Network Working Group
Request for Comments: 1821
Category: Informational

M. Borden
E. Crawley
Bay Networks
B. Davie
Bellcore
S. Batsell
NRL

August 1995

Integration of Real-time Services in an IP-ATM Network Architecture

Status of the Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

The IETF is currently developing an integrated service model which is designed to support real-time services on the Internet. Concurrently, the ATM Forum is developing Asynchronous Transfer Mode networking which similarly provides real-time networking support. The use of ATM in the Internet as a link layer protocol is already occurring, and both the IETF and the ATM Forum are producing specifications for IP over ATM. The purpose of this paper is to provide a clear statement of what issues need to be addressed in interfacing the IP integrated services environment with an ATM service environment so as to create a seamless interface between the two in support of end users desiring real-time networking services.

Table of Contents

1.0 Introduction

The traditional network service on the Internet is best-effort datagram transmission. In this service, packets from a source are sent to a destination, with no guarantee of delivery. For those applications that require a guarantee of delivery, the TCP protocol will trade packet delay for correct reception by retransmitting those packets that fail to reach the destination. For traditional computer-communication applications such as FTP and Telnet in which correct delivery is more important than timeliness, this service is satisfactory. However, a new class of application which uses multiple media (voice, video, and computer data) has begun to appear on the Internet. Examples of this class of application are video teleconferencing, video-on-demand, and distributed simulation. While these applications can operate to some extent using best-effort delivery, trading packet delay for correct reception is not an acceptable trade-off. Operating in the traditional mode for these applications results in reduced quality of the received information and, potentially, inefficient use of bandwidth. To remedy this problem the IETF is developing a real-time service environment in which multiple classes of service are offered [6]. This environment will greatly extend the existing best-effort service model to meet the needs of multimedia applications with real-time constraints.

At the same time that this effort is underway in the IETF, Asynchronous Transfer Mode (ATM) is being developed, initially as a replacement for the current telephone network protocols, but more recently as a link-layer protocol for computer communications. As it was developed from the beginning with telephone voice applications in mind, a real-time service environment is an integral part of the protocol. With the approval of UNI 3.1 by the ATM Forum, the ATM standards now have several categories of service. Given the wide acceptance of ATM by the long-line carriers, the use of ATM in the Internet is, if not guaranteed, highly likely. The question now becomes, how can we successfully interface between the real-time services offered by ATM and the new,integrated service environment soon to be available in the IP protocol suite. The current IP over ATM standards assume no real-time IP protocols. It is the purpose of this RFC to clearly delineate what the issues are in integrating real-time services in an IP-over-ATM network [10,15,19,20,21].

In the IP-over-ATM environment, as in many others, multicast routing adds an additional set of challenges. While the major focus of this paper is quality of service (QoS) issues, it is unwise at best to ignore multicast when considering these issues, especially since so many of the applications that motivate the provision of real time QoS also require efficient multicast support. We will therefore try to keep considerations of multicast in the foreground in the following discussion.

One of the primary motivations for this document is a belief by the authors that ATM should, if possible, be used as more than a leased line replacement. That is to say, while it is possible for the Internet to be overlaid on constant bit rate (CBR), permanent virtual circuits (PVCs), thus reducing IP over ATM to a previously solved problem, we believe that this is unlikely to be the most efficient way to use ATM services as they are offered by carriers or as they appear in LANs. For example, a carrier offering a CBR service must assume that the peak bit rate can be used continuously with no degradation in quality and so resources must be allocated to the connection to provide that service, even if the peak rate is in fact rarely used. This is likely to make a CBR service more expensive that a variable bit rate service of the same peak capacity. Another way to view this is that the new IP service model will allow us to associate information about the bandwidth requirements of applications with individual flows; surely it is not wise to discard this information when we request a service from an ATM subnet.

While we believe that there is a range of capabilities in ATM networks that can be effectively used by a real-time Internet, we do not believe that just because ATM has a capability, the Internet must use it. Thus, our goal in this RFC is to begin to explore how an Internet with real time service capability might make most effective use of ATM networks. Since there are a number of problems to be resolved to achieve this effective use, our major goal at this point is to describe the scope of the problems that need to be addressed.

2.0 Problem Space Overview

In this section we aim to describe in high level terms the scope of the problem that will be explored in more detail in later sections.

2.1 Initial Assumptions

We begin by assuming that an Integrated Services Internet, i.e., an Internet with a range of qualities of service to support both real-time and non-real-time applications, will eventually happen. A number of working groups are trying to make this happen, notably

* the Integrated Services group (int-serv), which is working to define a new IP service model, including a set of services suited to a range of real-time applications;

* the Resource reservation Setup Protocol group (rsvp), which is defining a resource reservation protocol [7] by which the appropriate service for an application can be requested from the network;

* the Internet Streams Protocol V2 group (ST-II), which is updating [27], a stream-oriented internet protocol that provides a range of service qualities.

In addition, the IETF IP over ATM working group and the ATM Forum Multiprotocol over ATM group are working to define a model for protocols to make use of the ATM layer.

Since these groups have not yet generated standards, we will need to do some amount of extrapolation to predict the problems that may arise for IP over ATM. We also assume that the standards being developed in the ATM Forum will largely determine the service model for ATM. Again, some extrapolation may be needed. Given these assumptions, this paper aims to explore ways in which a future Integrated Services Internet might make effective use of ATM as it seems likely to be deployed.

2.2 Topologies Under Consideration

Figure 1 shows a generic internetwork that includes ATM and non-ATM subnetworks. This paper aims to outline the problems that must be addressed to enable suitable quality of service to be provided end-to-end across such a network. The problem space, therefore, includes

* communication across an 'ATM-only' network between two hosts directly connected to the ATM network;

* communication between ATM-connected hosts which involves traversing some non-ATM subnets;

* communication between a host on a non-ATM subnet and a host directly connected to ATM;

* communication between two hosts, neither of which has a direct ATM connection, but which may make use of one or more ATM networks for some part of the path.

```
                     [H]
                      |                           [H]
              ________|________________________    |
              |                                |   |
      ________|__                        ______|___|____
      |         |                        |             |
      |  ATM   [R]                      [R]  ATM       |
      |  Cloud  |                        |   Cloud     |___[H]
      |         |     Non-ATM Internet   |             |
      |         |                       [R]            |
      |________[R]                       |_____________|
       |      |                                |
       |      |                                |
      [H]     |________________________________|
                                        |
                                        |
                                       [H]
```

[H] = Host
[R] = Router

Figure 1

In the last case, the entities connected to the ATM network are IP routers, and it is their job to manage the QoS provided by the ATM network(s) in such a way that the desired end-to-end QoS is provided to the hosts. While we wish to describe the problem space in a way that covers all of these scenarios, the last is perhaps the most general, so we will use it for most illustrative purposes. In particular, we are explicitly not interested in ways of providing QoS that are applicable only to a subset of these situations. We claim that addressing these four situations is sufficiently general to cover other situations such as those in which several ATM and non-ATM networks are traversed.

It is worth mentioning that the ATM networks in this case might be local or wide area, private or public. In some cases, this distinction may be significant, e.g., because there may be economic implications to a particular approach to providing QoS.

2.3 Providing QoS in IP over ATM - a walk-through

To motivate the following discussion, this section walks through an example of what might happen when an application with a certain set of QoS needs starts up. For this example, we will use the fourth case mentioned above, i.e., two hosts connected to non-ATM networks, making use of an ATM backbone.

A generic discussion of this situation is made difficult by the fact that the reservation of resources in the Internet may be sender or receiver initiated, depending on the specifics of the setup protocol. We will attempt to gloss over this distinction for now, although we will return to it in Section 4. We will assume a unicast application and that the traffic characteristics and the QoS requirements (such as delay, loss, throughput) of the application are known to at least one host. That host launches a request for the desired QoS and a description of the expected traffic into the network; at some point this request hits a router at the edge of the ATM network. The router must examine the request and decide if it can use an existing connection over the ATM network to honor the request or whether it must establish a new connection. In the latter case, it must use the QoS and traffic characterizations to decide what sort of ATM connection to open and to describe the desired service to the ATM network. It must also decide where to open the connection to. Once the connection is opened, the request is forwarded across the ATM network to the exit router and then proceeds across the non-ATM part of the network by the normal means.

We can see from the above description that there are several sets of issues to be discussed:

* How does the IP service model, with certain service classes and associated styles of traffic and QoS characterization, map onto the ATM service model?

* How does the IP reservation model (whatever it turns out to be) map onto ATM signalling?

* How does IP over ATM routing work when service quality is added to the picture?

These issues will be discussed in the following sections.

3.0 Service Model Issues

There are several significant differences between the ways in which IP and ATM will provide QoS. When IP commits to provide a certain QoS to an application according to the Internet service model, it must be able to request an appropriate QoS from the ATM network using the ATM service model. Since these service models are by no means the same, a potentially complex mapping must be performed for the IP layer to meet its commitments. The details of the differences between ATM and IP and the problems presented by these differences are described below.

We may think of a real-time service model as containing the following components:

* a way to characterize traffic (sometimes called the Tspec);

* a way to characterize the desired quality of service (the Rspec).

We label these components as traffic characterization and QoS characterization. Each of these components is discussed in turn in the following sections.

As well as these aspects of the service model, both ATM and IP will have a number of mechanisms by which the model is implemented. The mechanisms include admission control, policing, and packet scheduling. A particularly important mechanism is the one by which end-nodes communicate their QoS needs and traffic characteristics to the network, and the network communicates admission control decisions to the end-nodes. This is referred to as resource reservation or signalling, and is the subject of Section 4. In fact, it seems to be the only mechanism where significant issues of IP/ATM integration arise. The details of admission control, policing and packet scheduling are largely internal to a single network element and we do not foresee significant problems caused by the integration of IP and ATM. For example, while there may be plenty of challenges in designing effective approaches to admission control for both IP and ATM, it is not apparent that there are any special challenges for the IP over ATM environment. As the walk-through of Section 2.3 described, a reservation request from a host would at some point encounter the edge of the ATM cloud. At this point, either a new connection needs to be set up across the ATM cloud, or the router can decide to carry the requested traffic over an existing virtual circuit. If the ATM cloud cannot create a new connection as requested, this would presumably result in an admission control failure which would cause the router to deny the reservation request.

3.1 Traffic Characterization

The traffic characterization provided by an application or user is used by the network to make decisions about how to provide the desired quality of service to this application and to assess the effect the new flow will have on the service provided to existing flows. Clearly this information feeds into the admission control decision process.

In the Internet community, it is assumed that traffic will in general be bursty and that bursty traffic can be characterized by a ‘token bucket’. While ATM does not expect all traffic to be bursty (the Continuous Bit Rate class being defined specifically for non-bursty traffic), it uses an essentially equivalent formulation for the characterization of traffic that is bursty, referred to as the Generic Cell Rate Algorithm (GCRA). However, ATM in some classes also requires specification of peak cell rate, whereas peak rates are not currently included in the IP traffic characterizations. It may be possible to use incoming interface speeds to determine an approximate peak rate.

One of the functions that must be performed in order to carry IP traffic over an ATM network is therefore a mapping from the characterization of the traffic as supplied to IP to a characterization that is acceptable for ATM. While the similarity of the two characterizations suggests that this is straightforward, there is considerable flexibility in the mapping of parameters from IP to ATM. As an extreme example, a router at the edge of an ATM cloud that expects to receive bursts of IP packets on a non-ATM interface, with the bursts described by some token bucket parameters, could actually inject ATM cells at a constant rate into the ATM network. This may be achieved without significant buffering if the ATM link speed is faster than the point-to-point link speed; alternatively, it could be achieved by buffering out the burstiness of the arriving traffic. It seems more reasonable to map an IP flow (or a group of flows) with variable bandwidth requirements onto an ATM connection that accommodates variable bit rate traffic. Determining how best to map the IP traffic to ATM connections in this way is an area that warrants investigation.

A potential complication to this process is the fact that the token bucket parameters are specified at the edge of the IP network, but that the specification of the GCRA parameters at the entry to an ATM network will frequently happen at a router in the middle of an IP network. Thus the actual burstiness that is encountered at the router may differ from that described by the IP token bucket parameters, as the burstiness changes as the traffic traverses a network. The seriousness of this problem needs to be understood to permit efficient resource utilization.

3.2 QoS Characterization

In addition to specifying the traffic that they will submit to the network, applications must specify the QoS they require from the network. Since the goal is to carry IP efficiently over ATM networks, it is necessary to establish mechanisms by which QoS specifications for IP traffic can be translated into QoS specifications that are meaningful for an ATM network.

The proposed method of QoS specification for the Internet is to specify a ‘service class’ and some set of parameters, depending on the service class. The currently proposed service classes are

* guaranteed, which provides a mathematically guaranteed delay bound [23];

* predictive delay, which provides a probabilistic delay bound [24];and

* controlled delay, which merely tries to provide several levels of delay which applications may choose between [25].

These are in addition to the existing ‘best-effort’ class. More IP service classes are expected in the future. ATM has five service classes:

* CBR (constant bit rate), which emulates a leased line, providing very tightly constrained delay and designed for applications which can use a fixed bandwidth pipe;

* VBR (variable bit rate)-real-time which attempts to constrain delay for applications whose bandwidth requirements vary;

* VBR-non-real-time, intended for variable bandwidth applications without tight delay constraints;

* UBR (unspecified bit rate) which most closely approximates the best effort service of traditional IP;

* ABR (available bit rate) which uses a complex feedback mechanism to control loss.

Each class requires some associated parameters to be specified, e.g., CBR requires a peak rate. Observe that these classes are by no means in direct correspondence with the IP classes. In some cases, ATM classes require parameters which are not provided at the IP level, such as loss rate, to be specified. It may be necessary to assume reasonable default values in these cases.

The major problem here is this: given traffic in a particular IP service class with certain QoS parameters, how should it be sent across an ATM network in such a way that it both meets its service commitments and makes efficient use of the ATM network’s resources? For example, it would be possible to transport any class of IP traffic over an ATM network using the constant bit rate (CBR) ATM class, thus using the ATM network like a point-to-point link. This would allow IP to meet its service commitments, but would be an inefficient use of network resources in any case where the IP traffic was at all bursty (which is likely to be most cases). A more reasonable approach might be to map all IP traffic into a variable bit rate (VBR) class; certainly this class has the flexibility to accommodate bursty IP traffic more efficiently than CBR.

At present, the IETF is not working on any service classes in which loss rate is considered as part of the QoS specification. As long as that is the case, the fact that ATM allows target loss rates to be specified is essentially not an issue. However, we may certainly expect that as the IP service model is further refined, service classes that include specifications of loss may be defined. At this point, it will be necessary to be able to map between loss rates at the IP level and loss rates at the ATM level. It has already been shown that relatively small loss rates in an ATM network can translate to high loss rates in IP due to the fact that each lost cell can cause the loss of an entire IP packet. Schemes to mitigate this problem, which include the proposed approach to implementing the ABR class, as well as other solutions [22], have been proposed. This is clearly likely to be an important issue in the future.

4.0 Resource Reservation Styles

ATM uses a signalling protocol (Q.2931) both to establish virtual connections and to allocate resources to those connections. It has many of the characteristics of a 'conventional' signalling protocol, such as being sender-driven and relying on hard-state in switches to maintain connections. Some of the key characteristics are listed in the table below. In the current standards, the QoS associated with a connection at setup time cannot be changed subsequently (i.e., it is static); in a unicast connection, resources are allocated in both directions along the path, while in the multicast case, they are allocated only from the sender to the receivers. In this case, all senders receive the same QoS.

Two protocols have been proposed for resource reservation in IP. The first (chronologically) is ST-II, the other is RSVP. Each of these, and its relationship to ATM, is discussed in the following sections.

4.1 RSVP

IP has traditionally provided connectionless service. To support real-time services in a connectionless world, RSVP has been proposed to enable network resources to be reserved for a connectionless data stream. ATM, on the other hand, provides a connection-oriented service, where resource reservations are made at connection setup time, using a user-network interface (UNI) and a network-network interface (NNI) signalling protocol.

Category	RSVP	ATM (UNI 3.0)
Orientation	Receiver-based	Sender-based
State	Soft state (refresh/time-out)	Hard state (explicit delete)
QoS SetupTime	Separate from route establishment	Concurrent with route establishment
QoS Changes?	Dynamic QoS	Static QoS (Fixed at setup time)
Directionality	Unidirectional resource allocation	Bidirectional allocation for unicast Unidirectional allocation for multicast
Heterogeneity	Receiver heterogeneity	Uniform QoS to all receivers

The principles used in the design of RSVP differ from those of ATM in the following respects:

* Resource reservations in IP hosts and routers are represented by soft state, i.e., reservations are not permanent, but time out after some period. Reservations must be refreshed to prevent time-out, and may also be explicitly deleted. In ATM, resources are reserved for the duration of a connection, which must be explicitly and reliably deleted.

* The soft state approach of RSVP allows the QoS reserved for a flow to be changed at any time, whereas ATM connections have a static QoS that is fixed at setup time.

* RSVP is a simplex protocol, i.e., resources are reserved in one direction only. In ATM, connections (and associated reservations) are bi-directional in point-to-point calls and uni-directional in point-to-multipoint calls.

* Resource reservation is receiver-initiated in RSVP. In ATM, resources are reserved by the end system setting up the connection. In point-to-multipoint calls, connection setup (and hence resource reservation) must be done by the sender.

* RSVP has explicit support for sessions containing multiple senders, namely the ability to select a subset of senders, and to dynamically switch between senders. No such support is provided by ATM.

* RSVP has been designed independently of other architectural components, in particular routing. Moreover, route setup and resource reservation are done at different times. In ATM, resource reservation and route setup are done at the same time (connection setup time).

The differences between RSVP and ATM state establishment, as described above, raise numerous problems. For example, since point-to-point connections are bidirectional in ATM, and since reservations can be made in both directions, receiver-initiated resource reservations in RSVP can be simulated in ATM by having the receiver set up the connection and reserve resources in the backward direction only. However, this is potentially wasteful of connection resources since connections are only ever used to transfer data in one direction even though communication between the two parties may be bidirectional. One option is to use a 'point-to-multipoint' ATM connection with only one receiver. Of course, the fact that the RSVP reservation request is made by the receiver(s) means that this request must be somehow communicated to the sender on the ATM network. This is somewhat analogous to the receiver-oriented join operation of IP multicast and the problems of implementing it over ATM, as discussed in Section 6. In general, the efficiency of any proposed connection management scheme needs to be investigated in both unicast and multicast contexts for a range of application requirements, especially at a large scale.

The use by RSVP of 'soft state' as opposed to explicit connections means that routers at the ATM network's edges need to manage the opening and closing of ATM connections when RSVP reservations are made and released (or time out). The optimal scheme for connection setup and tear-down will depend on the cost of setting up a connection versus the cost of keeping the connection open for possible future use by another stream, and is likely to be service class-dependent. For example, connections may be left open for reuse by best-effort traffic (subject to sufficient connections being available), since no resources are explicitly reserved. On the other hand, connections supporting the real-time service classes are likely to be expensive to leave open since resources may be allocated even when the connection is idle. Again, the cost incurred will depend on the class. For example, the cost of an open, idle ‘guaranteed’ QoS connection is likely to be significantly more expensive than a connection providing predictive or controlled delay service. Note that connections can be reused for traffic of the same class with compatible QoS requirements, and that it may sometimes be possible to use a ‘higher quality’ class to substitute for a lower quality one.

Another characteristic of RSVP which presents problems for ATM is the use of PATH messages to convey information to receivers before any reservation is made. This works in IP because routing is performed independently of reservation. Delivery of PATH messages across an ATM network is therefore likely to require a mechanism for setting up connections without reservations being made. The connection also needs to be of sufficient quality to deliver PATH messages fairly reliably; in some circumstances, a low quality best effort service may be inadequate for this task. A related issue is the problem of advertising services prior to reservations. The OPWA model (one pass with advertising) requires network elements to advertise the QoS that they are able to provide so that receivers can decide what level of reservation to request. Since these advertisements may be made prior to any resources having been reserved in the ATM network, it is not clear how to make meaningful advertisements of the QoS that might be provided across the ATM cloud.

Finally, the multiparty model of communication is substantially different in RSVP and ATM. Emulating RSVP receiver-initiation using ATM point-to-multipoint connections is likely to cause severe scaling problems as the number of receivers becomes large. Also, some functions of RSVP are not currently provided by ATM. For example, there is no support for different receiver requirements and capabilities-all receivers in a session receive the same QoS, which is fixed at the time the first receiver is added to the multicast tree. It is likely that ATM support for multi-party sessions will be enhanced in later versions of the standards. It is necessary for such support to evolve in a manner compatible with RSVP and IP multicast routing protocols if large ATM clouds are to be deployed successfully.

4.2 ST-II

ST-II [27] and ST2+ [12] (referred to generically as ST hereafter) have data distribution and resource reservation schemes that are similar to ATM in many respects.

* ST is connection oriented using "hard state". Senders set up simplex data flows to all receivers closely matching point-to-multipoint connections in ATM. Routing decisions are made when

the connection is made and are not changed unless there is a failure in the path. Positive acknowledgment is required from all receivers. ST2+ [12] adds a receiver-based JOIN mechanism that can reduce the burden on senders to track all receivers.

* ST reserves network resources at connection setup time. The ST CONNECT message contains a flowspec indicating the resources to be reserved for the stream. Agents along the path may change the flowspec based on restrictions they may need to impose on the stream. The final flowspec is returned to the sender in the ACCEPT message from each receiver or target.

Category	RSVP	ATM (UNI 3.0)
Orientation	Sender-based	Sender-based
State	Hard state (explicit disconnect)	Hard state (explicit delete)
QoS SetupTime	Concurrent with stream setup	Concurrent with route establishment
QoS Changes?	Dynamic QoS	Static QoS (Fixed at setup time)
Directionality	Unidirectional resource allocation	Bidirectional allocation for unicast Unidirectional allocation for multicast
Heterogeneity	Receiver heterogeneity	Uniform QoS to all receivers

These similarities make mapping ST services to ATM simpler than RSVP but the mapping is still not trivial. The task of mapping the ST flowspec into an ATM service class still has to be worked out. There may be policy issues related to opening a new VC for each stream versus aggregating flows over an existing VC.

Additionally, ST has some differences with UNI 3.1 that can cause problems when integrating the two protocols:

* In ST, changes to active stream reservations are allowed. For example, if the flowspec received from the target is not sufficient for the stream, the sender can send a CHANGE message, requesting a different QoS. UNI 3.1 does not allow changes to the QoS of a VC after it is set up. Future ATM UNI specifications are contemplating allowing changes to a VC after set up but this is still preliminary. In the meantime, policies for over reservation or aggregation onto a larger VC may be needed.

* ST uses simplex streams that flow in only one direction. This is fine for UNI 3.1 point-to-multipoint connections since the data flow is only in one direction. When mapping a point-to-point ST connection to a standard point-to-point ATM VC, the reverse flow connection is wasted.

This can be solved simply by using only point-to-multipoint VCs, even if there is only one receiver.

4.3 Mapping IP flows to ATM connections

In general, there will be a great deal of flexibility in how one maps flows at the IP level to connections at the ATM level. For example, one could imagine setting up an ATM connection when a reservation message arrives at the edge of an ATM cloud and then tearing it down as soon as the reservation times out. However, to minimize latency or perhaps for economic reasons, it may be preferable to keep the ATM connection up for some period in case it is needed. Similarly, it may be possible or desirable to map multiple IP flows to a single ATM connection or vice versa.

An interesting situation arises when a reservation request is received for an existing route across the cloud but which, when added to the existing reservations using that connection, would exceed the capacity of that connection. Since the current ATM standards do not allow the QoS of a connection to be changed, there are two options: tear down the old connection and create a new one with the new, larger allocation of resources, or simply add a new connection to accommodate the extra traffic. It is possible that the former would lead to more efficient resource utilization. However, one would not wish to tear down the first connection before the second was admitted, and the second might fail admission control because of the resources allocated to the first. The difficulties of this situation seem to argue for evolution of ATM standards to support QoS modification on an existing connection.

5.0 End System Issues

In developing an integrated IP-ATM environment the applications need to be as oblivious as possible of the details of the environment: the applications should not need to know about the network topology to work properly. This can be facilitated first by a common application programing interface (API) and secondly by common flow and filter specifications [18].

An example of a common API that is gaining momentum is the BSD sockets interface. This is a UNIX standard and, with Winsock2, has also become a PC standard. With the IETF integrated service environment just beginning to appear in the commercial marketplace, the ability to standardize on one common interface for both IP and ATM applications is still possible and must be seriously and quickly pursued to insure interoperability.

Since the IP integrated service and ATM environments offer different QoS service types, an application should specify sufficient information in its flow specification so that regardless of the topology of the network, the network can choose an acceptable QoS type to meet the applicationUs needs. Making the application provide sufficient information to quantify a QoS service and allowing the network to choose the QoS service type is essential to freeing the application from requiring a set network topology and allowing the network to fully utilize the features of IP and ATM.

6.0 Routing Issues

There is a fundamental difference between the routing computations for IP and ATM that can cause problems for real-time IP services. ATM computes a route or path at connection setup time and leaves the path in place until the connection is terminated or there is a failure in the path. An ATM cell only carries information identifying the connection and no information about the actual source and destination of the cell. In order to forward cells, an ATM device needs to consult a list of the established connections that map to the next hop device, without checking the final destination.

In contrast, routing decisions in IP are based on the destination address contained in every packet. This means that an IP router, as it receives each packet, has to consult a table that contains the routes to all possible destinations and the routing decision is made based on the final destination of the packet. This makes IP routing very robust in the face of path changes and link failures at the expense of the extra header information and the potentially larger table lookup. However, if an IP path has been selected for a given QoS, changes in the route may mean a change in the QoS of the path.

6.1 Multicast routing

Considerable research has gone into overlaying IP multicast models onto ATM. In the MARS (Multicast Address Resolution Server) model [1], a server is designated for the Logical IP Subnet (LIS) to supply the ATM addresses of the hosts in the IP multicast group, much like the ATM ARP server [15]. When a host or router wishes to send to a multicast group on the LIS, a query is made to the MARS and a list of the ATM address of the hosts or routers in the group is returned. The sending host can then set up point-to-point or point-to-multipoint VCs to the other group members. When a host or a router joins an IP multicast group, it notified the MARS. Each of the current senders to the group is then notified of the new group member so that the new member can be added to the point to multipoint VC's.

As the number of LIS hosts and multicast groups grows, the number of VCs needed for a one-to-one mapping of VCs to multicast groups can get very large. Aggregation of multicast groups onto the same VC may be necessary to avoid VC explosion. Aggregation is further complicated by the QoS that may be needed for particular senders in a multicast group. There may be a need to aggregate all the multicast flows requiring a certain QoS to a set of VCs, and parallel VCs may be necessary to add flows of the same QoS.

6.2 QoS Routing

Most unicast and multicast IP routing protocols compute the shortest path to a destination based solely on a hop count or metric. OSPF [16] and MOSPF [17] allow computation based on different IP Type of Service (TOS) levels as well as link metrics, but no current IP routing protocols take into consideration the wide range of levels of quality of service that are available in ATM or in the Integrated Services models. In many routing protocols, computing all the routes for just the shortest path for a large network is computationally expensive so repeating this process for multiple QoS levels might be prohibitively expensive.

In ATM, the Private Network-to-Network Interface (PNNI) protocol [13] communicates QoS information along with routing information, and the network nodes can utilize this information to establish paths for the required QoS. Integrated PNNI (I-PNNI) [9] has been proposed as a way to pass the QoS information available in ATM to other routing protocols in an IP environment.

Wang & Crowcroft [28] suggest that only bandwidth and delay metrics are necessary for QoS routing and this would work well for computing a route that required a particular QoS at some setup time, but this goes against the connectionless Internet model. One possible solution to the exhaustive computation of all possible routes with all possible QoS values would be to compute routes for a common set of QoS values and only then compute routes for uncommon QoS values as needed, extracting a performance penalty only on the first packets of a flow with an uncommon QoS. Sparse multicast routing protocols that compute a multicast path in advance or on the first packets from a sender (such as CBT [5] and MOSPF [17]) could also use QoS routing information to set up a delivery tree that will have adequate resources.

However, no multicast routing protocols allow the communication of QoS information at tree setup time. Obtaining a tree with suitable QoS is intended to be handled by RSVP, usually after the distribution tree has been set up, and may require recomputation of the distribution tree to provide the requested QoS.One way to solve this problem is to add some "hints" to the multicast routing protocols so they can get an idea of the QoS that the multicast group will require at group initiation time and set up a distribution tree to support the desired QoS. The CBT protocol [5] has some TBD fields in its control headers to support resource reservation. Such information could also be added to a future IGMP [11] JOIN message that would include information on the PIM Rendezvous Point (RP) or CBT Core.

Another alternative is to recompute the multicast distribution tree based on the RSVP messages but this has the danger of losing data during the recomputation. However, this can leave a timing window where other reservations can come along during the tree recomputation and use the resources of the new path as well as the old path, leaving the user with no path to support the QoS desired.

If unicast routing is used to support multicast routing, we have the same problem of only knowing a single path to a given destination with no QoS information. If the path suggested by unicast routing does not have the resources to support the QoS desired, there are few choices available. Schemes that use an alternate route to "guess" at a better path have been suggested and can work for certain topologies but an underlying routing protocol that provides QoS information is necessary for a complete solution. As mentioned earlier, I-PNNI has the potential to provide enough information to compute paths for the requested QoS.

6.3 Mobile Routing

In developing an integrated IP-ATM network, potential new growth areas need to be included in the planning stages. One such area is mobile networking. Under the heading of mobile networks are included satellite extensions of the ATM cloud, mobile hosts that can join an IP subnetwork at random, and a true mobile network in which all network components including routers and/or switches are mobile.

The IP-ATM real-time service environment must be extended to include mobile networks so as to allow mobile users to access the same services as fixed network users. In doing so, a number of problems exist that need to be addressed. The principle problems are that mobile networks have constrained bandwidth compared to fiber and mobile links and are less stable than fixed fiber links. The impact of these limitations affect IP and ATM differently. In introducing one or more constrained components into the ATM cloud,the effects on congestion control in the overall network are unknown. One can envision significant buffering problems when a disadvantaged user on a mobile link attempts to access information from a high speed data stream. Likewise, as ATM uses out of band signalling to set up the connection, the stability of the mobile links that may have significant fading or complete loss of connectivity could have a significant effect on ATM performance.

For QoS, fading on a link will appear as a varying channel capacity. This will result in time-dependent fluctuations of available links to support a level of service. Current routing protocols are not designed to operate in a rapidly changing topology. QoS routing protocols that can operate in a rapidly changing topology are required and need to be developed.

7.0 Security Issues

In a quality of service environment where network resources are reserved, hence potentially depriving other users access to these resources for some time period, authentication of the requesting host is essential. This problem is greatly increased in a combined IP-ATM topology where the requesting host can access the network either through the IP or the ATM portion of the network. Differences in the security architectures between IP and ATM can lead to opportunities to reserve resources without proper authorization to do so. A common security framework over the combined IP-ATM topology would be desirable. In lieu of this, the use of trusted edge devices requesting the QoS services are required as a near term solution.

Significant progress in developing a common security framework for IP is underway in the IETF [2]. The use of authentication headers in conjunction with appropriate key management is currently being considered as a long range solution to providing QoS security [3,8]. In developing this framework, the reality of ATM portions of the Internet should be taken into account. Of equal importance, the ATM Forum ad-hoc security group should take into account the current work on an IP security architecture to ensure compatibility.

8.0 Future Directions

Clearly, there are some challenging issues for real-time IP-ATM services and some areas are better understood than others. For example, mechanisms such as policing, admission control and packet or cell scheduling can be dealt with mostly independently within IP or ATM as appropriate. Thus, while there may be hard problems to be solved in these areas that need to be addressed in either the IP or ATM communities, there are few serious problems that arise specifically in the IP over ATM environment. This is because IP does not particularly care what mechanisms a network element (such as an ATM network) uses to provide a certain QoS; what matters is whether the ATM service model is capable of offering services that can support the end-to-end IP service model. Most of the hard problems for IP over ATM therefore revolve around the service models for IP and ATM. The one piece of mechanism that is important in an IP/ATM context is signalling or resource reservation, a topic we return to below.

The following paragraphs enumerate some of the areas in which we believe significant work is needed. The work falls into three areas: extending the IP over ATM standards; extensions to the ATM service model; and extensions to the IP service model. In general, we expect that practical experience with providing IP QoS over ATM will suggest more enhancements to the service models.

We need to define ways of mapping the QoS and traffic characterizations (Tspecs and Rspecs) of IP flows to suitable characterizations for ATM connections. An agreement is needed so that some sort of uniform approach is taken. Whatever agreement is made for such mappings, it needs to be done so that when traversing several networks, the requested QoS is obtained end-to-end (when admission is possible). Practical experience should be gained with these mappings to establish that the ATM service classes can in fact provide suitable QoS to IP flows in a reasonably efficient way. Enhancement of the ATM service classes may be necessary, but experience is needed to determine what is appropriate.

We need to determine how the resource reservation models of IP (RSVP and ST-II) interact with ATM signalling. Mechanisms for establishing appropriate connection state with suitable QoS in ATM networks that are part of a larger integrated services Internet need to be defined. It is possible that the current IP/ATM mechanisms such as ARP servers and MARS can be extended to help to manage this state.

There is a need for better QoS routing. While this functionality is needed even in the pure ATM or pure IP environment, there is also an eventual need for integrated QoS routing between ATM and IP. Further research and practical experience is needed in the areas of QoS routing in IP in order to support more than the shortest best-effort path, especially when this path may traverse ATM networks. In many IP networks, there are multiple paths between a given source and destination pair but current routing technologies only pay attention to the current shortest path. As resources on the shortest path are reserved, it will be necessary and viable to explore other paths in order to provide QoS to a flow.

Enrichment of the ATM model to support dynamic QoS would greatly help the IP over ATM situation. At present, the QoS objectives for ATM are established at call set-up and then fixed for the duration of a call. It would be advantageous to have the ability to provide a dynamic QoS in ATM, so that an existing call could be modified to provide altered services.

Another possible area of enhancement to the ATM service model is in the area of multicasting. The multicast QoS offered is equal for all receivers, and thus may be determined by the least favorable path through the tree or by the most demanding receiver. Furthermore, there is no current provision for multipoint to multipoint connections. This limitation may rule out some of the services envisioned in the IP service model.

There are areas of potential enrichment of the IP model as well. While the receiver-based approach of RSVP has nice scaling properties and handles receiver heterogeneity well, it is not clear that it is ideal for all applications or for establishing state in ATM networks. It is possible that a sender-oriented mode for RSVP might ease the IP/ATM integration task.

Since the widespread availability of QoS raises new security concerns (e.g., denial of service by excessive resource reservation), it seems prudent that the IP and ATM communities work closely to adopt compatible approaches to handling these issues.

This list is almost certainly incomplete. As work progresses to define IP over ATM standards to support QoS and to implement integrated services internetworks that include ATM, more issues are likely to arise. However, we believe that this paper has described the major issues that need to be taken into consideration at this time by those who are defining the standards and building implementations.

9.0 References

1. Armitage, G., "Support for Multicast over UNI 3.1 based ATM Networks", Work in Progress, Bellcore, February 1995.

2. Atkinson, R., "Security Architecture for the Internet Protocol", RFC 1825, NRL, August 1995.

3. Atkinson, R., "IP Authentication Header", RFC 1826, NRL, August 1995.

4. Ballardie, A., and J. Crowcroft, "Multicast-Specific Security Threats and Counter-Measures", Proceedings of ISOC Symposium on Network and Distributed System Security, San Diego, Feb. 1995, pp. 2-16.

5. Ballardie, T., Jain, N., Reeve, S. "Core Based Trees (CBT) Multicast, Protocol Specification", Work In Progress, University College London, Bay Networks, June, 1995.

6. Braden, R., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, ISI/MIT/Xerox PARC, July 1994.

7. Braden, R., Zhang, L., Estrin, Herzog, D., and S. Jamin, "Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification", Work in Progress, ISI/PARC/UCS, July 1995.

8. Braden, R., Clark, D., Crocker, S., and C. Huitema, "Report of IAB Workshop on Security in the Internet Architecture", RFC 1636, ISI, MIT, TIS, INRIA, June 1994.

9. Callon, R., and B. Salkewicz, An Outline for Integrated PNNI for IP Routing", ATM Forum/ 95-0649, Bay Networks, July 1995.

10. Cole, R., Shur, D., and C. Villamizar, "IP over ATM: A Framework Document", Work in Progress, AT&T Bell Laboratories/ ANS, April 1995.

11. Deering, S., "Host Extensions for IP Multicasting", STD 5, RFC 1112, Stanford University, August 1989.

12. Delgrossi, L., and L. Berger, Editors, "Internet Stream Protocol Version 2 (ST-2) Protocol Specification - Version ST2+", RFC 1819, ST2 Working Group, August 1995.

13. Dykeman, D., Ed., "PNNI Draft Specification", ATM Forum/94-0471R8, IBM Zurich Research Lab, May 1995.

14. Goyal, P., Lam, S., and Vin, H., "Determining End-to-End Delay Bounds in Heterogeneous Networks," 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April, 1995.(Available via URL http://www.cs.utexas.edu/users/dmcl)

15. Laubach, M., "Classical IP and ARP over ATM", RFC 1577, HP, January 1994.

16. Moy, J., "OSPF Version 2", RFC 1583, Proteon, March 1994.

17. Moy, J., "Multicast Extensions to OSPF," RFC 1584, Proteon, March 1994.

18. Partridge, C., "A Proposed Flow Specification", RFC 1363, BBN, September 1992.

19. Perez, M., Liaw, F., Mankin, A., Hoffman, E., Grossman, D. and A. Malis, "ATM Signaling Support for IP over ATM", RFC 1755, ISI, Fore, Motorola Codex, Ascom Timeplex, February 1995.

20. Perkins, D., and Liaw, Fong-Ching, "Beyond Classical IP-Integrated IP and ATM Architecture Overview", ATM Forum/94-0935, Fore Systems, September 1994.

21. Perkins, D. and Liaw, Fong-Ching, "Beyond Classical IP-Integrated IP and ATM Protocol Specifications", ATM Forum/94-0936, Fore Systems, September 1994.

22. Romanow, A., and S. Floyd, "The Dynamics of TCP Traffic over ATM Networks", Proceedings of ACM SIGCOMM U94, London, August 1994, pp.79-88.

23. Shenker, S., and C. Partridge. "Specification of Guaranteed Quality of Service", Work in Progress, Xerox/BBN, July 1995.

24. Shenker, S., and C. Partridge. "Specification of Predictive Quality of Service", Work in Progress, Xerox/BBN, March 1995.

25. Shenker, S., C. Partridge and J. Wroclawski. "Specification of Controlled Delay Quality of Service", Work in Progress, Xerox/BBN/MIT, June 1995.

26. Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-time Applications", Work in Progress, GMD/ISI/Xerox/LBL, March 1995.

27. Topolcic, C., "Experimental Internet Stream Protocol, Version 2 (ST-II)", RFC 1190, BBN, October 1990.

28. Wang, Z., and J. Crowcroft, "QoS Routing for Supporting Resource Reservation", University College of London white paper, 1995.

10. Authors' Addresses

Eric S. Crawley
Marty Borden
Bay Networks
3 Federal Street
Billerica, Ma 01821
508-670-8888
esc@baynetworks.com
mborden@baynetworks.com

Bruce S. Davie
Bellcore
445 South Street
Morristown, New Jersey 07960-6438
201-829-4838
bsd@bellcore.com

Stephen G. Batsell
Naval Research Laboratory
Code 5521
Washington, DC 20375-5337
202-767-3834
sgb@saturn.nrl.navy.mil