allocated to groups of users.

A *virtual network* is the subset of network resources allocated to a class of users. Requests for VC or VP connectivity, or changes in resource allocation are managed as a set entirely within the virtual network. The virtual network concept simplifies resource management as only a subset of network resources needs to be managed. The TP layer can be viewed as spawning multiple virtual networks, each with its own management system. Virtual networks can be dedicated according to QoS/service/traffic class or different customer groups. Each virtual network can be provisioned and operated to support different requirements. Virtual networks can be spawned from other virtual networks by suballocating resources, thus enabling scalability and software reuse.

The CITR project developed an NRM architecture, shown in Fig. 2. The architecture was chosen to align with TINA functional layering and TMN's network and element management layers. The managing systems include a connection session manager, a set of layer network managers, and a set of element managers. Interactions between these managing entities consist of invocations, responses, and notifications. A small-scale prototype of the NRM management system was implemented and is described in [7]. The GENESIS project at Columbia University and the TEMPEST project at Cambridge University [8] also employ virtual networks for resource management. These projects differ from ours in regard to the definition of the virtual network as a collection of virtual paths as opposed to a connected set of virtual links. The importance of this distinction is discussed next.

Packet-Level Resource Management Using Virtual Networks

Currently there are three broad classes of approaches to QoS, which we refer to as *overlay application layer virtual networks*, *market managed networks*, and *virtual QoS networks* (VQNs). The overlay approach attempts to provide improved QoS performance by building application-specific overlay virtual networks on the existing best effort infrastructure and use QoS routing or content-based routing to improve service quality and/or efficiency. The market managed and VQN approaches are more fundamental in that they impact deeper layers. These latter two classes differ in respect to the degree of dynamic coupling between charging and QoS provisioning. We focus here on VQNs.

Virtual QoS networks are dimensioned to provide a desired grade of service (GoS), usually expressed as a connection or session blocking probability based on forecast demand for each network-specified QoS category, where the QoS is quantified in terms of packet level performance metrics such as packet loss probability, packet delay, or delay variation. In this case one still charges for quality, but the charge is relatively static compared to market managed networks. A feature of this class of approaches is that tariffs are set to recover embedded costs and related to the "value of service" to the user. Aggregate network capacity is provisioned to handle busy hour traffic load and satisfy packet level performance metrics within each QoS category. To accommodate variations of traffic mix across service categories, the network capacity is partitioned into a set of QoS virtual networks, VNs, where the allocated capacity to a given VN is adaptively adjusted to satisfy QoS and fairness requirements for the expected near-term traffic demand. The proposed approaches in this class differ in regard to the granularity of traffic flows being managed, ranging from per-flow-aware schemes to aggregate path flow approaches to aggregate link flow schemes per QoS class.

The similarity between the VC/VP concepts of ATM and the nested LSP construct of MPLS suggests that many of the techniques developed for ATM resource management and network survivability should find analogs in MPLS environments. In particular, the *virtual network concept* employed in our multilevel resource management scheme for ATM networks [4, 5, 9] is very naturally extended to MPLS networks where the label stacking functionality provides the means of building virtual links in the VNs that may be defined for either private users, (virtual private networks, VPNs), or different QoS or CoS classes (VQNs). Below we propose an architecture for MPLS networks. This architecture is similar to those investigated in [10, 11], but differs in regard to the way in which bandwidth is managed.

The architecture supports both connectionless IP frames suitable for short duration flows as well as LSPs for longer duration flows and integrates this traffic efficiently while meeting QoS requirements. We refer the reader to [4, 5, 9], for a detailed description in the ATM case. Here we concentrate on the extension of this architecture to MPLS networks.

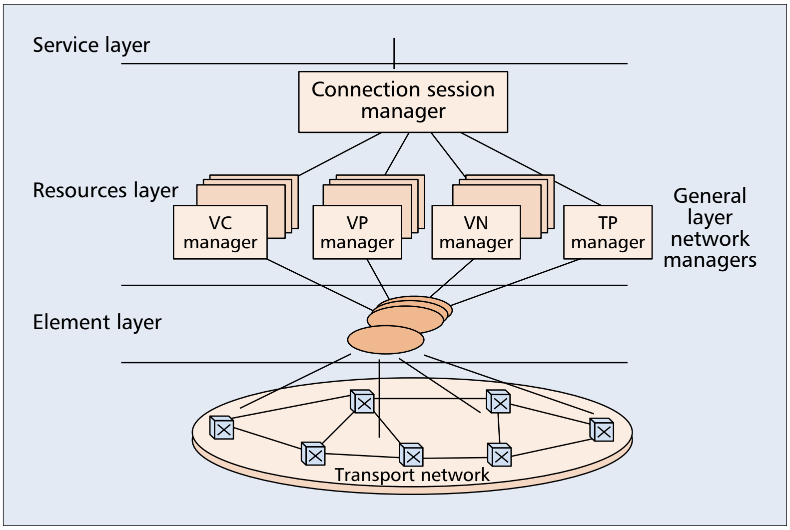

Figure 2. *NRM architecture.*

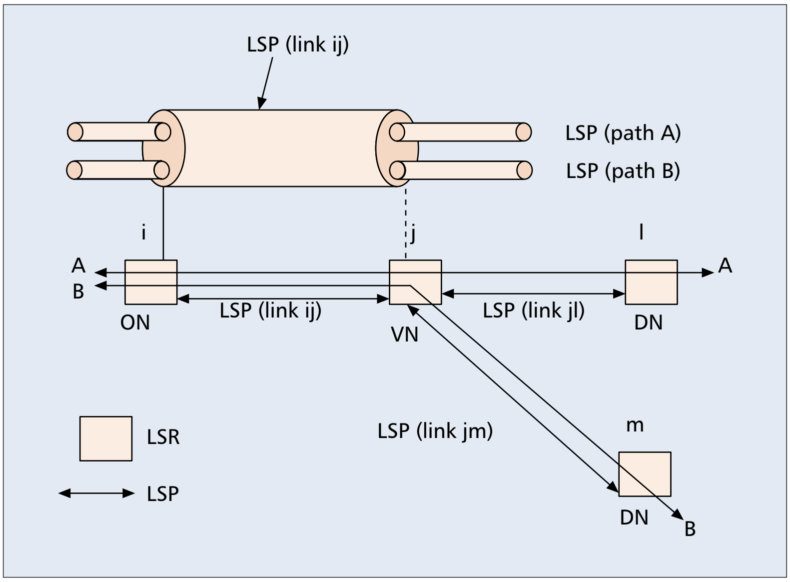

Figure 3. *Path and link label switched paths.*

We first illustrate how label stacking is used to define a multilevel hierarchy of microflows, path flows, and link flows. The path LSPs can comprise a set of microflows between the same endpoints. In Fig. 3 we illustrate the relationship between path LSPs and link LSPs, where several end-to-end or path LSPs of the same QoS class can be contained within a given link LSP for that QoS class. The label associated with the path LSP is used by the label switched router (LSR) for routing. The label associated with an inbound link LSP must be removed (*popped* in MPLS label stacking terminology) by the LSR to unbundle its constituent path LPSs before label swapping. Following the label swapping necessary to route the path LSP, the appropriate label for the outgoing link LSP is added (pushed) on the label stack. These operations at

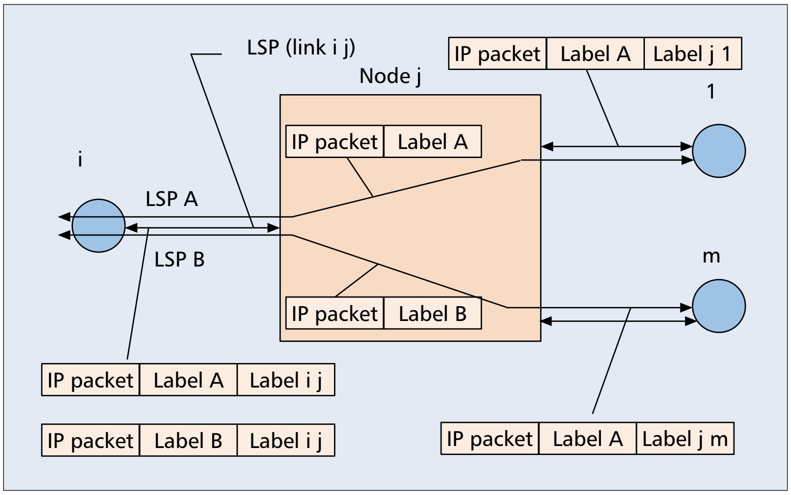

Figure 4. *Label switched router implementation of link level bandwidth management.*

a via node LSR, necessary to implement our link level bandwidth management approach, are depicted in Fig. 4. Note that bandwidth is not assigned to individual end-to-end LSPs, only to the link LSPs. A tally is kept of the aggregated flow assigned to the links of the CoS virtual networks to control flow admission for flow setup requests. Alternatively, the aggregate effective bandwidth can be measured online [9] for tighter control, while reducing the burden of the user in defining the traffic characteristics of the setup request.

At this point we wish to emphasize that our approach to virtual network resource management can be applied to any switching technique, queue management, and routing scheme. Of course the dimensioning of the VQNs depends on the specific multiplexing, scheduling, and routing procedures actually used. In the following we illustrate how CoS VQNs are defined in the specific case of a DiffServ/MPLS domain.

Figure 5 from [11] illustrates how the bandwidth is managed to implement DiffServ at an LSR node. There are three CoSs in the model considered, corresponding to the DiffServ categories:

- Expedited forwarding with the highest queuing priority
- Assured forwarding with second queuing priority
- Best effort with lowest priority

By applying the routing control available with the MPLS end-to-end LSPs, as well as access control, through either CAC for reserved flows or traffic conditioning for connectionless flows, we can offer QoS to admitted flows, across an MPLS domain by trading off access blocking with packet, loss/delay. In the model investigated by Ash [11] bandwidth reservation gives preference to the preferred traffic by allowing it to seize any idle bandwidth in a link, while nonpreferred traffic is only allowed to seize bandwidth if there is a minimum of idle bandwidth available. The minimum bandwidth threshold is called the *reservation level*. Bandwidth reservation can be static or dynamic.

Bandwidth management, can be performed at three distinct levels of granularity, namely:

1) Per flow (or connection demand) as investigated by [10, 11]

2) Per aggregated bandwidth demand , or bandwidth pipe linking origin nodes (ONs) and destination nodes (DNs) within the same QoS class. The collection of such bandwidth pipes for all ON-DN pairs for a given QoS class is termed a VNET for that class in [11].

3) Per (QoS virtual network) VQN = set of virtual link LSPs as proposed in [5] and investigated further in [12]. Note that this differs from the VNET described in [11] in that bandwidth is

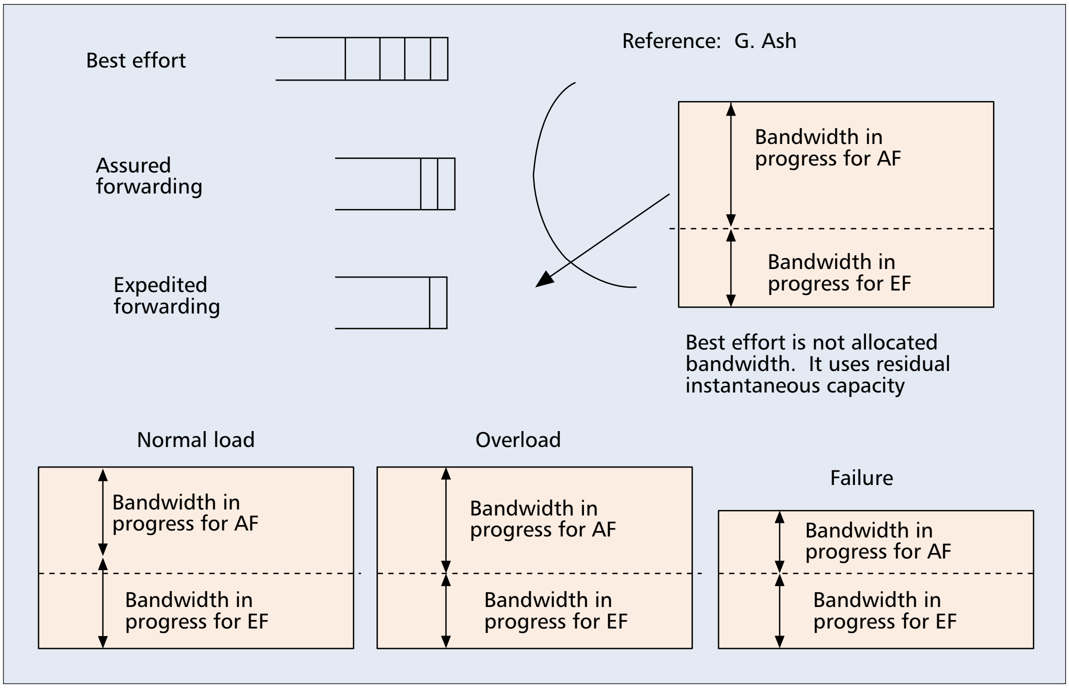

Figure 5. *Bandwidth management for DiffServ and MPLS CAC.*

managed on a virtual link basis as opposed to an end-to-end path or *bandwidth pipe*, which may contain one or more concatenated virtual links. The advantage of introducing these link LSPs is that it allows statistical multiplexing over the aggregate flow on a link for a specific QoS class, thereby achieving multiplexing gains relative to the bandwidth pipe approach.

In order to ensure fair efficient access for heterogeneous effective bandwidth requests to virtual link bandwidth, two alternatives have been investigated. Virtual link dynamic bandwidth reservation techniques described in [13] can be employed for any path routing and CAC method. Several bandwidth sharing techniques were investigated within a game theoretic framework. Dynamic trunk reservation achieves near-optimal performance while being less complex than the exact admission policy derived from Markov decision theory. As pointed out in [11], bandwidth reservation is also used in nonhierarchical routing schemes to prevent instability. In the case of revenue maximum routing [12], the call class revenue parameters can be used to ensure fair efficient access without the need for explicit link level admission policies.

CAC is performed at ONs for arriving connection requests on end-to-end LSPs by successively checking whether all virtual links on the path traversed by the end-to-end LSP in the requested VQN can support the request. Crankback is done when a request cannot be met, and the ON selects an alternate path to attempt connection. The process repeats until a path is found or the connection request is denied. Note that this setup procedure is decentralized and is a control plane function involving signaling. The bandwidth allocation to virtual links, on the other hand, need not be a control plane function and should be operated at a slower timescale. Virtual link bandwidth allocation can be either centrally managed or decentralized.

Ash [11] reports on the merits of the VNET approach (path-based bandwidth management in our terminology) compared to managing microflows. Preliminary studies reported in [12] compare path and link bandwidth management for ATM networks. It is shown that managing bandwidth on a virtual link basis compares favorably to path bandwidth management in terms of the required transmission capacity for several network design examples.

In all three approaches discussed above the routing of connections is determined at setup time, and remains fixed for the duration of the connection. Routing and bandwidth management techniques developed for multirate circuit-switched networks can be used where the effective bandwidth is employed in the allocation of available bandwidth to the requested connection in place of real bandwidth. Routing algorithms may be fixed, or time-, state-, or event-dependent. We are focusing on event-dependent algorithms with learning because of the low signaling overhead and adaptive capability. Reference [11] has shown the merits of EDR in the context of a path-managed architecture.

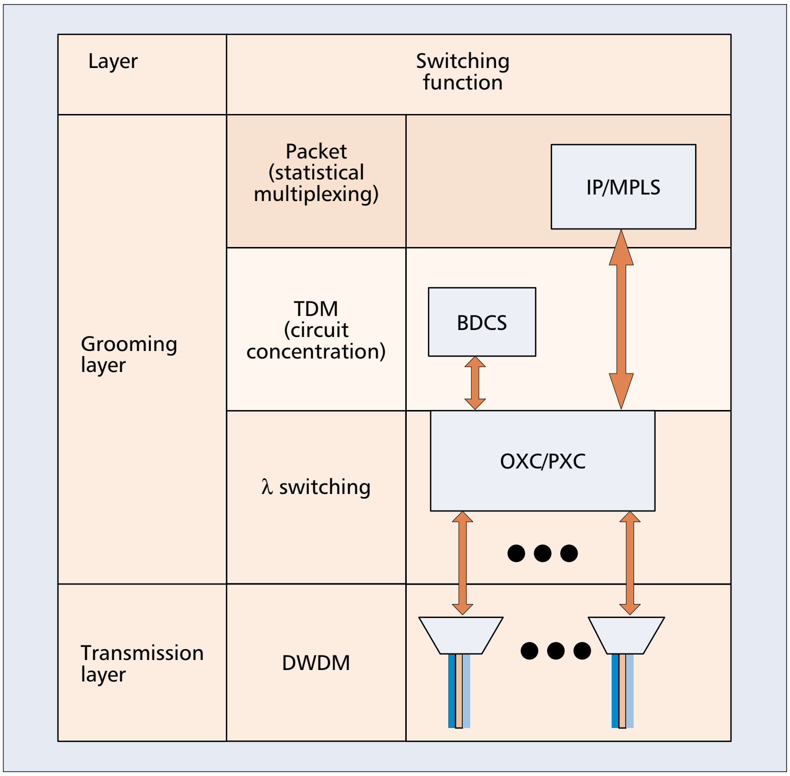

Figure 6. *Multilayer transport network management.*

Multilayer Transport Networks: GMPLS

The last decade has seen a tremendous increase in the scale of transport networks. At the top layer, the volume of packet traffic has followed the explosive growth of the Internet. At the lower layer the large-scale deployment of optical fiber transmission technology has resulted not only in new fiber connectivity, but also in vast wavelength connectivity from dense wavelength-division multiplexing (DWDM) gear. Present transport networks incorporate most of the layers shown in Fig. 6, which include connectivity at the fiber, optical waveband, individual wavelengths, SONET, ATM, and IP/MPLS. These multilayer networks pose new challenges in control of connections and management of network resources.

Generalized MPLS (GMPLS) is under development to address the need for generic connection control and resource management protocols. The proliferation of DWDM to gain capacity in optical fiber transmission led to the need to control the interconnection of wavelength paths, or lambdas, within an optical network. The MPLS notion of a label is implicitly carried by the color or wavelength, so an initiative to extend the MPLS protocols to wavelength connections was initially dubbed multiprotocol lambda switching (MPλS). The initiative recognized that MPLS provided for connection control regardless of whether the connection involved circuit or packet switching.

It was soon realized that having a unified control plane for the optical circuit-switched network and the packet-switched network would

The virtual network concept can play a useful role in the management and control of this new generation of multilayer networks. GMPLS extends the label stacking capability of MPLS to span the range from packet to fiber connectivity.

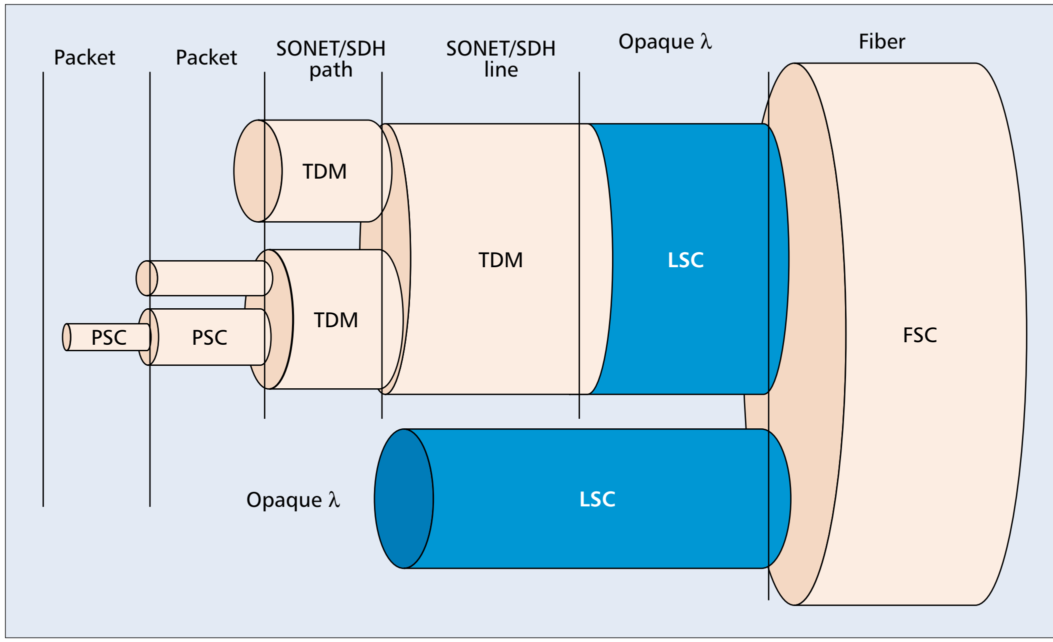

Figure 7. *Label stacking in GMPLS.*

unify and simplify management and resource control. Otherwise, MPLS would be effectively an overlay on MPλS, with duplication of signaling and routing functions making coordination of the two layers more difficult. Incidentally, this development was reminiscent of the attempts to overlay IP networks on ATM, which led to the initial MPLS proposal. Subsequently GMPLS has been proposed as a unified control plane for all switching categories including TDM circuit switching, packet switching, and optical switching. This unified control plane facilitates integration of management and control functions with a potential for more efficient resource utilization and improved network survivability while keeping signaling overhead to a minimum.

The potential for integrated multilevel resource management has been around since the early days of circuit switching and rearrangeable transmission networks, employing crossconnects or slow switches. It arose again in the context of ATM in the form of VC and VP layers, as well as in ATM/SDH. Current IP traffic and the availability of DWDM has given rise to the existing hierarchy IP/ATM/SDH/DWDM. Here, ATM provides functionality for resource management, while SDH provides for fast restoration capability, and finally DWDM provides the speed. This multilayered architecture is undesirable as there is duplication of functionality at several layers, which results in increased cost and performance degradation. With the advent of MPLS and GMPLS the optical Internet architecture is evolving toward IP/GMPLS/DWDM that will reduce this redundancy and improve performance at reduced cost.

The virtual network concept can play a useful role in the management and control of this new generation of multilayer networks. GMPLS extends the label stacking capability of MPLS to span the range from packet to fiber connectivity, as shown in Fig. 7. The client-server relationship between layers as well as within layers continues to govern the allocation of resources. Clearly, the notion of a virtual network captures quite accurately the relationship between layers, and its use can simplify network management by constraining the scope of resources that need to be managed in a given setting. The recursive use of virtual networks to describe sets of network resources that are managed together to provide connectivity services for client layers provides a powerful tool that can give a high degree of scalability.

Open Problems and Issues

There are a number of questions that need further study in connection with our virtual network architecture. The dimensioning of the VQNs depends strongly on the scheduling and routing method used. So far we have considered class-based priority queuing and two different routing algorithms: Dynamically Controlled Routing (DCR) and Revenue Maximum routing. Other event-dependent approaches are under study with reduced signaling overhead. Apart from this we note that any of the reported scheduling and QoS routing methods can be applied within the VN framework, and such studies and comparisons should be undertaken.

Our previous work on the ATM resource management architecture considered only a single domain, where the traffic demands for different service classes were endogenous quantities. For this single domain design, static service level agreements (SLAs) are implicit in the QoS and GoS parameters associated with the traffic demands offered to each service class. Reference [14] proposed a VQN dimensioning model that explicitly includes SLAs for a single domain network. In [15] a robust revenue-based bandwidth management scheme has been studied for adapting to deviations from the designed traffic demand. While the SLAs are crafted offline, implying static operation, it would appear that this general framework could be extended to include dynamic SLAs, such as those employed in market managed networks. A very important area requiring investigation involves multidomain networks, where several MPLS islands interwork via SLAs. End-to-end QoS depends critically on such an extension. It would be desirable to compare an optimized VQN approach including dynamic SLAs with existing proposals for market managed networks.

It is anticipated that future networks will involve packet services supported over a flexible circuit-switched all-optical core. GMPLS and the VN concept offer an opportunity to effectively coordinate resource management in these multilayer networks. We would anticipate that network performance would benefit from multilayer coordination and such studies should be undertaken.

The virtual network concept is a useful abstraction for network resource management. We have proposed and studied a multilevel resource management architecture for ATM employing the virtual network concept. In this article we have shown how this approach can be profitably extended to future networks employing MPLS and GMPLS. The versatility of the virtual network construct is another example of how traffic engineering results can transcend the technology for which they were initially conceived.

Acknowledgments

The research reported in this article is an extension of the research carried out in the Canadian Institute for Telecommunications Research (CITR) major project on Broadband Network Architecture and Resource Management. Researchers from the University of Alberta, INRS-Telecommunications, University of Toronto, and University of Waterloo have participated in this program. The topics of research included a reference model for network resource management, virtual network configuration management, ABR services management and cell-level congestion control, and multipoint call routing and virtual path design. The industrial partners were Aliant Inc., Alcatel Canada, Bell Canada, Miranda Media Networks, and Nortel Networks.

This research was supported in part by a grant from CITR under the Networks of Centres of Excellence (NCE) program of the Government of Canada and the industrial partners.

References

[1] K. Nichols, V. Jacobson, and L. Zhang, "A Two-Bit Differentiated Services Architecture for the Internet," ftp://ftp.ee.lbl.gov/papers/dsarch.pdf, Nov. 1997.

[2] G. Armitage, "MPLS: The Magic behind the Myths," *IEEE Commun. Mag.*, vol. 38, no. 1, Jan. 2000.

[3] D. O. Awduche *et al.*, "Overview and Principles of Internet Traffic Engineering," IETF RFC 3272, May, 2002.

[4] G. Woodruff *et al.*, "ATM Network Resources Management Using Layer and Virtual Network Concepts," *5th IFIPIEEE Int'l. Symp. Integrated Net. Mgmt.*, San Diego, CA, May 1997, pp 314–26.

[5] Z. Dziong, Y. Xiong, and L. Mason, "Virtual Network Concept and its application for Resource Management in ATM Based Networks," *Proc. IFIP/IEEE Broadband Commun. '96*, 1996, pp. 223–34.

[6] F. P. Kelly, "Notes on Effective bandwidth," *Stochastic Networks: Theory and Applications*, F. P. Kelly, S. Zachary, and I. B. Ziedens, Eds., Oxford Univ. Press, 1996, pp. 141–68.

[7] W. F. Ng *et al.*, "MIBlets: A Practical Approach to Virtual Network Management," *Integrated Net. Mgmt.*, 1999.

[8] S. Rooney *et al.*, "The Tempest: A Framework for Safe, Resource-Assured, Programmable Networks," *IEEE Commun. Mag.*, vol. 36, no. 10, Oct 1998, pp. 42–53.

[9] Z. Dziong, *ATM Resource Management*, McGraw Hill, 1997.

[10] J. Roberts and S. Queslati-Boulahia, "Quality of Service by Flow-Aware Networking," *Phil. Trans. R. Soc. London A. 2000*, pp. 2197–07

[11] G. R. Ash, "Performance Evaluation of QoS-Routing Methods for IP-Based Multi-Service Networks," *Comp. Commun.*, to be published; http://www.research.att.com/~jrex/jerry/

[12] Z. Dziong, J. Zhang J., and L. G. Mason, "Virtual Network vs. Virtual Path Design for ATM Networks," *Proc. 15th Int'l. Teletraffic Cong.*, Washington, DC, 1997.

[13] Z. Dziong and L. G. Mason, "Fair-Efficient Call Admission for Broadband Networks — A Game Theoretic Framework," vol. 4, no. 1, *IEEE/ACM Trans. Net.*, Feb. 1996.

[14] D. Mitra, J. A. Morrison, and K. G. Ramakrishnan, "Virtual Private Networks: Joint Resource Allocation and Routing Design," *Proc. IEEE INFOCOM '99*, pp. 480–90.

[15] E. Bouillet, D. Mitra, and K. G. Ramakrishnan, "The Structure and Management of Service Level Agreements in Networks," *IEEE JSAC*, vol. 20, no. 4, May 2002.

Biographies

ALBERTO LEON-GARCIA [F] was co-founder and CTO of AcceLight Networks in Ottawa, Canada, which developed leading edge optical switching equipment. He holds the Jeffrey Skoll Chair in Computer Networks and Innovation, and is co-director of the Master of Engineering in Telecommunications program which he founded in 1997. He is author of *Probability and Random Processes for Electrical Engineering* and co-author with Indra Widjaja of *Communication Networks: Fundamental Concepts and Key Architecture*.

LORNE G. MASON [SM] (mason@ece.mcgill.ca) is a visiting professor of electrical and computer engineering at McGill University, Montreal, Quebec, Canada, and Professor Honouraire at INRS-Telecommunications (a research institute affiliated with the Université du Québec). He was co-recipient of the 1993 STENTOR Award for collaborative research in telecommunications for his contributions to state-dependent routing. He chaired the joint IFIP/IEEE conference Broadband Communications '96 and was technical co-chairman for the IEEE ATM '98 workshop. He is the Canadian representative to the IFIP TC6 committee.

It is anticipated that future networks will involve packet services supported over a flexible circuit switched all-optical core. GMPLS and the VN concept offer an opportunity to effectively coordinate resource management in these multi-layer networks.