Internet Draft
Expires December 13, 1996
File: draft-ietf-issll-atm-support-00.ps
S. Berson
ISI
June 1996

IP Integrated Services Support in ATM

June 13, 1996

Status of Memo

Abstract

Integrated Services in the Internet is rapidly becoming a reality. Meanwhile, ATM technology is appearing in the marketplace. It is an important problem to integrate ATM networks into this emerging Integrated Services Internet. "Classical" IP over ATM[11, 8] is now widely deployed, effectively solving the problem of "best effort service" in internets with ATM links. A key remaining issue is to provide the Internet Integrated Services (IIS) model over internets with ATM links. This draft provides guidelines for using ATM VCs with QoS as part of an Integrated Services Internet.

1 Introduction

[NOTE: This draft discusses RSVP and ATM specifically, but most of the issues and approaches are not specific to either RSVP or ATM, but rather would apply to any setup protocol and any Non-Broadcast Multiple Access (NBMA) network. In those cases where an issue is specific to either RSVP or ATM, it will be pointed out in the text. It is the intention of the author in the future to split this draft into two documents, one dealing with RSVP-specific issues, and the other with general integrated services issues.]

The Internet currently has one class of service normally referred to as "best effort." This service is typified by first-come, first-serve scheduling at each hop in the network. Best effort service has worked exceptionally well for electronic mail, World Wide Web (WWW) access, file transfer (e.g. ftp), etc. For real-time traffic such as voice and video, the current Internet has performed well only across unloaded networks. In order to provide guaranteed quality real-time traffic, new classes of service and new protocols are being introduced in the Internet[10, 4], while retaining the existing best effort service. Signalling for these new services is done by RSVP[3, 12], the Resource ReSerVation Protocol.

Due to extensive telephone company support, ATM is rapidly becoming an important link layer technology. As this draft was being written, the current version of ATM signalling was UNI 3.1[9], and UNI 3.1 is expected to be predominant in the marketplace in the near term. However, a new version (UNI 4.0) is expected to be standardized soon, and this draft will be updated as appropriate.

One of the important features of ATM technology is the ability to request a point-to-point Virtual Circuit (VC) with a specified Quality of Service (QoS). An additional feature of ATM technology is the ability to request point-to-multipoint VCs with a specified QoS. Point-to-multipoint VCs allows leaf nodes to be joined and removed from the VC dynamically and so provide a mechanism for supporting IP multicast. It is only natural that RSVP and the Internet Integrated Services (IIS) model would like to utilize the QoS properties of any underlying link layer, and this draft concentrates on ATM.

Classical IP over ATM[11] has solved part of this problem, supporting IP unicast traffic over ATM. Classical IP over ATM is based on a Logical IP Subnetwork (LIS), which is a separately administered IP subnetwork. Hosts within a LIS communicate using the ATM network, while hosts from different subnets communicate only by going through an IP router (even though it may be possible to open a direct VC between the two hosts over the ATM network). Classical IP over ATM provides an Address Resolution Protocol (ATMARP) for ATM edge devices to resolve IP addresses to native ATM addresses. For any pair of IP/ATM edge devices (i.e. hosts or routers), a single VC is created on demand and shared for all traffic between the two devices. A second part of the RSVP and IIS over ATM problem, IP multicast, is close to being solved with MARS[1], the Multicast Address Resolution Server. The MARS server generalizes ATMARP by allowing an IP address to resolve into a list of native ATM addresses, rather than just a single address.

There is work in progress on alternative IP over ATM models (e.g. NHRP[7]) but similar solutions to those described here for the classical model are still expected to apply. But further study is needed.

A key remaining issue for IP over ATM is the integration of RSVP signalling and ATM signalling in support of the Internet Integrated Services (IIS) model. There are two main areas involved in supporting the IIS model, QoS translation and VC management. QoS translation concerns mapping a QoS from the IIS model to a proper ATM QoS, while VC management concentrates on how many VCs are needed and which traffic flows are routed over which VCs. This draft concentrates on VC management, and we assume that the QoS for a single RSVP flow can be acceptably translated to an ATM QoS. This assumption is reasonable since the traffic parameters for both IIS and ATM QoS are very similar.

The remainder of this draft is organized as follows. Section 2 describes how source traffic is mapped to RSVP flows and then to VCs. Section 3 describes the "single homogeneous VC" model where there is only one VC per RSVP flow even when there are different receivers requesting different reservations. In particular, this section describes how the QoS of the single VC must be dynamically changed over ATM and how to limit excessive signalling using timers. Section 4 examines the case where a single RSVP flow can be fed into several VCs. In particular, this section describes the need for multiple VCs particularly for simultaneously supporting a single QoS service as well as best effort service. Section 5 describes alternatives for VCs used for RSVP signalling messages. Section 6 discusses some additional minor issues. Section 7 concludes the draft and describes continuing and future work.

2 RSVP

[NOTE - this section is RSVP specific]

RSVP[3, 12] provides many important features for signalling in the Internet including receiver initiated reservations, "soft state" in routers, heterogeneous reservations within a multicast session, dynamic change of reservations, and multiple reservation styles or modes of sharing. Of particular interest for RSVP and ATM[1] are reservation styles, dynamic change of reservations, and heterogeneous reservations within a multicast session. The remainder of this section describes these features.

2.1 RSVP messages

In RSVP resources are reserved by use of two types of messages, **PATH** and **RESV**. Each RSVP traffic source sends out PATH messages along the same route (unicast or multicast) as the data traffic will travel. PATH messages carry information about the source traffic parameters such as a mean bit rate, and token bucket size. Using information from PATH messages and obtained from higher layer protocols, receivers can then make reservations for specific resources by sending RESV messages along the reverse path of the PATH messages. Resources will be allocated on the proper outgoing link of all nodes along the route from traffic source(s) to the receiver. This is shown in figure 1 with two sources "S1" and "S2," and two receivers "R1" and "R2."

[1] Though ATM is sender oriented, it is worth noting that receiver initiated reservations were not an issue at all.

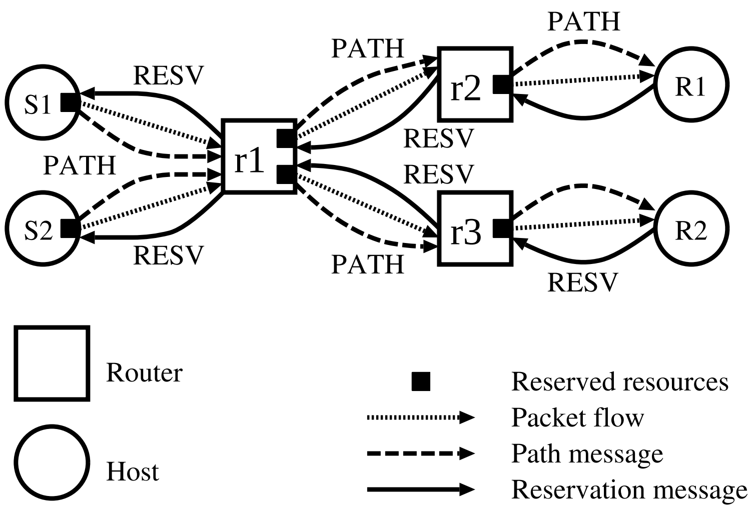

Figure 1: RSVP message flow

Multiple receivers can send RESV messages toward the same traffic source(s) as shown in figure 1. The quantities of resources requested by these receivers may be different. RSVP handles RESV messages for different quantities of resources by doing *merging*. A node where multiple different reservation messages arrive (from different receivers) is called a *merge node*. In order that RSVP scale well with the number of receivers of a session, only one merged RESV message is forwarded from a merge node towards a source, and this message carries the maximum of the incoming resource reservation requests. In figure 1, router "r1" is a merge node and would send to traffic sources "S1" and "S2" a reservation message requesting the maximum of the resources requested by receivers "R1" and "R2."

2.2 Reservation Styles

RSVP provides several different styles of reservations. The styles that are part of the current version 1 RSVP specification are shown in figure 2, while noting that additional styles may be defined in the future. The reservation style helps determine how many RSVP flows are needed for a multicast group and which source traffic uses which flow. Reservation styles can be "distinct", where each RSVP flow is used by exactly one sender, or "shared", where multiple senders use the same RSVP flow. Reservation styles can also have either "explicit sender selection", where a reservation is established only for senders explicitly listed in the reservation message, or can have "wildcard sender selection" where traffic from any sender is selected. Different reservation styles are suitable for different types of data traffic. Typically, shared type styles are best for audio conferencing since there would typically be only one person speaking at a time independent of the

Sender Selection	Reservations: Distinct	Reservations: Shared
Explicit	Fixed-Filter (FF) Style	Shared-Explicit (SE) Style
Wildcard	(None defined)	Wildcard-Filter (WF) Style

Figure 2: Reservation Styles

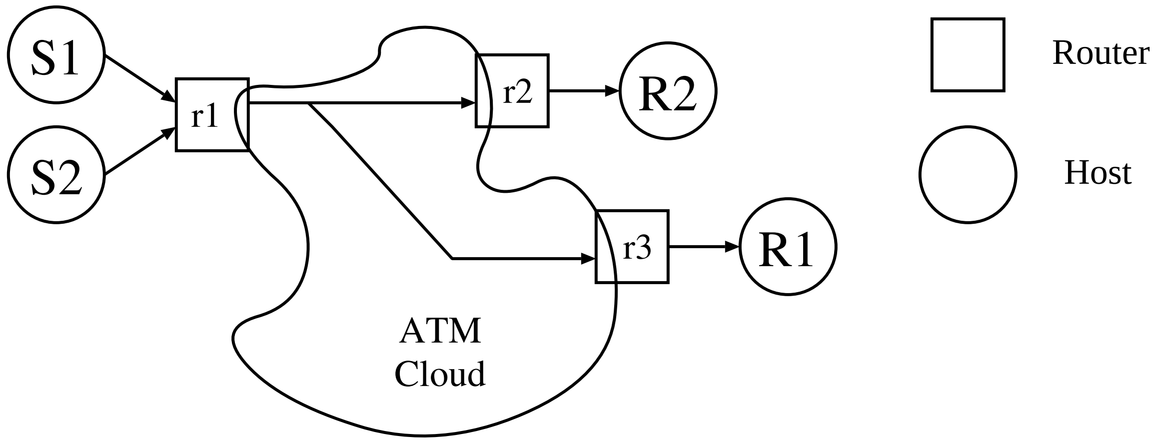

Figure 3: IP over ATM network

number of senders. Distinct type styles are typically used for video where each traffic source is continuously sending, and so sharing is difficult.

Three reservation styles have currently been defined. Fixed filter style has explicit sender selection and distinct reservations and is typically used for video conferencing where each video stream has its own RSVP flow. Wildcard filter style has wildcard sender selection and shared style and is typically used for audio conferencing where people naturally take turns speaking. Shared explicit style has explicit sender selection and shared reservations and is similar to wildcard style except it provides better security by explicitly listing senders. Currently, there is no style defined with wildcard sender selection and distinct reservations.

2.3 RSVP flows

As an example, consider a system with two senders (S1 and S2), two receivers (R1 and R2), and three routers (r1, r2, and r3) as shown above in figure 3. Assuming that the two receivers are requesting the same QoS, then the number of RSVP flows at the ATM edge device "r1" needed to support different reservation styles for this configuration can be determined. For a wildcard filter style, only one (point-to-multipoint) flow is needed as all senders for this group can use the shared reservation across the ATM network. For styles with explicit sender selection, assume that both senders are selected. For fixed filter style, three flows are needed, one QoS flow for each of the two senders explicitly listed and one best effort service flow to be shared among any other senders. For the shared explicit style, two flows are needed, one QoS flow to be shared among all the (i.e. two) explicitly listed senders, and one best effort service flow to be shared among the remaining senders.

More generally, the following numbers of RSVP flows are created for different filter styles. For a reservation with wildcard filter (WF) style, there is only one RSVP flow, since all senders to a (multicast or unicast) session are part of the same reservation. For shared explicit (SE) style, there are two RSVP flows, one for the senders explicitly listed who receive reserved service, and one for any other senders who receive ordinary best effort service. Finally, the number of RSVP flows created by a fixed filter (FF) style reservation depends on the number of senders listed in the reservation message. If there are n senders listed in the message, then there are $n+1$ RSVP flows created, one for each of the n senders, and one additional flow for all senders that are not listed.

2.4 RSVP flows and VCs

There are different approaches to mapping RSVP flows to ATM VCs. Two approaches are discussed in this draft, while a third remains for future work. Section 3 examines the VC per flow model, where each RSVP flow is mapped to a single ATM VC. This ATM VC can be a point-to-point or point-to-multipoint as appropriate. Section 4 examines the multiple VCs per RSVP flow model where a single flow can be forwarded into several VCs each with a different QoS. Other models allowing aggregation of RSVP flows into VCs (VC for multiple flows model) are being studied and will be briefly discussed in section 7.

3 Single VC per RSVP Flow

One approach for mapping RSVP flows into VCs is to have a single (point-to-point or point-to-multipoint) VC for each RSVP flow. The potential problem with this scheme is that different receivers might request different qualities of service. In the "single VC per RSVP flow" model, this heterogeneity problem is handled by using the maximum of the requested resources of all the receivers of a session. The remainder of this section discusses the issue of dynamically changing the QoS of a VC. The following section allows more heterogeneity in reservations by using additional VCs.

3.1 Dynamic QoS

RSVP provides dynamic quality of service (QoS) in that the resources that a reservation requests may change at any time. There are several common reasons for a change of reservation QoS. First, an existing receiver can request a new larger (or smaller) QoS if the current received quality is unacceptable. Second, a sender may change its traffic specification (TSpec), which can trigger a change in the reservation requests of the receivers. Third, a new sender can start sending to a multicast group with a larger traffic specification than existing senders, triggering larger reservations. Finally, a new receiver can make a reservation that is larger than existing reservations. If the merge node for the larger reservation is an ATM edge device, a new larger reservation must be set up across the ATM network.

Since ATM service, as currently defined, does not allow renegotiating the QoS of a VC, dynamically changing the reservation means tearing down an established VC, and creating a new VC with the new QoS. Tearing down a VC and setting up a new VC in ATM are complex operations that involve a non-trivial amount of processor time, and have a substantial latency.

To prevent excessive signalling load on an ATM network, timers can be used. Each ATM edge device is configured with a time parameter τ that gives the minimum amount of time the edge device will wait between successive changes of the QoS of a particular VC. Thus if the QoS of a VC is changed at time t, all messages that would change the QoS of that VC that arrive before time $t + \tau$ would be queued. If several messages changing the QoS of a VC arrive during the interval, redundant messages can be discarded. At time $t + \tau$, the remaining change(s) of QoS, if any, can be executed. This timer approach would apply more generally to any network structure, and might be worthwhile to incorporate into RSVP.

The sequence of events for a single VC would be

1. Wait if timer is active
2. Establish VC with new QoS
3. Remap data traffic to new VC
4. Tear down old VC
5. Activate timer

3.2 Limitations

There are two major problems with the single VC per RSVP flow approach. The first problem is that a user making a small or no reservation would get a "free ride" across the ATM network on any person(s) making a larger reservation. The second problem is that a user may want to join an existing VC at the established QoS level, but, because of a lack of resources, not be able to join. The rejected user would still like to receive the traffic on a best effort basis, which is the standard method of operation in the Internet. Preserving the homogeneous single VC per RSVP flow model in this case would mean tearing down the QoS VC, and replacing it with a single best effort VC. Clearly it is unacceptable to tear down one customer's existing QoS reservation because a second customer was not able to join the existing VC. A solution to both the "free ride" problem and the best effort service problem involves having multiple VCs carrying identical traffic which is the subject of the next section.

4 Multiple VCs per RSVP Flow

The previous "single VC per RSVP flow" model had the advantage that each byte of data traffic was sent over the ATM network exactly once and each leaf of the VC received identical service. This homogeneous model is simple but does not take into account the situation that different users may demand (and be willing to pay for) different levels of service than others. This section discusses solutions to heterogeneous reservation requests from different receivers involving multiple VCs per RSVP flow. Two different approaches are discussed. The first approach requires at most two VCs per RSVP flow, one for best effort traffic, and the other for (homogeneous) QoS traffic. The other approach has one VC per QoS requested, and so potentially has an unlimited number of VCs per RSVP flow.

4.1 Two VCs per RSVP flow

RSVP supports heterogeneous QoS, meaning that different receivers of the same multicast group can request a different QoS. But most importantly, some receivers might have no reservation at all, but want to receive the traffic on a best effort service basis. The IP model allows receivers to join a multicast group at any time on a best effort basis, and it is important that ATM as part of the Internet continue to provide this service. We define the "limited heterogeneity" model as the case where the receivers of a multicast session are limited to use either best effort service or a single alternate quality of service. The alternate QoS can be chosen either by higher level protocols or by dynamic renegotiation of QoS as described in the previous section.

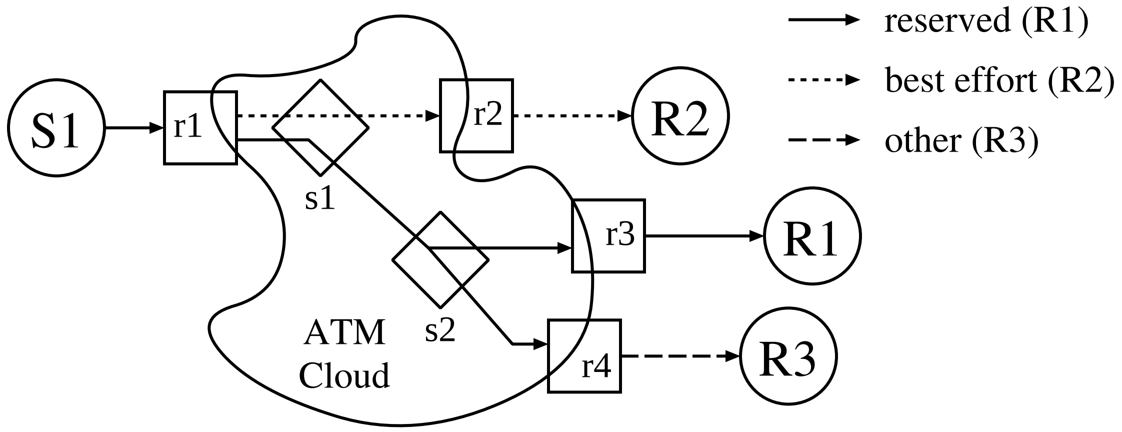

Figure 4: Limited heterogeneity

In order to support limited heterogeneity, each ATM edge device participating in a session would need at most two VCs. One VC would be a point-to-multipoint best effort service VC and would serve all best effort service IP destinations for this multicast group. The other VC would be a point to multipoint VC with QoS and would serve all IP destinations for this multicast group that have an RSVP reservation established. This is shown in figure 4 where there are three receivers, R2 requesting best effort service, while R1 and R3 request distinct reservations. One point-to-multipoint VC is set up for best effort traffic which serves R2. Another VC is set up for the QoS traffic to receivers R1 and R3, using the maximum of R1 and R3's reservation. Note that though the VC and hence the QoS for R1 and R3 are the same within the ATM cloud, the reservation outside the ATM cloud (from router r4 to receiver R3) uses the QoS actually requested by R3.

The disadvantage of the limited heterogeneity scheme is that each packet will need to be duplicated at the network layer and one copy sent into each of the 2 VCs. Looking at figure 4, there are two VCs going from router r1 to switch s1. Two copies of every packet will traverse the r1-s1 link. Depending on the network topology and group membership, there may be a large amount of duplicate traffic flowing over the ATM links.

Note that two separate VCs for each session will not normally be needed. First, if all the receivers are making reservations, no best effort VC is needed. Second, the best effort traffic for all sessions with the same ATM destinations can be multiplexed on the same best effort VC. In the ideal case, there will be only one best effort VC for all the best effort traffic for all sessions on a single node. This will limit the number of VCs needed.

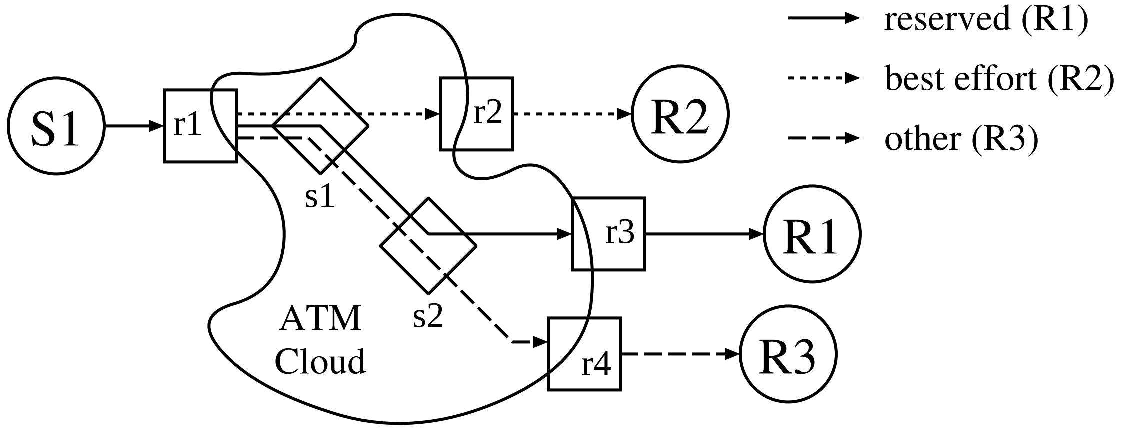

Figure 5: Limited heterogeneity

4.2 Many VCs per RSVP flow

Instead of arbitrarily restricting an RSVP flow to at most two VCs, rules can be specified as to when a new QoS VC can be created (e.g. when the user is willing to pay). This gives a lot of flexibility to a service provider. Full heterogeneity is possible where a separate VC could be created for each distinct QoS for a multicast session.

While we advocate the limited heterogeneity approach as in section 4.1, different service providers may choose alternate approaches. RSVP allows for local policy control [6] as well as admission control. Thus a user can request a reservation with a specific QoS and with a policy object that, for example, offers to pay for additional costs setting up a new VC. The policy module at the entry to a service provider can decide how to satisfy that request - either by merging the request in with an existing VC or by creating a new VC for this (and perhaps other) users. This policy can be on a user-provider basis where a user and a provider have an agreement on the type of service offered, or on a provider-provider basis, where two providers have such an agreement. With the ability to do local policy control, service providers can provide services best suited to their own resources and their customers desires.

This is shown in figure 5. Whereas, in figure 4, R1 and R3 shared the same VC across the ATM network; in figure 5, R1 and R3 have a separate VC, so each receives precisely the resources requested.

Note that while full heterogeneity gives users exactly what they request, it requires more resources of the network than limited heterogeneity. In figure 5, three copies of each packet are sent on the link from r1 to s1. Two copies of each packet are then sent from s1 to s2. As with limited heterogeneity, the exact amount of bandwidth dedicated to duplicate traffic depends on the network topology and group membership.

4.3 Making a reservation

For the limited heterogeneity case, making an RSVP reservation will mean that a leaf of a point to multipoint VC will need to leave the best effort service VC and join as a new leaf of the QoS VC. If no QoS VC exists, a new QoS VC is created with the receiver as a leaf. If the new QoS leaf cannot be created, an error message will be sent and the user will continue receiving best effort service. If there is a QoS VC, but the QoS needs to be increased for the new reservation, a new VC with the larger QoS will be requested (for all QoS receivers). If the new VC request fails, the old QoS VC will remain, the new reservation will fail, and the new user will continue to receive best effort service. An RSVP reservation teardown will mean leaving the QoS VC and joining the best effort service VC. If no best effort VC exists, then one would be created.

For the full heterogeneity model, making an RSVP reservation is similar to the limited heterogeneity case. The difference is that a change in reservation may attempt to switch a leaf from one QoS VC to another QoS VC.

5 RSVP VCs

[NOTE - this section is RSVP specific]

One last important issue is providing a data path for the RSVP messages themselves. As mentioned in section 2, there are two main types of messages, PATH and RESV. PATH messages are sent to a multicast address, while RESV messages are sent to a unicast address. Other RSVP messages are handled similar to either PATH or RESV. So ATM VCs used for RSVP signalling messages need to provide both unicast and multicast functionality.

There are several different approaches for how to assign VCs to use for RSVP signalling messages. The main approaches are:

- use same VC as data
- single VC per session
- single point-to-multipoint VC multiplexed among sessions
- multiple point-to-point VCs multiplexed among sessions

There are several different issues that affect the choice of how to assign VCs for RSVP signalling. One issue is the number of additional VCs needed for RSVP signalling. Related to this issue is the degree of multiplexing on the RSVP VCs. In general more multiplexing means less VCs. An additional issue is the latency in dynamically setting up new RSVP signalling VCs. A final issue is complexity of implementation. The remainder of this section discusses the issues and tradeoffs among these different approaches and suggests guidelines for when to use which alternative.

5.1 Mixed data and control traffic

In this scheme RSVP signalling messages are sent on the same VCs as is the data traffic. The main advantage of this scheme is that no additional VCs are needed beyond what is needed for the data traffic. An additional advantage is that there is no ATM signalling latency for PATH messages (which follow the same routing as the data messages). However there can be a major problem when data traffic on a VC is nonconforming. With nonconforming traffic, RSVP signalling messages may be dropped. While RSVP is resilient to a moderate level of dropped messages, excessive drops would lead to repeated tearing down and re-establishing QoS VCs, a very undesirable behavior for ATM. Due to these problems, this is not a good choice for providing RSVP signalling messages, even though the number of VCs needed for this scheme is minimized.

One variation of this scheme that is hopeful but requires further study is to have a packet scheduling algorithm (before entering the ATM network) that gives priority to the RSVP signalling traffic. In this way, the ATM Cell Loss Priority (CLP) bit could be used to make sure that RSVP signalling messages only rarely get dropped.

5.2 Single RSVP VC per RSVP Flow

In this scheme, there is a parallel RSVP signalling VC for each RSVP flow. This scheme results in twice the minimum number of VCs, but means that RSVP signalling messages have the advantage of a separate VC. This separate VC means that RSVP signalling messages have their own traffic contract and compliant signalling messages are not subject to dropping due to other noncompliant traffic (such as can happen with the scheme in section 5.1). The advantage of this scheme is its simplicity - whenever a data VC is created, a separate RSVP signalling VC is created. The disadvantage of the extra VC is that extra ATM signalling needs to be done.

This scheme requires twice the minimum number of VCs and also additional latency, but is quite simple. This approach would tend to work well on hosts.

5.3 Multiplexed point-to-multipoint RSVP VCs

In this scheme, there is a single point-to-multipoint RSVP signalling VC for each unique ingress router and unique set of egress routers. This scheme allows multiplexing of RSVP signalling traffic that shares the same ingress router and the same egress routers. This can save on the number of VCs, by multiplexing, but there are problems when the destinations of the multiplexed point-to-multipoint VCs are changing. Several alternatives exist in these cases, that have applicability in different situations. First, when the egress routers change, the ingress router can check if it already has a point-to-multipoint RSVP signalling VC for the new list of egress routers. If the RSVP signalling VC already exists, then the RSVP signalling traffic can be switched to this existing VC. If no such VC exists, one approach would be to create a new VC with the new list of egress routers. Other approaches include modifying the existing VC to add an egress router or using a separate new VC for the new egress routers. When a destination drops out of a group, an alternative would be to keep sending to the existing VC even though some traffic is wasted.

The number of VCs used in this scheme is a function of traffic patterns across the ATM network, but is always less than the number used with the Single RSVP VC per data VC. In addition, existing best effort data VCs could be used for RSVP signalling. Reusing best effort VCs saves on the number of VCs at the cost of higher probability of RSVP signalling packet loss. One possible place where this scheme will work well is in the core of the network where there is the most opportunity to take advantage of the savings due to multiplexing.

5.4 Multiplexed point-to-point RSVP VCs

In this scheme, multiple point-to-point RSVP signalling VCs are used for a single point-to-multipoint data VC. This scheme allows multiplexing of RSVP signalling traffic but requires the same traffic to be sent on each of several VCs. This scheme is quite flexible and allows a large amount of multiplexing. Since point-to-point VCs can set up a reverse channel at the same time as setting up the forward channel, this scheme could save substantially on signalling cost. In addition, signalling traffic could share existing best effort VCs. Sharing existing VCs reduces the total number of VCs needed, but might cause signalling traffic drops if there is congestion in the ATM network.

This pt-pt scheme would work well in the core of the network where there is much opportunity for multiplexing. Also in the core of the network, RSVP VCs can stay permanently established either as Permanent Virtual Circuits (PVCs) or as long lived Switched Virtual Circuits (SVCs). The number of VCs in this scheme will depend on traffic patterns, but in the core of a network would be approximately $n(n-1)/2$ where n is the number of IP nodes in the network. In the core of the network, this will typically be small compared to the total number of VCs.

5.5 QoS for RSVP VCs

There is an issue for what QoS, if any, to assign to the RSVP VCs. Two solutions have been covered in section 5.1 and in the shared best effort VC variation in section 5.4. For other RSVP VC schemes, a QoS (possibly best effort) will be needed. What QoS to use partially depends on the expected level of multiplexing that is being done on the VCs, and the expected reliability of best effort VCs. Since RSVP signalling is infrequent (typically every 30 seconds), only a relatively small QoS should be needed. This is important since using a larger QoS risks the VC setup being rejected for lack of resources. Falling back to best effort when a QoS call is rejected is possible, but if the ATM net is congested, there will likely be problems with with RSVP packet loss on the best effort VC also. Additional experimentation is needed in this area.

6 Additional issues

There is an issue with VCs timing out as described in RFC1755, section 3.4 on VC Teardown. RFC1755 recommends tearing down a VC that is inactive for a certain length of time - 20 minutes is recommended. This timeout could mean that a valid VC with QoS that has been established with RSVP might be torn down unexpectedly. While this behavior is acceptable for best-effort traffic, it is important that if a reservation has been established on a VC, that the VC not be torn down. If there is no choice about the VC being torn down, the RSVP daemon must be notified, so a reservation failure message can be sent.

Most of the previous discussion has been concerned with meshes of point-to-multipoint connections. An alternate approach is to use a MultiCast Server[1] (MCS). An MCS provides a relay node that forwards the multicast (e.g. PATH) messages. It is expected that all the previous discussion applies for an MCS as well as meshes, but further experimentation is needed.

7 Conclusions and Future Work

We have described a scheme for deploying "classical" RSVP over IP over ATM. We have a prototype system using the limited heterogeneity model running at ISI and we are experimenting with it now. There are a number of other issues that are subjects of continuing research. These issues (and others) are covered in [2], and are briefly repeated here.

One key issue is how to translate an Internet Integrated Services (IIS) QoS to an ATM QoS. Fortunately, this issue seems to be getting easier as the IETF and ATM Forum QoS definitions are getting more similar as they evolve. However there are still potential differences in traffic shaping and policing models. Additional coordination between the IETF and the ATM Forum in testing and standardizing a QoS translation is desirable now.

A second issue is to consider the multiple RSVP flows per VC (or aggregation) model. With this model, large VCs could be set up between IP routers in an ATM network. These VCs could be managed the same as how IIS point-to-point links (e.g. T-1, DS-3) are managed now. Traffic from multiple sources over multiple multicast groups might be multiplexed on the same VC. This approach has a number of advantages. First, there is typically no signalling latency as VCs would be in existence when the traffic started flowing, so no time is wasted in setting up VCs. Second, the heterogeneity problem in full over ATM has been reduced to a solved problem. Finally, the dynamic QoS problem for ATM has also been reduced to a solved problem. The problem with this approach is that the choice of what QoS to use for which of the large VCs is difficult. If chosen poorly, there might be many VC setups failing while other bandwidth is unused. An additional problem is that the ability of ATM to do QoS dependent routing is wasted.

ATM is currently close to a new standard (UNI 4.0) that promises certain advantages over the current UNI 3.1 version. Current work in the ATM Forum and ITU promises additional advantages. A final issue is to keep current with changes in ATM, and to keep this document up-to-date. Some specific issues are covered in [2, 5].

References

[1] Armitage, G., "Support for Multicast over UNI 3.0/3.1 based ATM Networks," Internet Draft.

[2] Borden, M., Crawley, E., Krawczyk, J, Baker, F., and Berson, S., "Issues for RSVP and Integrated Services over ATM," Internet Draft.

[3] Braden, R., Zhang, L., Berson, S., Herzog, S., and Jamin, S., "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," Internet Draft.

[4] Clark, D., Shenker, S., and Zhang, L., "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanisms," SIGCOMM '92.

[5] Demirtjis, A., Berson, S., Edwards, B., Maher, M., Braden, B., and Mankin, A., "RSVP and ATM Signalling," ATM Forum Contribution 96-0258.

[6] Herzog, S., "Building Blocks for Accounting and Access Control in RSVP," Internet Draft.

[7] Katz, D., Piscitello, D., Cole, B., Luciani, J., "NBMA Next Hop Resolution Protocol (NHRP)," Internet Draft.

[8] Perez, M., Liaw, F., Grossman, D., Mankin, A., Hoffman, E., and Malis, A., "ATM Signalling Support for IP over ATM," RFC 1755.

[9] *ATM User-Network Interface (UNI) Specification - Version 3.1*, Prentice Hall.

[10] Braden, R., Clark, D., Shenker, S. "Integrated Services in the Internet Architecture: an Overview," RFC 1633, June 1994.

[11] Laubach, M., "Classical IP and ARP over ATM," RFC1577, January 1994.

[12] Zhang, L., Deering, S., Estrin, D., Shenker, S., Zappala, D., "RSVP: A New Resource ReSerVation Protocol," IEEE Network, September 1993.

Security Considerations

Security considerations have not been addressed in this draft.

Author's Address

```
Steven Berson
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292

Phone: (310) 822-1511
EMail: berson@isi.edu
```