ELSEVIER

Computer Communications 25 (2002) 100–108

computer communications

www.elsevier.com/locate/comcom

Review

A review of quality of service mechanisms in IP-based networks — integrated and differentiated services, multi-layer switching, MPLS and traffic engineering

Ray Hunt*

Department of Computer Science, University of Canterbury, Private Bag 4800, Christchurch, New Zealand

Received 30 August 2000; revised 23 March 2001; accepted 23 March 2001

Abstract

ISPs are facing the challenge of offering improved quality of service (QoS) to their customers. No longer is best effort delivery with no service guarantee acceptable for many applications. Although ATM has provided a limited solution by way of service classes, such a solution pre-supposed an underlying ATM network which — in the case of pure IP traffic — may not be the case. Rather IP traffic requires a degree of engineering into service classes (differentiated services, DiffServ) as well as a break from traditional layer three-based routing. Although access to virtually unlimited bandwidth via WDM and Photonic Networks may potentially offer a solution to the QoS issue, access to such services on a universal basis is not a services class paradigm and using a label switching technique is seen as an appropriate medium term solution. Further, label switching offers a simple and efficient mechanism for IP traffic engineering, multi-service functionality and scalability. This paper examines a number of service classifications and solutions, which aim to provide a realistic QoS solution. In particular it addresses Integrated and DiffServ, multi-layer switching and MPLS, which forms the basis of DiffServ as it allows ISPs to deliver new services not easily supportable by conventional IP routing infrastructure. Finally the paper makes some important observations about traffic engineering.

Keywords: QoS (quality of service); Differentiated services (DiffServ); Integrated services (IntServ); BM (bandwidth manager); Multi-later switching; MPLS (multi-protocol label switching); Traffic engineering

1. Introduction

IP-based networks only provide Best Effort Service, which implies that there is no guarantee as to delay margins or actual delivery times. The problem with today's generic IP is that it only provides point-to-point connectivity, operates on a first-come-first-served basis, provides only best effort services and is subject to variable queuing delays and congestion losses. Neither is it possible to share bandwidth on a particular link between applications with different performance requirements.

The IETF has proposed many service models and mechanisms to meet the demand for QoS. These proposals include: an integrated services (IntServ)/RSVP model [1,2], differentiated services (DiffServ) model [3], multi-protocol label switching (MPLS) [4,5], and Traffic Engineering [5]. Further, a variety of service classes have been proposed which offer low delay and low jitter for applications such as Internet telephony and video conferencing as well as a range of services for guaranteed (real-time), controlled load (premium) and best effort services.

2. Integrated services/RSVP

RSVP was originally designed as a signalling protocol for applications to reserve network resources [2]. It represents a fundamental change to existing Internet architecture where all flow-based state information exists in the end systems. The Integrated Services/RSVP (IntServ) model comprises three classes of service [6]:

- best effort — for time delay independent applications,
- guaranteed — for applications requiring fixed delays,
- predictive — for applications requiring probabilistic delays.

Routers are required to reserve resources to provide QoS for specified flows. When a host application needs to transmit real-time data requiring a specific QoS level (e.g. a

* Tel.: +64-3-3642347; fax: +64-3-3642569.
E-mail address: ray@cosc.canterbury.ac.nz (R. Hunt).

0140-3664/02/$ - see front matter
PII: S0140-3664(02)00334-6

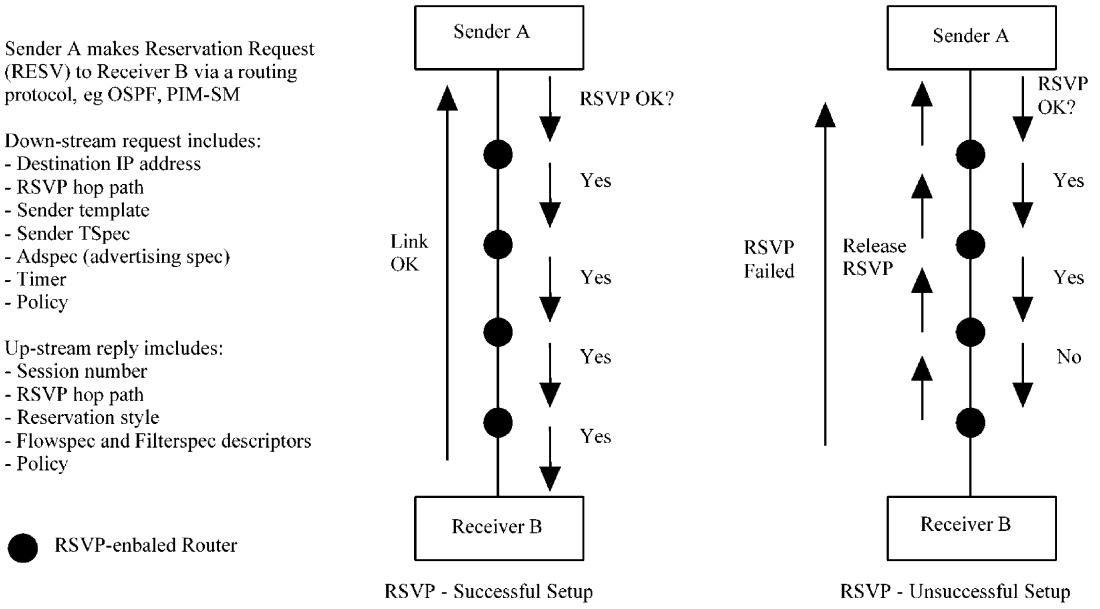

Fig. 1. Examples of successful and unsuccessful RSVP setups.

guaranteed bandwidth) a host application sends an Reservation Request (RESV) to the receiver via routing protocols such as OSPF, PIM-SM, etc. Fig. 1 shows examples of successful and unsuccessful RSVP setups.

At each node on the path several decisions are made. Firstly, the node must make sure that the requested bandwidth is available (admission control). The node must then check with its policy control module to ensure that the receiver has enough rights to request the specified level of service. Also at each node a check is made to see if the flows can be merged with those from other nodes. When the receiver has successfully reserved resources over the entire path a success message is returned. If a node rejects a reservation, the request is denied and the resources already reserved at intermediate nodes are released.

Although an interesting solution for the QoS issue, RSVP suffers from scalability problems as information for individual traffic flows must be stored throughout the network. This results in difficult management issues at times of congestion rerouting, etc. More recently RSVP has been extended to reserve resources for aggregate flows (Sections 4 and 5.4) and to setup explicit routes with QoS perimeters for network signalling [7,8].

3. Differentiated services

3.1. The differentiated services (DiffServ) framework

The DiffServ (DS) framework provides a methodology for offering a range of IntServ without the requirement for the substantial overhead needed for per-flow state information in every router as is the case with the IntServ Model. Potentially DiffServ has been available in IPv4 by way of the type of service (TOS) field but todate it has been rarely used.

DiffServ defines a set of packet forwarding criteria — per hop behaviour (PHB) [9]. Packets are handled based upon the DS field and therefore a variety of classes can be defined thus creating a priority scheme. However individual flows within a DiffServ class cannot be differentiated.

By the use of classification, policing, shaping and scheduling, a variety of services can be provided [10–12].

Table 1
Comparison of Inter Serv and Diff Serv architecure

	IntServ	DiffServ
Number of new service classes	2	Limited by DS field
State information	Proportional to no. of flows	Proportional to no. of service classes
Scalability	No	Yes
Deployment and implementation	Difficult	Easy [a] can be incrementally deployed [b]

[a] Classification, marking, policing and shaping operations are only needed at the boundary of the networks and ISP core routers need only to implement behaviour aggregate (BA) Classification (the process of sorting packets based upon their DS fields).

[b] Incremental deployment is possible for assured services. Routers which cannot handle the DS field fall back to providing best effort service.

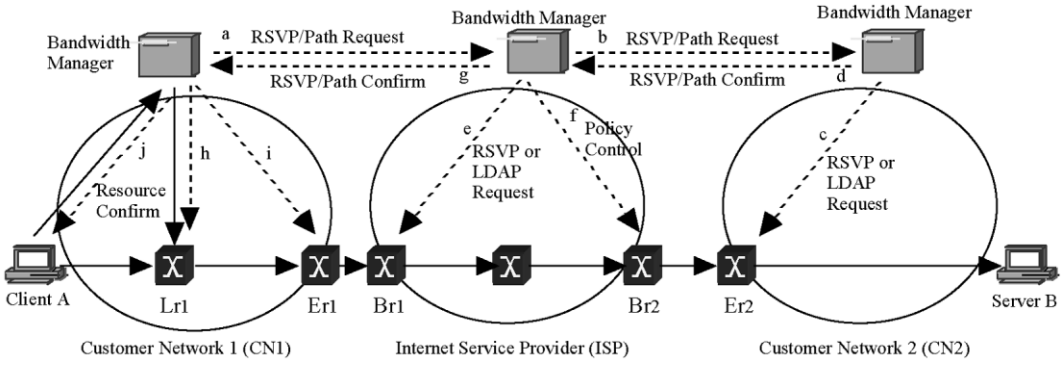

Fig. 2. Establishment of service between client A and server B using a BM and RSVP signalling [17].

However, DiffServ only defines the DiffServ and per-hop behaviour fields. It is up to the implementor (e.g. ISP) to implement appropriate handling mechanisms. Services currently being defined include:

- Assured services — for applications requiring better reliability than best effort service.
- Premium services — for applications requiring low delay and low jitter.

DiffServ provides a very different service offering than IntServ which can be seen from Table 1.

3.2. Assured differentiated service

Assured Service [11,12] is designed for customers who require an improved QoS over best effort as well as performance parameters to be met by the ISP.[1] Assured service resembles ATM's ABR or VBR services. Such an service level agreement (SLA) will allocate bandwidth but applications must share this bandwidth in accordance with their own policy. The ISP's ingress router performs classification and policing. If the traffic rate does not exceed the SLA bit rate then it is said to be *in* profile. Excess packets are *out* profile and are handled by random early detection (RED) and RIO (RED with *in* and *out*) queue management discipline [13]. RED drops packets randomly while RIO forms two RED queues — one for *in* packets and one for *out* packets. There are two thresholds for each queue in this case. When the queue size is below the first threshold no packets are dropped but when the queue size is between the two thresholds only *out* packets are randomly dropped. When the queue size exceeds the second threshold indicating possible network congestion then both *in* and *out* packets are randomly dropped, but *out* packets are dropped more aggressively. The DS field contains an A-bit which distinguishes the *in* (A bit = 1) and *out* (A bit = 0) flows.

[1] This is not unsimilar to a Frame Relay service which provides a user selected CIR with best effort based upon credit allocation beyond the CIR level and up to another user selected level — line bit level.

3.3. Premium differentiated service

For applications which require a specific maximum or average bit rate then a Premium Service is required. This is a low delay, low jitter service. Traffic rates in excess of this SLA will result in packet discard. This service is appropriate for voice over IP, video conferencing and certain VPNs [14] and resembles ATM's CBR service.

For Premium Service it is necessary to support both static and dynamic SLAs in order for customers to request a different service level on the fly without having first subscribed to them although some admission control mechanism is needed. When Premium Service traffic arrives (P-bit set) traffic may need to be reshaped before it leaves the customer's network to ensure that it conforms to the SLA profile.

Various schemes have been proposed to ensure a fair and even balance between Premium and Assured traffic flows. Examples include:

- Control Premium to Assured traffic flow to a specified ratio, e.g. 20%.
- Packet rates in excess of the SLA can be discarded at the network ingress.
- Implement a weighted fair queuing [15] scheme between Premium and Assured queues.

The Premium Service queue should normally be empty or at least very short thus ensuring low delay and jitter. However uneven traffic flows can cause a problem for Premium Services.

4. Service and resource allocation using bandwidth management

Although an SLA establishes an agreement with an ISP, it is still necessary for an intranet service customer to decide upon how such resources are shared. Although individual hosts can make arbitrary decisions, a more intelligent scheme establishes a bandwidth manager (BM) [16] (to make decisions for all hosts based upon management policies). This device can simply be software running on a router or even a dedicated host. For Premium traffic,

which commonly requires a dynamic SLA, the BM would use protocols such as RSVP and/or LDAP to establish classification, shaping policies, etc. at the boundary router as well as with corresponding BMs in ISP(s) and the destination intranet.

Within the ISP's domain, static SLAs cause no problems. Routers can be configured with classification and policy-based shaping rules once at the beginning of the flow. Such static allocation — although inflexible for the customer — permits easy avoidance of congestion. Dynamic SLAs require careful resource allocation with the BMs via RSVP signaling, which is illustrated by way of the example shown in Fig. 2. A service is established between Client A and Server B each connected to private customer networks and interconnected by transit ISP network(s).

The BM in CN1 sends an RSVP/PATH message to the ISP's BM which makes admission control decisions and (if successful) sends a further RSVP/PATH message to CN2's BM. If successful, CN2's BM will set classification and policy rules for router Er_2 using RSVP and/or LDAP services. CN2's BM also sends a confirmation to the ISP's BM via an RSVP/RESV message.

On receipt of this message the ISP's BM sets the classification and policy rules on router Br_1 as well as policing and reshaping rules on router Br_2. Finally an RSVP/RESV confirmation is returned to CN1's BM. At this stage CN1's BM knows that the resources have been allocated and it then sets the classification and shaping rules on router Lr_1. If the traffic received by Lr_1 is not conformant it can be reshaped. CN1's BM will also set the policy and reshaping rules for Er_1. Finally once this is complete Client A receives an RSVP/RESV message from CN1's BM and it can commence transmission. If at any stage any of the BMs reject an RSVP/PATH request, then Client A is advised accordingly.

Although this IntServ/RSVP mechanism is similar in principal to that described in Section 2, there are four main differences:

- It is the sender that requests resources not the receiver.
- A request can be rejected when the BM receives the PATH message from the sender. In IntServ a request is rejected only when a router receives the RESV message from the receiver.
- A BM can aggregate multiple requests and make a single request to the next BM.
- Each domain behaves like a single node represented by the BM. ISP core routers are not involved in this process.

Where traffic traverses an IEEE 802 customer LAN QoS mechanisms will have to be in place if end-to-end QoS-based SLAs can be guaranteed. For example if CN1 in Fig. 2 is a traditional Ethernet LAN, then IEEE 802.1p will have to be implemented to provide classes of service across this LAN. IEEE 802.1p defines a Subnet Bandwidth Manager (SBM) which accepts RESVs from hosts and routers and updates IEEE 802.1p end-points. In this protocol a 3-bit VLAN tag field is used to classify/prioritise frames as they traverse this switched LAN [18–20]. This proposal defines a signalling protocol for LAN-based admission control for RSVP flows. The intention is that this mechanism — when combined with policing at the end systems as well as traffic control and priority queuing at the network layer — will provide a close approximation to premium/assured differentiated service flows. It is essential to separate flows under RSVP control from best-effort flows and the SBM architecture provides a mechanism to do this via the IEEE 802.1p priority system.

The SBM might be implemented on an intelligent LAN switch, which supports the signalling protocol mechanisms consistent with IEEE 802.1p. The SBM is configured with information about the maximum bandwidth that can be reserved on each segment under its control. Although this information can be gained statistically, dynamic discovery methods will likely be used in the future.

5. Multi-layer switching

5.1. Evolution of multi-layer switching in IP-based networks

In order to achieve a well-engineered IP network that can provide the flow requirements for the DiffServ model described in Section 3, the conventional IP routing architecture had to change.

Multi-layer switching specifies an integration of layer 2 switching with layer 3 routing. Networks started to be constructed using an *overlay* model in which a logical IP router topology operates over and is independent of an underlying layer 2 switching technology such as Frame Relay or ATM. There were however complexities in operating this model. For example PVCs between routers had to be manually configured. Further, use of SVCs mandated the resolution of IP to ATM addresses. Although this approach derives the benefits of both layer 2 and layer 3 architectures difficulties arose in the complexity of mapping between two separate topologies, address spaces, routing protocols, signalling protocols and resource allocation systems.

Further evolution occurred to the *peer* model in which integrated Switch/Routers maintained single IP addressing space and ran a single IP routing protocol — just like a network of routers. Some work was required to map IP traffic to layer-2 switched path via IP switching control protocols. This work resulted in the evolution of multi-layer switching solutions. In particular MPLS represents an important effort designed to decrease the complexity of combining layer 2 switching and layer 3 routing into an integrated system. This is further discussed in Section 5.3.

5.2. Forwarding and control mechanisms

Multi-layer switching solutions are characterised by two components: control and forwarding; and label swapping.

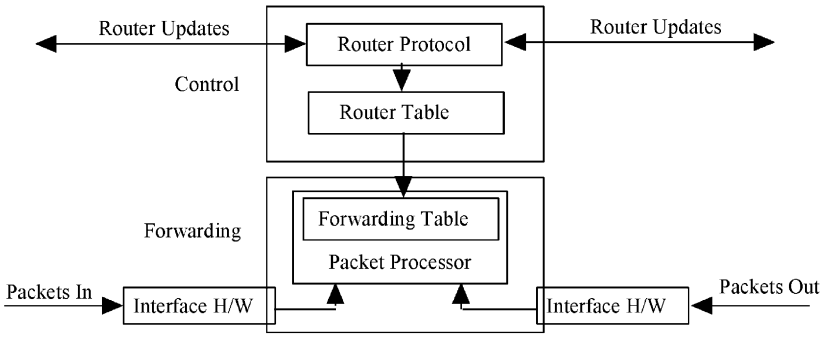

Fig. 3. Control and forwarding routing functionality.

5.2.1. Control and forwarding

The control and forwarding components are common to all switching methodologies (including MPLS) as shown in Fig. 3. The control component uses routing protocols such as OSPF, IS–IS and BGP4 to exchange control information and maintain forwarding tables with its neighbours. This forwarding table provides information necessary for a routing decision thus forming a switched path between the input and output ports. The control components are separated from the forwarding component and thus each can be modified independently of the other.

5.2.2. Label swapping

The forwarding component of virtually all multi-layer switching solutions is based upon a label-swapping forwarding algorithm. This is the same algorithm used to forward data in ATM and Frame Relay networks. Signalling and label distribution are fundamental to the operation of a label swapping forwarding algorithm. A label is a short fixed-length value carried in the packet's header and is used to identify a forwarding equivalent class (FEC). A label is similar to a connection identifier such as that used in ATM (VPI/VCI) or in Frame Relay (DLCI), as it has only local significance and maps traffic to a specific FEC. FEC represents a set of packets that are forwarded over the same path even if their ultimate destinations are different.

Label swapping forwarding algorithms require that a packet classification occur at the network entry point and that an initial label be assigned to every packet. In Fig. 4 it can be seen that the entry label switch receives an unlabelled packet with a destination address of 204.137.98.1. The label switch performs a longest-match routing table lookup and then maps the packet to an FEC-204.137.98/24. The ingress label switch then assigns a label with a value of seven to the packet and forwards it to the next hop in the label switching path (LSP).

This LSP path is equivalent to a virtual circuit as it defines entry to exit points through the network and all packets follow this path. Within the network, label switches ignore a packet's network layer header and forward the packet using the label-swapping algorithm. At the exit point from the network the forwarding component searches its forwarding table and if the next hop is not a label switch then the exit switch discards the label and forwards the packet using conventional longest-match IP forwarding.

5.3. Multi-protocol label switching (MPLS)

MPLS [21,22] is a forwarding scheme which primarily evolved from Cisco's *Tag Switching* [23]. The motivation for MPLS is to use a fixed-length label to decide upon packet handling. It is also a useful tool for Traffic Engineering [5]. MPLS is the latest step in the evolution of multi-layer switching for IP-based networks. It is an IETF standards-based approach built on the efforts of the various proprietary multi-layer switching solutions.

MPLS uses the control-driven model to initiate the assignment and distribution of label bindings for the establishment of LSPs. An LSP is created by concatenating one or more label switched hops, allowing a packet to be forwarded from one label-switching router (LSR) to another LSR across the MPLS domain. An LSR is a router that

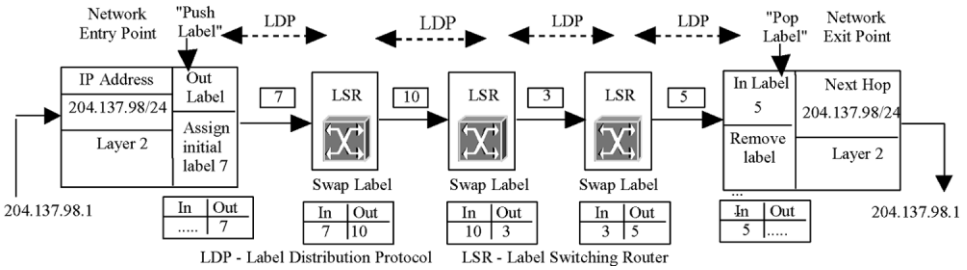

Fig. 4. Packet traversing a LSP.

supports MPLS-based forwarding. Also full duplex traffic requires two LSPs.

MPLS defines new IP signalling and label distribution protocols (LDPs), as well as extensions to existing protocols in order to support multi-vendor interoperability. MPLS does not implement any of the ATM Forum signalling or routing protocols so the complexity of coordinating two different protocol architectures is eliminated. In this way, MPLS brings significant benefits to IP-based networks. An LSR examines only the label in forwarding the packet. The network protocol can be IP or others which is why it is called *Multi-Protocol* Label Switching.

MPLS requires a protocol to distribute labels to setup *Label Switched Paths* (LSPs) and this is defined as a LDP [24]. An LSP is similar to an ATM VC and is unidirectional. MPLS LSRs use the protocol to negotiate the semantics of each label, i.e. how to handle a packet with a particular label from the peer. LSP setup can be control driven (triggered by control traffic such as routing updates) or data driven (triggered by the request of a flow or a *Traffic Trunk*). In MPLS, a traffic trunk is an aggregation of flows with the same service class that can be sent over a LSP. The LSP between two routers can be the same as the layer 3 hop-by-hop route, or the sender LSR can specify an *Explicit Route* (ER) for the LSP. The ability to setup ERs is one of the most useful features of MPLS. A forwarding table indexed by labels is constructed as the result of label distribution. Each forwarding table entry specifies how to process packets carrying the indexing label.

Packets are classified and routed at the ingress LSRs of a MPLS-capable domain. MPLS headers are then inserted. When an LSR receives a labelled packet, it will use the label as the index to look up the forwarding table. This is faster than the process of parsing the routing table in search of the longest match carried out in IP routing [25]. The packet is processed as specified by the forwarding table entry. The incoming label is replaced by the outgoing label and the packet is switched to the next LSR. This label-switching process is similar to ATM's VCI/VPI processing. Inside an MPLS domain, packet forwarding, classification and QoS service are determined by the labels and the class of service (CoS) fields. This makes core LSRs simple. Before a packet leaves a MPLS domain, its MPLS label is removed.

MPLS LSPs can be used as tunnels. When a packet enters the start point of a tunnel, its path is completely determined. With MPLS, a packet's path is completely determined by the label assigned by the ingress LSR. There is no need to enumerate every intermediate router of the tunnel. MPLS is therefore more efficient in terms of header overhead than other tunnelling mechanisms. Thus MPLS has the advantages of providing fast packet classification and forwarding as well as an efficient tunnelling mechanism.

5.4. Service architecture based on MPLS

MPLS can be used together with DiffServ to provide QoS in IP-based networks [26]. In such an architecture it is likely that for each ingress–egress pair a separate LSP is created for each traffic class. In this case, a total number of $CN(N-1)/2$ LSPs are needed, where C is the number of traffic classes and N is the number of boundary routers. To reduce the number of LSPs, all ingress router LSPs to a single egress router can be merged into a *Sink Tree*. The total number of Sink Trees needed is CN. It is also possible to use a single Sink Tree to transmit packets of different traffic classes, and use the CoS bits to differentiate packet classes. In this case, the number of Sink Trees is reduced to N. In this architecture, as the number of transiting flows increases, the number of flows in each LSP or Sink Tree also increases although the number of LSPs or Sink Trees themselves need not increase which makes the architecture more scalable. The operation of the routers are basically the same in this architecture as in the DS field-based architecture previously described in Section 3.1.

Whether a particular ISP's architecture is DS field-based or MPLS-based is transparent to other ISPs. Therefore, the DS field-based architecture and the MPLS-based architecture can easily inter-operate. Each customer domain still needs a BM (Section 4) to allocate services, and to request resources on behalf of the customer domain when the SLA is dynamic. Since LSPs are configured within the ISPs, resource requests can be easily hidden from the core routers by tunnelling them from the ingress to the egress routers. Therefore, BMs may not be needed in MPLS-based ISP networks.

6. Traffic engineering

Traffic engineering is the process of controlling traffic flow through the network and the techniques described in Sections 2–5 all constitute traffic engineering in one form or another. Discussion so far has concentrated on techniques designed to replace routing with switching or label swapping as well as for classifying traffic for forwarding according to various scenarios. The decision on network routes has (so far) been left to the traditional distance vector and link state routing protocols. However QoS-based routing (and constraint based routing, CBR in particular) [27] is a technique gaining favour. It dynamically determines routes based upon constraints such as delay and bandwidth requirements.

6.1. Architectural issues of traffic engineering

Switching techniques designed to offer traffic classification and speed — and as far as possible to minimise routing large volumes of traffic is central to traffic engineering. However such techniques are very much what ATM was designed to do. The momentum behind MPLS and DiffServ in particular, arises from the difficulties in interfacing with resource allocation systems in underlying protocols such as ATM. Further, it cannot be assumed that such a

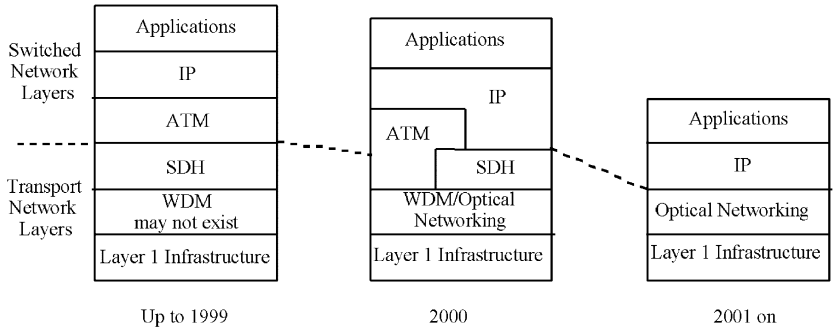

Fig. 5. Progressive simplification of the IP stack.

classification-based underlying network even exists on an end-to-end basis. Trying to link together multiple classification-based frame level protocols to achieve seamless end-to-end connectivity is fraught with problems.

Fig. 5 shows a progressive move to simplification of the underlying architecture.

ATM and SDH are giving way to a simplified structure of IP over optical transport. Fig. 6 shows the equivalent protocol layering with a strong focus towards IP over MPLS over optical networks. Although this simplification of the transport architectures makes a lot of sense, it is having an interesting side effect. The 'thinning' of layers 1 and 2 is resulting in a 'mushrooming' of layer 3 as shown in Fig. 7. Mechanisms to carry out control and management functions such as multicasting, congestion management, transport configuration, protection switching, path management, security, VPN tunneling, caching, filtering, etc. are all still required. Simplifying the underlying layers still means that certain essential control and management functions must be carried out in other parts of the protocol stack.

6.2. Constraint-based routing

Network congestion can result from a shortage of resources or an uneven traffic distribution. Current dynamic routing protocols such as RIP-2, IGRP, OSPF, etc. are based upon well known Bellman-Ford and Dijkstra's algorithms and use relatively simple metrics to determine the shortest path. More recent developments such as the equal-cost multi-path options used in OSPF version 2 [28] and IS–IS [29] assist in distributing the load across multiple paths. QoS routing and CBR in particular, is a more recent development which calculates routes where multiple constraints exist and offers a number of alternative paths that meet the QoS requirements [30–32]. CBR might be used to assign bandwidth or service class characteristics to an LSP or one may want to ensure that alternative routing via separate physical paths is available. Thus CBR takes into account the network topology, flow specifications, availability of links, and other specified policies. Metrics used by CBR includes hop count, bandwidth, transit delay, jitter, availability, monetary cost, etc. This requires routers to distribute link state information and to determine optimum paths accordingly.

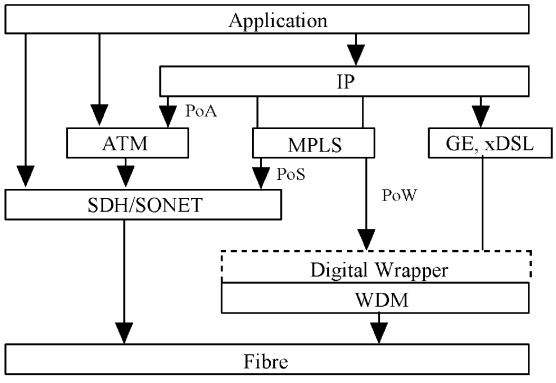

Fig. 6. Evolution in protocol layering.

CBR can be based upon traffic classes, traffic trunks, flow-based and topology-based metrics as well as source and destination addresses. The finer the granularity of the parameters, the better the result although this implies a need for greater bandwidth in order to distribute link-state data. CBR offers support to DiffServ in selecting routes to meet the QoS requirements. It offers support to RSVP in determining optimal paths for resource reservation by taking QoS requirements into consideration.

CBR also operates well with MPLS as even though CBR determines the route based upon resources and topology information and MPLS uses its LDP to setup LSPs, the two benefit each other. The statistics resulting from the setting up of an MPLS's LSP can assist CBR in determining traffic flows between ingress/egress pairs. Thus CBR can calculate the routes for setting up LSPs. Physical paths can be determined by an off-line configuration program but the benefits of using an online method of QoS routing such as CBR are significant. The forwarding states are installed across the network using RSVP signalling.

From an ISP's point of view the network administrator configures the LSP based upon individual constraints and then the network — using CBR — determines the optimal path for the collection of all LSPs given these constraints. Together MPLS and CBR are valuable tools for traffic engineering and [30] provides the full specification of the constraints.

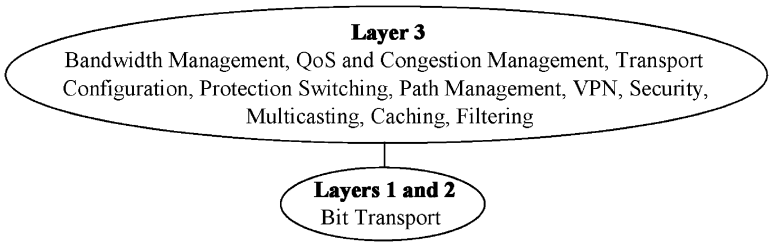

Fig. 7. Enlarging the layer 3 at expense of layers 1 and 2.

7. Conclusions

Unpredictable performance has been the most significant problem in the deployment of IP-based networks. Demand for IP networks has outstripped all predictions — yet the lack of quality of service (QoS) and methodology to traffic engineer these networks has been the single greatest impediment to their deployment. IP infrastructure has to go through a revolution in order to provide the network services and SLAs demanded by the industry. It is simply unacceptable to deploy best-effort networks for many of today's applications.

The issue is further complicated by the heterogeneous nature of networks involving a number of local area and access architectures (wired and wireless) and an equal number and variety of wide area architectures. The concatenation of these architectures must provide a framework for the delivered of end-to-end IP service with mechanisms designed to meet customers' stringent service level requirements. Further, the nature of communications has expanded to include multicast networks — particularly in support of multimedia distribution services. Multicast QoS has yet to be realised on such networks.

IntServ, DiffServ and MPLS are all-important stepping stones in the evolution of a new IP infrastructure. Although ATM is the only networking service to offer classes of service, the realisation that ATM would never be likely to be deployed end-to-end meant that new protocols and architectures had to be designed.

Important work has been — and continues to be undertaken — by the IETF working groups. In particular the Internet Traffic Engineering working group focuses on performance optimisation of traffic, which involves the design, provisioning and tuning of IP Networks. It also addresses issues such as constraint-based routing, resource allocation, and the measurement of inter and intra-domain traffic flows.

The DiffServ working group focuses on methodologies to provide classification of traffic flows to support various applications. This involves well-defined building blocks from which a variety of aggregate behaviours can be built including per-hop behaviour and code-point specification.

The MPLS working group is responsible for standardising a base technology for using label swapping over various link level technologies such as packet-over-SDH/SONET, Frame Relay, ATM and IEEE802 LAN architectures. Of significant importance is the provisioning of MPLS over WDM which has the potential to provide a very high speed service over a relatively simple architecture as illustrated in Figs. 5–7. These architectures must be scalable as well as supporting unicast and multicast traffic flows.

Important work still remains to be done in areas such as per-domain behaviour; traffic conditioning; policy definition, infrastructure and enforcement; QoS routing; and QoS multicasting. Finally, a security framework for many of these new architectures has yet to be designed.

References

[1] R. Branden, D. Clark, S. Shenker, Integrated Services in the Internet Architecture: An Overview, RFC1633, 1994.
[2] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, Resource ReSerVation Protocol (RSVP) — version 1 Functional Specification, RFC 2205, 1997.
[3] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, An Architecture for Differentiated Services, RFC 2475, 1998.
[4] C. Semeria, J. Stewart, Optimizing Routing Software for Reliable Internet Growth, Juniper Networks White Papers, www.juniper.net/techcenter/techpapers/200003.pdf, 2000.
[5] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, J. McMaus, Requirements for Traffic Engineering over MPLS, RFC 2702, 1999.
[6] J. Wroclawski, Specification of the Controlled-Load Network Element Service, RFC 2211, 1997.
[7] T. Li, Y. Rekhter, Provider Architecture for Differentiated Services and Traffic Engineering (PASTE), RFC 2430, 1998.
[8] J. Boyle, R. Cohen, D. Durham, S. Herzog, R. Rajan, A. Sastry, Common Open Service Protocol Usage for RSVP, RFC 2749, 2000.
[9] K. Nichols, S. Blake, F. Baker, D. Black, Definition of the Differentiated Services Field (DS Field) in the IPv4 and. IPv6 Headers RFC 2474, 1998.
[10] C. Metz, IP Switching: Protocols & Architecture, McGraw-Hill, New York, 1999.
[11] K. Nichols, V. Jacobson, L. Zhang, A Two-Bit Differentiated Services Architecture for the Internet, RFC 2638, 1999.
[12] M. Goyal, Performance Analysis of Assured Forwarding, Internet Draft ⟨draft-goyal-diffserv-afstdy-00.txt⟩, February 2000.
[13] B. Braden, et al., Recommendation on Queue Management and Congestion Avoidance in the Internet, RFC 2309, 1998.
[14] K. Muthukrishnan, A. Malis, A Core MPLS IP VPN Architecture, RFC 2917, September 2000.
[15] H. Zhang, Service disciplines for guaranteed performance service in packet-switching networks, Proc. IEEE 83 (10) (1995) 1374–1396.
[16] Y. Bernet, et al., A Framework for Integrated Services Operation over DiffServ Networks, RFC 2998, November 2000.

[17] X. Xiao, L. Ni, Internet QoS: the big picture, IEEE Network 13 (2) (1999) 8–18.
[18] R. Yavatkar, D. Hoffman, Y. Bernet, F. Baker, M. Speer, SBM (Subnet Bandwidth Manager): A Protocol for RSVP-based Admission Control over IEEE 802-style networks, RFC 2814, 2000.
[19] M. Seaman, A. Smith, E. Crawley, J. Wroclawski, Integrated Service Mappings on IEEE 802 Networks, RFC 2815, 2000.
[20] A. Ghanwani, J. Pace, V. Srinivasan, A. Smith, M. Seaman, A Framework for Integrated Services over Shared and Switched IEEE 802 LAN Technologies, RFC 2816, 2000.
[21] E. Rosen, A. Viswanthan, R. Callon, Multiprotocol Label Switching Architecture, RFC 3031, January 2001.
[22] E. Rosen, et al., MPLS Label Stack Encoding, RFC 3032, January 2001.
[23] Y. Rekhter, B. Davie, D. Katz, E. Rosen, G. Swallow, Cisco Systems Tag Switching Architecture Overview, RFC2105, 1997.
[24] L. Andersson et al., Label Distribution Protocol Specification, RFC 3036, January 2001.
[25] S. Nilsson, G. Karlsson, Fast Address Lookup for Internet Routers, ACM SIGCOMM'97, Cannes, France, www.acm.org/sigcomm/sigcomm97, 1997.
[26] F. Le Faucheur et al., MPLS Support of Differentiated Services, Internet Draft ⟨draft-ietf-mpls-diff-ext-08.txt⟩, February 2001.
[27] E. Crawley, A Framework for QoS-based Routing in the Internet, RFC 2386, 1998.
[28] J. Moy, OSPF Version 2, RFC 2178, 1998.
[29] T. Li, H. Smit, IS–IS Extensions for Traffic Engineering, Internet draft ⟨draft-ietf-isis-traffic-02.txt⟩, September 2000.
[30] R. Callon, Constraint-Based LSP Setup using LDP, Internet Draft ⟨draft-ietf-mpls-cr-ldp-05.txt⟩, February 2001.
[31] K. Kompella, Carrying Constraints in RSVP, Internet Draft ⟨draft-kompella-mpls-rsvp-constraints-00.txt⟩, February 2001.
[32] A. Banerjee, Impairment Constraints for Routing in All-Optical Networks, Internet Draft ⟨draft-banerjee-routing-impairments-00.txt⟩, 2001.