Computer Communications 22 (1999) 411–418

Support QoS in IP over ATM

Gung-Chou Lai, Ruay-Shiung Chang*

Department of Computer Science and Information Engineering, National Dong Hwa University, Hualien, Taiwan

Received 2 June 1997; accepted 27 May 1998

Abstract

An integrated service internet running real-time and multimedia applications is rapidly becoming a reality. Meanwhile, ATM technology is appearing in the marketplace. It is an important problem to integrate ATM networks into this integrated service internet. One of the approaches, classical IP over ATM, is now widely deployed, effectively solving the problem of internetworking and interoperability. A key remaining issue is to provide quality-of-service (QoS) guarantees for internet traffic running through ATM subnets. This paper describes a priority scheme, named User Priority, for providing an IP integrated service with QoS over ATM switched virtual circuits (SVCs) to obtain better performance of packet delivery. The User Priority is defined as a three-bit field which uses a type of service (TOS) field in the IP datagram header. This yields eight different service classes with a value of 7 for the highest priority and 0 for the lowest priority. Class 6 and 7 services are for real-time traffic and have their own VCs. Packets with Classes 0 through 5 are sent an aggregate VC. This method is different from the IP over ATM scheme where only one VC is set up between two communicating IP hosts. Thus, packets can be treated differently according to their priorities such that they can enjoy the various QoS guarantees provided by ATM networks. It is backward-compatible with existing IP implementations. These newer options need only to be implemented on the end systems that want to take advantage of them.

Keywords: IP over ATM; LAN emulation; Next-hop routing protocol; Resource reservation; Quality of service

1. Introduction

The asynchronous transfer mode, commonly given the acronym ATM, is the most widely studied and implemented form of cell networking [1]. ATM began as a technology designed specifically to address the needs of the international telecommunications carrier community. It has evolved over the past few years, and the various protocols and interfaces are defined in a set of standards created by the International Telecommunications Union (ITU). This gives network designers a solid base on which to build ATM networks. ATM is the underlying transmission system for the ITU's next-generation ISDN, broadband (B) ISDN. B-ISDN is designed to provide subscriber communications services over a wide range of bit rates from a few megabits to several gigabits. The current ATM standards are designed to allow subscribers access to the telephone networks at speeds of up to 622 Mbits s^{-1}, and it is expected that eventually, gigabit speeds will also be supported as the underlying ATM transmission system is clearly capable of these speeds.

The major selling point of ATM is that it is the first technology that can deliver different types of traffic, such as voice, video and data, over a single digital transport mechanism. ATM can also handle scalable amounts of bandwidth as a result of its switching architecture, which can support multimedia applications and network growth for years to come. As the Internet integrated service (IIS) is becoming important, the ATM will play an important role as a backbone network technology for the Internet.

However, in a very competitive market, ATM cannot be the sole technology used; it is going to cooperate with existing network technologies in the Internet environment. It is hoped that the combined networks will provide quarantees of quality of service (QoS), which is required by network users and for the performance of the Internet. These QoS guarantees, however, come at a price. Contrary to common misconceptions, ATM is a very complex technology [2], perhaps the most complex ever developed by the networking industry. While the structure of ATM cells and cell switching do facilitate the development of hardwired and high-performance ATM switches, the deployment of ATM networks requires an infrastructure which consists of layers of highly complex protocols and softwares. Therefore, one of the challenges that ATM faces is to interoperate with the

* Corresponding author. Fax: 886-2-7376777.
E-mail address: rschang@csie.ndhu.edu.tw (R.-S. Chang)

0140-3664/99/$ - see front matter
PII: S0140-3664(98)00212-6

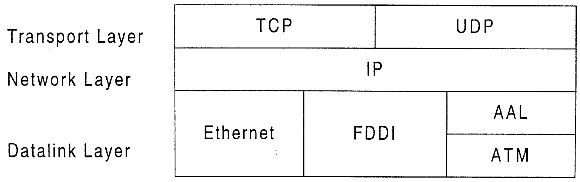

Fig. 1. Internet protocol suite and datalink layer.

vast number of TCP/IP networks. Using IP and ATM together presents some interesting problems because they differ in fundamental ways, from their respective models of data forwarding (connectionless vs. connection-oriented) to support for the preferential treatment of packets (no support vs. the potential for support guarantees). In this paper, we will introduce some strategies and propose a priority scheme to support QoS for IP datagrams carried over the interconnected ATM and TCP/IP networks in IETF IP over ATM [3–5]. The implications of various IP-over-ATM strategies on network performance, particularly the aspects relating to QoS, virtual circuit (VC) multiplexing, and VC management are also addressed.

The paper is structured as follows. In Section 2, the concept of IP over ATM is introduced along with some related work. The protocol design principles are described in Section 3. In Section 4, protocol implementation using ATM VCs with guarantee of performance to carry IP datagrams is shown. Finally, conclusions are given and further studies are discussed in Section 5.

2. The concepts of IP over ATM

In the current Internet, the solution to forward data through an heterogeneous internetwork is provided by the Internet protocol (IP). The IP is almost entirely independent of the subnet technology used—it just makes a few assumptions about the nature of individual subnets. IP packets can traverse many different types of subnets (including ATM networks) without either the senders or the receivers being aware of the details of the networks encountered along the path. Unlike the ATM, the IP is a datagram protocol and does not require the establishment of connections before data are sent.

As the ATM and the Internet is likely to coexist in the future, it is desirable that hosts attaching to these two types of networks can exchange data. One approach is to use an ATM network (with an appropriate adaptation layer) as a datalink layer, similar to Ethernet and FDDI, as shown in Fig. 1.

This method is commonly referred to as IP over ATM (IPOA). An interesting consideration in this approach is how to preserve the QoS in the IP conversation. In addition, the issue of ATM QoS will impact on the multiplexing and VC management. In fact, the performance of individual IP conversations and the resource reservation for a given VC will be a trade-off among different multiplexing policies. Different VC management strategies will impact upon resource reservations and delays.

Much work relating to IP over ATM concerns various paradigms for these services, such as IETF classical IP over ATM [6], ATM Forum LAN emulation [7] and multiprotocol over ATM [8], and how they affect the issues of addressing and routing. Multiplexing and VC management in IP over ATM have been studied for the best-effort service, but QoS issues are still not addressed yet. Although various solutions for supporting QoS or performance guarantees in the internetwork have been proposed, such as the resource reservation protocol (RSVP) [3], they did not deal with the specific characteristics of ATM subnets. In this paper, we will focus on support of the QoS in the IETF classical IP over ATM. We will first introduce the IETF classical IP over ATM architecture in detail. Some multiplexing policies and several types of VC management will then be discussed.

The paradigm for IETF classical IP over ATM is shown in Fig. 2. Nowadays this protocol is commonly used in IP over ATM by multiple ATM switch vendors.

In Fig. 2, several IP members, the LIS (logical IP subnet), have the same IP network/subnet number and address mask. Members of an LIS directly connect to the ATM network, and resolve IP addresses to ATM addresses via ATMARP and vice versa via InATMARP when using switched virtual circuits (SVCs). If a permanent virtual circuit (PVC) is used, members of an LIS will need InATMARP to resolve VCs to IP addresses. Members of an LIS would be able to communicate to each other via the ATM. Two hosts belonging to different subnets but attached to the same ATM network can only communicate via the router that is a member of both subnets. In any case, only one VC will be set up between two communicating hosts. All the traffic, with its priorities and characteristics, will be transmitted through this VC. In this paper, we modify the IP over ATM protocol such that more than one VC will be established to accommodate different kinds of data such that each VC can enjoy the QoS guarantees provided by the underlying ATM networks.

While classical IP over ATM is potentially inefficient in that a path between ATM-attached hosts may require forwarding through a router, it has the advantage of preserving the original semantics of IP subnets. Another approach, taken by the Routing Over Large Clouds (ROLC) Working Group of the IETF, seeks to remove the potential inefficiency of the classical model. In the ROLC model, hosts attached to the same ATM network can communicate directly, even if they do not belong to the same LIS. Since part of the original IP routing model dictates that hosts on different subnets must communicate via a router (rather than directly), this method forces changes to the way that IP routing and forwarding are performed. A next-hop routing protocol (NHRP) [9] is used to send data between subnets directly across the ATM network.

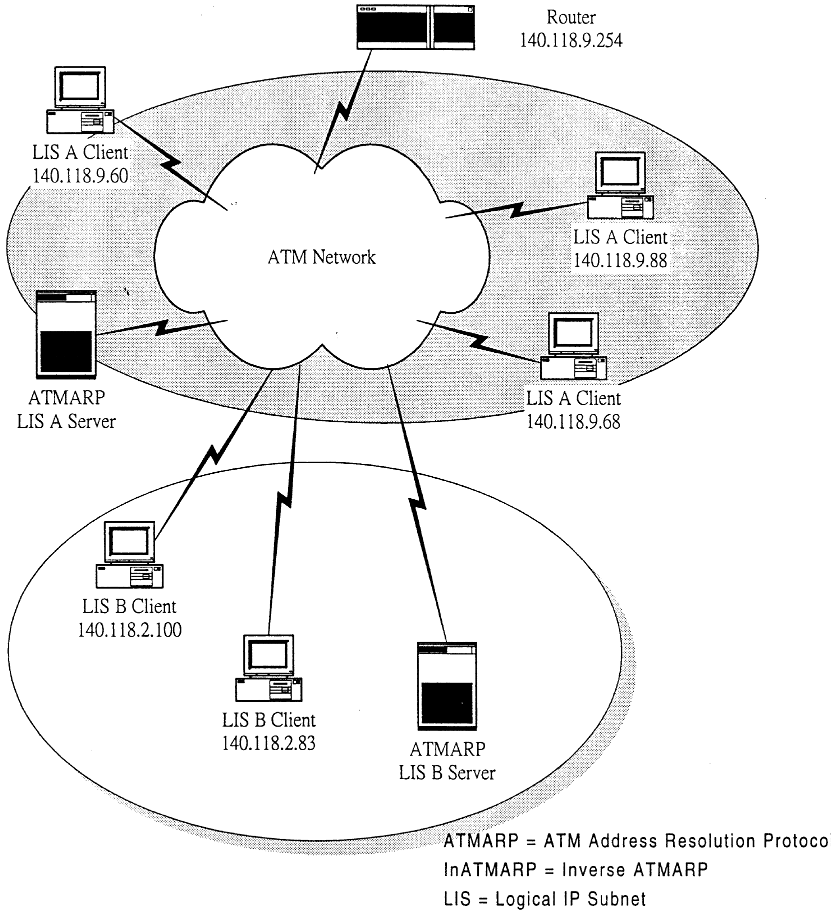

Fig. 2. Model of IETF classical IP over ATM.

The protocol hierarchy of IP over ATM is shown in Fig. 3. It is the encapsulation and transmission of IP network or link layer packets across an ATM adaptation layer (AAL) 5 [10] connection. We know that audio and video applications generally run over UDP, which does not provide reliable data transport. These applications are time-sensitive and need not retransmit packets when they are lost.

2.1. Multiplexing strategies

Three different multiplexing policies [11] can be considered to support the QoS in IP over ATM.

- A VC per pair of routers carrying all traffic passing through the pair of routers, regardless of source or destination host.

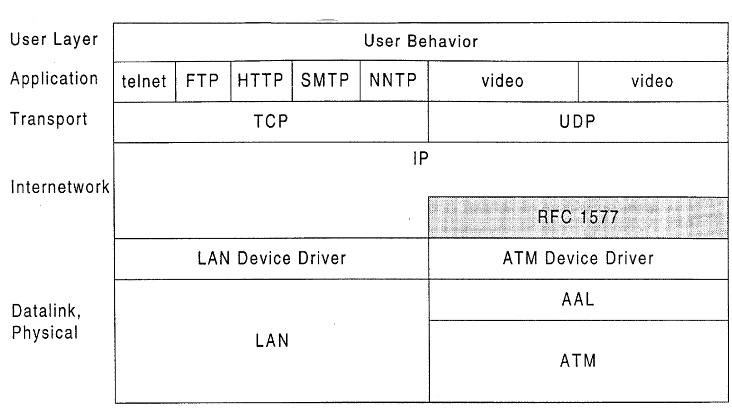

Fig. 3. Protocol stack of IP over ATM.

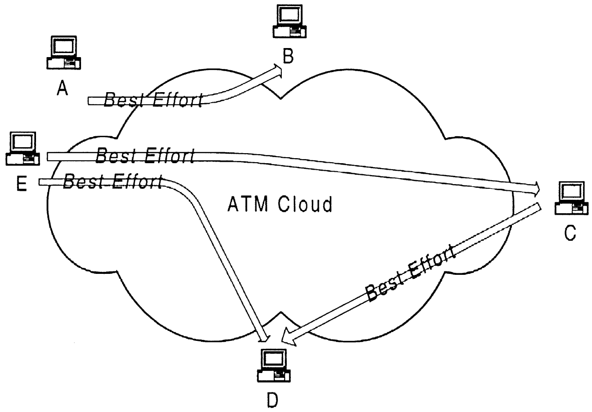

Fig. 4. Virtual circuit connection in classical IP over ATM.

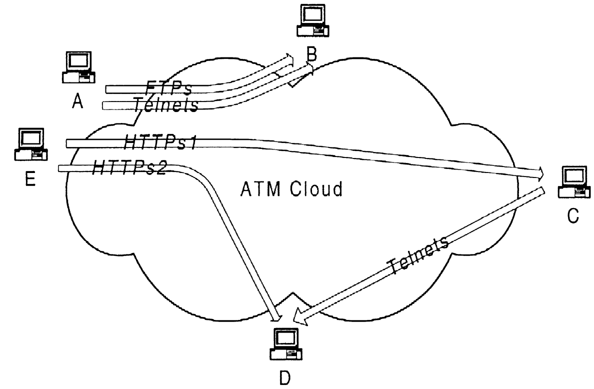

Fig. 5. Schematic showing that sessions in the same application type aggregate to one virtual circuit per host pair.

- A VC per IP conversation (e.g. TCP connection or UDP flow).
- A VC per application type (e.g. one VC for all telnets passing through a pair of hosts).

When performing router multiplexing, it is difficult to preserve a meaningful QoS for each VC, because the nature of the aggregate traffic between the routers is unknown. Therefore, all the policies that use a multiplexing policy of employing a VC per router pair do not use any of ATM's QoS features. A VC per IP conversation is limited by the VPI/VCI quota, and may reserve too many and useless resources on a given VC. However, multimedia applications, such as digital audio and video, would like to have their own VCs because they do not want to be disturbed by other packets. In addition, one VC per application aggregates the same types of applications into one VC. This policy solves the problem described in the discussion of the policy of one VC per IP conversation. Thus, we propose a priority scheme that applies the latter two methods.

2.2. VC management policies

There are three possible VC management policies [11]. They are PVC, SVC and SVC/cache.

- PVC. A set of permanent virtual circuits is established to carry IP packets. These connections are never torn down.
- SVC. Switched virtual circuits, with some time-out policy to be determined, are used to carry IP packets.
- SVC/cache. This is similar to SVC, but with the additional feature that VCs used by other IP conversations can be cached and reused to carry the packets.

It should be observed that these three policies are not entirely independent. In an IP over ATM service using PVCs, it is impractical to set a QoS parameter for an unknown workload traversing a fixed set of VCs. Moreover, the sheer number of VCs required for a complete PVC set may force a multiplexing policy of one VC per router pair. Using a PVC to carry IP traffic would be wasteful and unnecessary. Thus, the only PVC policy that can be considered is QoS-oblivious (no QoS). SVC-with-cache saves connection set-up and tear-down time. Applying SVC/cache sounds better than using SVC only, but a problem is that the resource requirement of the latter IP conversation may not match the previous IP conversation's, such as telnet and FTP. Hence, both SVC and SVC/cache have their own advantages and drawbacks respectively.

3. Protocol design principles

Our goal is to extend the QoS features of an ATM network to IP applications. Although the IP in its current form has no provision for QoS support, the underlying ATM subnet has the capability to offer performance guarantees. We would therefore like Internet applications to gain some of the benefits of ATM performance guarantees, without the end-hosts or applications necessarily being aware of this capability.

As shown in Fig. 4, there is only one SVC between each host pair in classical IP over ATM. These SVCs provide a 'best-effort' service. They do not guarantee any QoS. To improve this, we provide a priority scheme to support the QoS.

Our idea is to group different application types into different VCs. For example, different FTP sessions can be aggregated into one FTP VC and so do telnets and HTTPs. Each VC circuit is used by one application type. As shown in Fig. 5, the same applications share the same VC because they have the same resource requirement characteristic. For example, HTTPs care response time and telnets focus on delay. All these VCCs are SVCs. These VCs are created on demand. In contrast, real-time traffic would like to have its own VC to transmit data, as shown in Fig. 6, such that the other packets cannot disturb it. Using the above idea, we propose a priority scheme by using the precedence bits in the type of service (TOS) field of the IP header to determine whether a flow should initiate a new VC or join

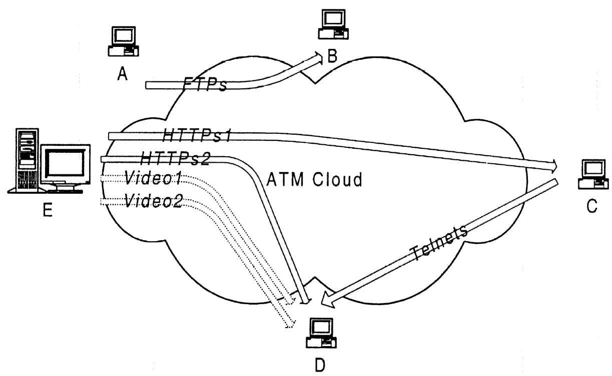

Fig. 6. Schematic showing that two video flows have their own VC, respectively.

an existing one. In contrast, traditional IP over ATM networks use only one SVC to transmit IP packets. There is no bandwidth or delay guarantee in classical IP over ATM networks.

4. Protocol implementation

In this section, we introduce the proposed priority scheme. It uses the TOS [12] field in the IP datagram header and is backward-compatible with existing IP implementation. These newer options need only be implemented on the end systems that wish to take advantage of them.

4.1. Specification of the TOS octet

As shown in Fig. 7, the precedence (named 'User Priority' in this paper) facility is one of the features of the TOS octet in the IP datagram header. The TOS octet consists of three fields.

The first field, 'precedence', is intended to denote the importance or priority of the datagram. This field is not defined and used in current IP implementation. We take this field as the 'User Priority' to determine the QoS types.

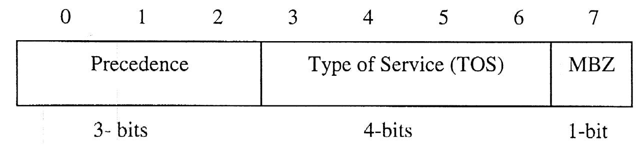

* MBZ: Must Be Zero.

Fig. 7. Type of service in IP datagram header.

The four TOS bits are 'minimize delay', 'maximize throughput', 'maximize reliability', and 'minimize monetary cost', respectively. Table 1 shows the recommendation values of the TOS [12]. Only one of these four bits can be turned on. If all four bits are zero, normal service is implied. RFC 1340 [13] specifies how these bits should be set by all the standard applications. RFC 1349 [12] contains some corrections to RFC 1340 and a more detailed description of the TOS feature. The TOS feature is not supported by most TCP/IP implementations today, although it is being set by newer systems starting with 4.3BSD Reno. Additionally, new routing protocols such as OSPF and IS–IS are capable of making routing decisions based on this field.

The last field, labeled MBZ (must-be-zero) above, is currently unused. The originator of a datagram sets this field to zero (unless participating in an Internet protocol experiment, which makes use of that bit). Routers and recipients of datagrams ignore the value of this field. The field is copied on fragmentation.

4.2. Specification of the User Priority octet

The three-bit User Priority field yields eight different service classes with value 7 denoting the highest priority and 0 the lowest priority. Table 2 defines the semantics of

Table 1
Recommended values for the type-of-service field

Application	Minimize delay	Maximize throughput	Maximize reliability	Minimize monetary cost	Hex value for TOS Octet
Telnet/Rlogin	1	0	0	0	0 × 10
FTP(control)	1	0	0	0	0 × 10
(Data)	0	1	0	0	0 × 08
(Any bulk data)	0	1	0	0	0 × 08
TFTP	1	0	0	0	0 × 10
SMT					0 × 10
(Command phase)	1	0	0	0	
(Data phase)	0	1	0	0	0 × 08
DNS(UDP query)	1	0	0	0	0 × 10
(TCP query)	0	0	0	0	0 × 00
(Zone transfer)	0	1	0	0	0 × 08
ICMP(error)	0	0	0	0	0 × 00
(Query)	0	0	0	0	0 × 00
(Any IGP)	0	0	1	0	0 × 04
SNMP	0	0	1	0	0 × 04
NNTP	0	0	0	1	0 × 02
BOOTP	0	0	0	0	0 × 00

Table 2
Recommendation values of user priority

Priority	Service	ATM QoS	Application
0	Best effort (BE)	UBR	Unspecified traffic
1	Bulk transfer (background)	ABR	NNTP, SMTP
2	Bulk transfer	ABR	FTP, HTTP
3	Interactive traffic	ABR	Telnet, HTTP
4	Internet control message	ABR	ICMP
5	Non-real-time (VBR)	NRT VBR	NRT digital video
6	Real-time (VBR)	RT VBR	Digital video
7	Real-time (CRB)	CBR	Digital audio

the User Priority field values. 0 is referred to as the default User Priority.

4.3. Selecting User Priority classes

The remaining question is how to set the User Priority field. Table 1 describes the priority values for various Internet applications. The next step is to determine the type of application. We distinguish between different applications by the well-known port numbers. The port number is included in the TCP or UDP protocol header [14]. Most of the TCP or UDP applications have the property that they are assigned ‘well-known’ ports. Assigning fixed port numbers to certain applications enables client processes to easily locate server processes. For example, a telnet client application knows that it can locate telnet servers on remote hosts on the TCP port 23. An ATM-attached router can check the source and destination port numbers of a TCP packet; if it sees a well-known port number in the TCP source port field, the packet is likely to be transmitted by a server process to a client process. Conversely, if a well-known port number appears in the TCP destination port field, the packet is likely to be transmitted by a client process to a server process. For non-default port numbers, a configuration table can be built first into IP over ATM.

4.4. Use of the User Priority field in the Internet protocol

For the User Priority facility to be useful, the User Priority field in IP packets must be filled in with reasonable values. When sending a datagram, the Internet protocol sets the User Priority according to the port number. There is no requirement that both the client and server in a connection use the same User Priority. That is called the ‘asymmetric transfer mode’. For example, the server sends packets on the QoS VC, but the client uses the best effort (BE) VC.

When the IP over ATM receives a new connection request, it needs to decide whether this request should initiate a new VC or join an existing one. Fig. 8 shows a flow chart of these procedures.

In Fig. 8, when joining an existing VC, the algorithm checks whether the number of sessions in this VC is already enough. When initiating a new VC, depending on the type of application a different QoS parameter is set before allocating the VC.

4.5. Mapping between the IPv6 priority and the IPv4 User Priority field

The Internet Engineering Task Force is currently designing a successor to IP, known as IPv6 (IP version 6) [15]. IPv6 addresses the primary limitations of IPv4, while retaining much of the same basic protocol architecture. Among the features of IPv6 are an expanded address space (128-bit addresses vs. 32-bit IPv4 addresses), ease of route aggregation for scalability, a redesigned packet header (see Fig. 9) for efficient packet processing, and explicit support for security and authentication.

The 4-bit priority field in the IPv6 datagram header enables a source to identify the desired delivery priority of its packets, relative to other packets from the same source. The priority values are divided into two ranges. Values 0 through 7 are used to specify the priority of traffic for which the source is provided congestion control, such as TCP traffic. Values 8 through 15 are used to specify the priority of traffic that does not back-off in response to congestion. For example, real-time packets are being sent at a constant bit rate.

For congestion-controlled traffic, the priorities shown in Table 3 are recommended for particular application categories.

For non-congestion-controlled traffic, the lowest priority value (8) should be used for those packets that the sender is most willing to have discarded under conditions of congestion, such as high-fidelity video traffic; and the highest value (15) should be used for those packets that the sender is least willing to have discarded, such as low-fidelity audio traffic. There is no relative ordering implied between the congestion-controlled priorities and the non-congestion-controlled priorities.

We can use the MBZ bit in the TOS field of the IPv4 datagram header to simulate the fourth bit in the IPv6 priority field. By using the MBZ bit, the User Priority field in the IPv4 datagram header can be extended to two range priorities as for the IPv6 priority field. Therefore, when IPv6 traffic is tunnelling through IPv4 networks, the priority concept in IPv6 can be still applied. Also, if the IPv4 packets

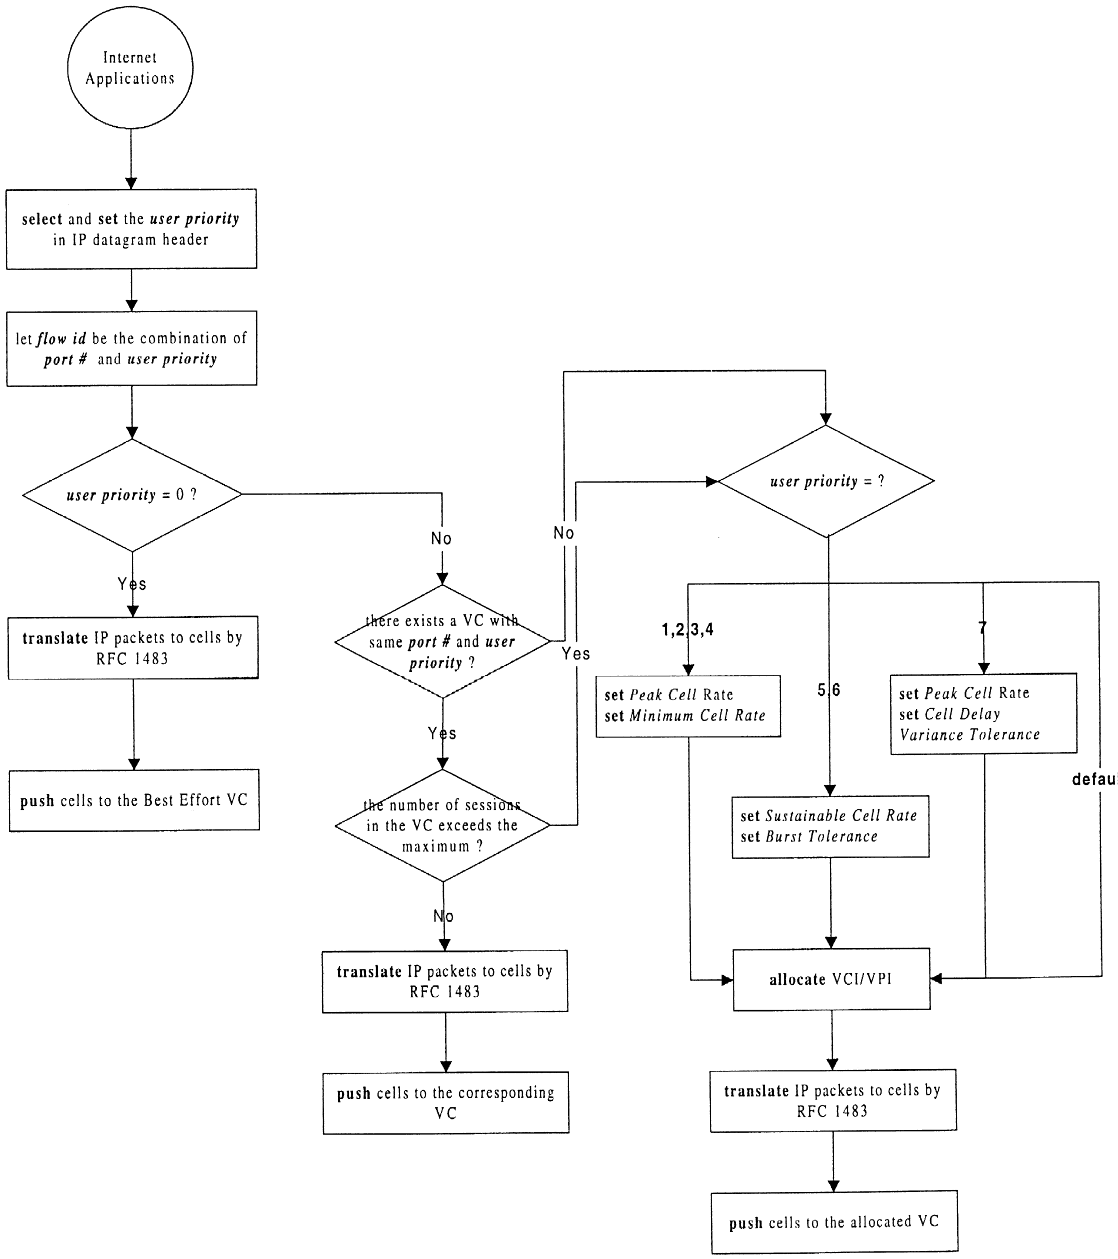

Fig. 8. How the QoS is implemented in IP over ATM.

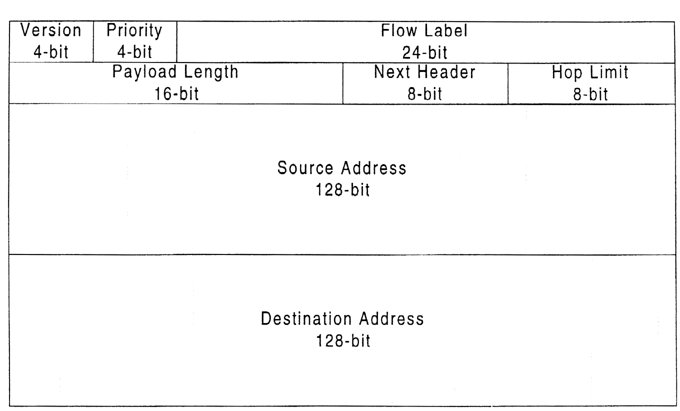

Fig. 9. IPv6 header format.

Table 3
Recommendation values of IPv6 priority

Priority	Description	Application
0	Uncharacterized traffic	Default
1	'Filler' traffic	Netnews
2	Unattended data transfer	E-mail
3	(Reserved)	–
4	Attended bulk transfer	FTP, NFS
5	(Reserved)	–
6	Interactive traffic	Telnet
7	Internet control traffic	Routing protocol, SNMP

finally pass an ATM subnet, the method proposed in this paper could be used to satisfy the QoS requested by the IPv6 packets.

5. Conclusions and future work

In this paper, we have proposed a priority scheme (named the User Priority) by extending the IP datagram header. The goal is to enable a better packet delivery performance in traditional IP over ATM networks with QoS guarantees. Today's networks consist of mostly IP traffic. They can be benefited by application of the method when passing through ATM networks.

Currently, the User Priority scheme does not support a dynamic change of the QoS. There are several commonly mentioned reasons for a change of the reserved QoS. First, an existing receiver can request a new, larger QoS. Second, a sender may change its traffic specification, which can trigger a change in the reservation requests of the receivers.

Finally, a new receiver can make a reservation that is larger than existing reservations. Since the ATM service, as currently defined in UNI 3. × [16] and UNI 4.0, does not allow renegotiation of the QoS of a VC, dynamically changing the reservation means creating a new VC with the new QoS, and tearing down an established VC. Tearing down a VC and setting up a new VC in ATM are time-consuming.

Furthermore, we need an enhanced signalling protocol. Setting up a connection is a hop-by-hop process in UNI 3. × and 4.0. A possible candidate is the connection request protocol (CRP) [17]. It uses a parallel connection set-up and resource management scheme. Its key feature is that it combines address resolution with connection set-up to improve performance. Further, it eliminates the need for IP end-points to support ATM signalling protocols, thereby significantly simplifying their configuration and management.

Acknowledgements

This research is supported in part by the NSC under contract numbers NSC86-2622-E-007-001 and NSC86-2213-E-011-056.

References

[1] Craig Patridge, Gigabit Networking, Addison-Wesley, 1994.
[2] Anthony Alles, ATM Internetworking, Engineering InterOp, Las Vegas, March 1995.
[3] Steve Berson, Classical RSVP and IP over ATM, USC Information Science Institute, April 1, 1996.
[4] R. Cole, D. Shur, C. Villamizar, IP over ATM: a framework document, Internet Request for Comments 1932, April 1996.
[5] M. Perez, F. Liaw, A. Mankin, E. Hoffman, D. Grossman, A. Malis, (Eds.), ATM signaling supporting for IP over ATM, Internet Request for Comments 1755, February 1995.
[6] M. Laubach (Ed.), Classical IP and ARP over ATM, Internet Request for Comments 1577, January 1994.
[7] ATM Forum, LAN Emulation Over ATM, Version 1.0, January 1995.
[8] ATM Forum, MPOA Baseline, Revision 9, September 3, 1996.
[9] James V. Luciani, Dave Katz, David Piscitello, Bruce Cole, NBMA next hop resolution protocol (NHRP), Internet draft draft-ietf-rolc-nhrp-10, October 1996.
[10] Juha Heinanen (Ed.), Multiprotocol encapsulation over ATM adaptation layer 5, Internet Request for Comments 1483, July 1993.
[11] B. Mah, On the use of quality service in IP over ATM, Technical Report UCB/CSD-95-884, University of California, Berkeley, CA, September 1995.
[12] P. Almquist, Type of service in the Internet protocol suite, Internet Request for Comments, 1349, July 1992.
[13] S.J. Reynolds, J. Postel, Assigned numbers, Internet Request for Comments 1340, July 1992.
[14] W. Richard Stevens, TCP/IP Illustrated, Vol. 3, Addison-Wesley, 1995.
[15] S. Deering, Internet protocol, version 6 (IPv6) specification, Internet Request for Comments 1883, December 1995.
[16] ATM User–network interface (UNI) specification, version 3.1, Prentice Hall.
[17] Clark Woodworth, Malathi Veeraraghavan, David Shur, Connection request protocol (CRP) for IP over ATM, IEEE Global Telecommunications Conference, November 1996.