# Internet Integrated Service over ATM: A Solution for Shortcut QoS Virtual Channels

*Roberto Cocca and Stefano Salsano, CoRiTeL — Consorzio di Ricerca sulle Telecomunicazioni*
*Marco Listanti, University of Rome "La Sapienza"*

**ABSTRACT** In principle, the interaction of RSVP and ATM should allow the IP level to benefit from some features of the ATM layer. The most interesting one is the native support of end-to-end quality of service provided by ATM. On the other hand, there are issues that must be clarified to define correct interworking: for example, the possible overlapping between the mechanisms used in the IP and ATM levels to support QoS, or the needed IP/ATM address resolution mechanism. This article proposes a solution to exploit ATM shortcut VCs supporting QoS in the Internet integrated services model. A straightforward enrichment to the RSVP protocol is defined, which only impacts the devices (hosts and routers) involved in the shortcut procedure. A mechanism for IP/ATM address resolution is provided, avoiding the use of other mechanisms such as NHRP. Special care has been taken to maintain compatibility with "traditional" RSVP hosts and routers.

Despite its tremendous growth, the Internet is still largely based on a very simple service model, called *best effort*, providing no guarantee on the correct and timely delivery of data packets. This simplicity has probably been one of the main reasons for the success of IP technology. The best-effort service model, combined with an efficient transport layer protocol (i.e., TCP), is perfectly suited for a large class of applications, referred to as *elastic*, which can adapt (even dynamically) to different performances offered by the network in terms of data throughput and end-to-end-delay. Web browsing and e-mail are typical examples of elastic applications. Real-time applications like video and audio conferencing typically require stricter guarantees on throughput and delay. The idea of extending Internet capabilities to provide support to real-time applications has led the Internet community to develop the Internet integrated services (IIS) architecture [1]. The guaranteed service and controlled load service models have been added to the best effort service model, and a signaling protocol called *Resource Reservation Protocol* (RSVP) [2] has been defined.

A second challenge to Internet technology is the aggregate switching throughput that should be provided by IP routers. Advances in optical technology allow higher and higher link bandwidths at decreasing costs. In this scenario, the processing capability of the routers could constitute the most relevant bottleneck. Although gigabit IP routers will be available in the near future, currently it is widely accepted that asynchronous transfer mode (ATM) technology offers potential advantages in both its capability in terms of aggregate switching throughput and its native ability to support quality of service (QoS) in both point-to-point and point-to-multipoint virtual channels (VCs). In fact, ATM is currently used as an efficient network technology for transport/switching in backbone networks.

A key point for the future of ATM is integration with IP technology. Current IP/ATM interworking solutions (LAN emulation, LANE [3]; classical IP over ATM, CLIP [4]; multiprotocol over ATM, MPOA [5]) do not natively support QoS. They focus on using ATM technology under the classical best-effort IP model.

This article investigates the interaction of the Internet integrated services (IIS) architecture and ATM. The rationale for this interaction is described in RFC 2382 [6]. According to this RFC, 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, and is extensively dealt with in [7]. This work concentrates on the issue of VC management, dealing in particular with shortcuts and address translation. VC management considers how many and what kinds of VCs are needed, and which traffic flows are routed over which VC. Different solutions are possible; for example, the first choice is between ATM permanent and switched virtual channels (PVCs, SVCs) or a mix of PVCs and SVCs. Another choice is related to the possibility of having a single VC for each IP flow or aggregating several flows in a single VC. Traditional hop-by-hop routing or shortcut VCs can be used.

It is worth spending some time on the topic of ATM shortcuts. Within the CLIP model, *logical IP subnetworks* (LISs) are defined as separately administered IP subnetworks. Hosts belonging to different LISs can communicate only by going through an IP router. An end-to-end path between two hosts can be composed of several router hops, even though it may be possible to open a direct VC between the two hosts over the ATM network. The possibility of using these direct VCs, referred to as *shortcut VCs*, is proposed in architectures such as MPOA [5]. The first problem to be solved in setting up a shortcut is to know the identity and ATM address of the remote end of the VC. Within MPOA a specific protocol, called *Next Hop Resolution Protocol* (NHRP) [8], allows the resolution of an IP address to an ATM address. Anyway, the basic mechanism allows redirection of all traffic directed to a specific IP address to a shortcut VC with no QoS requirements (e.g., using the ATM unspecified bit rate, UBR, transport class).

The interaction of NHRP with RSVP is still unclear. In the following section some critical issues related to this interaction are described. These considerations led us to propose a solution for the use of ATM shortcuts in the IIS architecture not based on NHRP. The solution avoids the use of IP/ATM address resolution mechanisms by a simple enrichment to the RSVP protocol. The fundamentals of this approach are described, and also the specification of the new classes and further details on the procedures to handle these classes in


0163-6804/99/$10.00 © 1999 IEEE

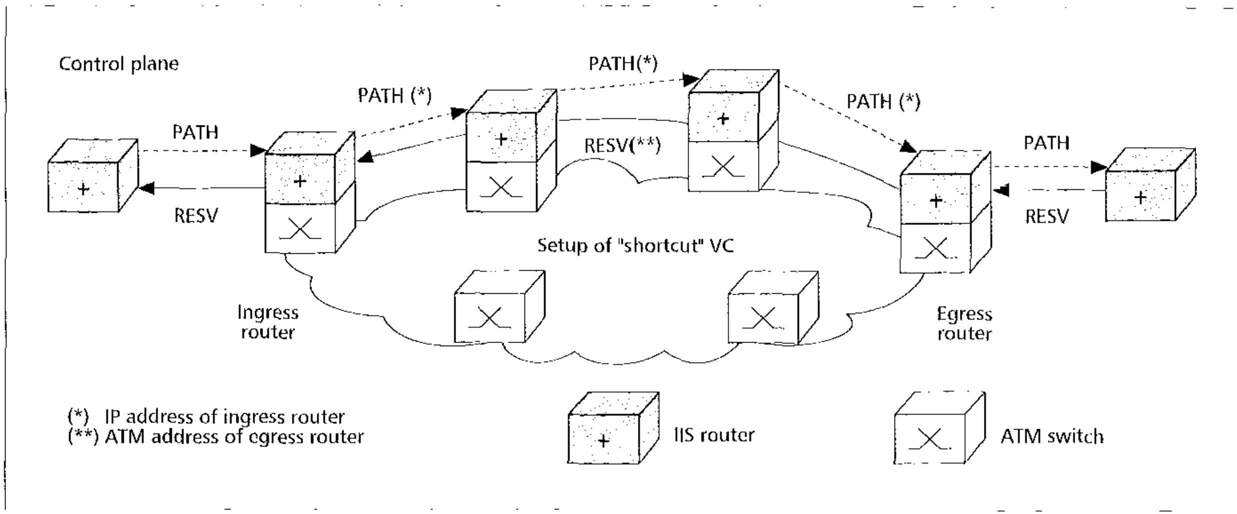


**Figure 1.** *Enrichments to RSVP for the support of ATM shortcuts.*

the routers. The article reports on a limited testbed implementation. The multicast case issues are listed, and finally the conclusions are given.

## Issues in the Interaction of NHRP and RSVP

NHRP, defined in [8], is used to establish unicast ATM VCs that bypass IP routers. NHRP provides for the mapping of an IP address to the corresponding ATM address.

There are some issues related to the use of NHRP: for example, the domino effect (the generation of multiple NHRP messages for the same data packet by subsequent routers), the possibility of stable routing loops in the router-to-router case, scalability issues, and burdensome management. These issues are analyzed in [8, 9].

Furthermore, new problems arise in considering a solution for the interworking of NHRP and RSVP. The interaction of these two approaches is pretty complex. NHRP is not QoS-oriented; for example, NHRP messages are not able to carry the traffic information to be used at each (NHRP) node as a "hint" to yield the "longest" possible QoS shortcut. Another issue is how to transport the RSVP control messages. These messages usually follow the hop-by-hop path according to the IP routing. When ATM shortcut VCs are available the RSVP message can use them, but side interaction with RSVP logic must be carefully considered.

A first approach could be to use NHRP and the ATM shortcuts whenever possible, including for transport of RSVP messages. The packet forwarding procedures should be carefully designed considering that different packets (best effort data packets for different destinations, packets of QoS flows, RSVP control messages) should go into different VCs. A different possible solution consists of using ATM shortcuts only for QoS flows. Best-effort packets should follow the hop-by-hop path. The NHRP procedures should start only after reception of the first RSVP RESV message of the session. A drawback of this solution is that unneeded resources in the intermediate nodes are reserved.

In the next sections an architectural solution for interworking between RSVP and ATM, making use of shortcut VCs, is proposed, but it does not make use of NHRP in order to overcome the problems listed above.

## Interaction of RSVP and ATM for the Support of Shortcut QoS VCs

Two new information elements (*classes* in RSVP terminology) must be added to the RSVP PATH and RESV messages. The related procedures to handle these classes will be described. Special care has been taken to maintain compatibility with "traditional" RSVP hosts and routers. No modifications are needed in the host and in the routers that do not take part in the shortcut procedure.

The procedure for the support of ATM shortcuts is shown in Fig. 1. It is assumed that a set of RSVP-capable routers are connected by an underlying ATM network. Classical IP over ATM is run; therefore, the packets will follow a hop-by-hop path. In this scenario, the goal is to determine the longest possible ATM shortcut, which directly connects the ingress and egress routers. The proposed procedure can be described as follows:

- The ingress router, while sending an RSVP PATH message toward the next hop internal to the ATM network, insert its own IP address using a new RSVP class called ATM_FHOP_IP_ADDRESS (FHOP stands for first hop).
- This information is stored in the PATH state information in each subsequent IP over ATM router in the core network. These routers forward the information unmodified. Therefore, the egress router also receives and stores the IP address of the ingress router.
- The egress router, when forwarding the RSVP PATH message on an interface outside the ATM core network, will not include the ATM_FHOP_IP_ADDRESS class.
- When receiving an RSVP RESV message, a router checks if ATM_FHOP_IP_ADDRESS information is stored in the relevant PATH state. In this case, the RSVP RESV message is forwarded using the stored IP address as the IP destination. According to the RSVP message processing rules, such a message will not be interpreted by the intermediate routers, which will simply forward it to the ingress router. In addition to the other information (e.g., Rspec), a newly defined class is added in this enriched RSVP RESV message in order to carry the egress router ATM address for the shortcut. This new class is called ATM_LHOP_ATM_ADDRESS (LHOP stands for last hop). In the example in Fig. 1 the egress router does find the ingress router IP address and sends the special RESV directly to it.

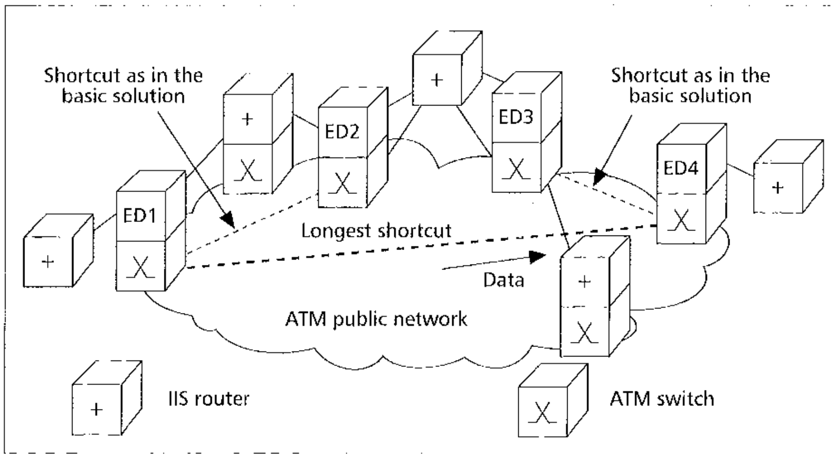


**Figure 2.** *A network topology which can lead to suboptimal shortcuts.*

- The ingress router will receive an RSVP RESV containing the ATM address of the egress router in the new class and, as usual, the traffic specification for the reservation. Therefore, all the information (traffic specs and ATM address) needed to set up a QoS shortcut VC is available, and the ingress router can send the ATM SETUP.

The semantics related to the use of the two newly defined classes can be further explained. When an RSVP router (e.g., the ingress router in Fig. 1) inserts the ATM_FHOP_IP_ADDRESS class in an RSVP PATH message, it means that the router is willing to set up a dedicated VC for the flow under consideration. Therefore, the router could choose to add this information only for special flows requiring a particular QoS (e.g., high bandwidth), which are worth being mapped in a specific VC. The traditional hop-by-hop procedure can be supported at the same time by the same router for all the other flows The addition of the ingress router IP address in an RSVP PATH is only offered to the next routers, which can choose to use or ignore this information when sending the RSVP RESV messages. If an egress router agrees to establish an ATM shortcut VC, it will accept the offer and insert its ATM address in the ATM_LHOP_ATM_ADDRESS class. The information needed by the egress router to correlate the incoming ATM connection setup with the RSVP session will be inserted in the *broadband higher-layer information* (BHLI) information element by the ingress router.

After the procedure is completed and the ATM QoS channel set up, there are two options related to the transport of the next *refresh* RSVP PATH messages. The ingress router will continue adding the ATM_FHOP_IP_ADDRESS class in the refresh PATHs. A first option is to continue using traditional IP routing as the first PATH. The main advantage of this solution is simplicity, because nothing has to be added to the traditional RSVP procedures. Another advantage is that changes in the IP routing can be handled without any risk of loops. The disadvantage is that the intermediate core routers will uselessly process PATH messages and store PATH states for the lifetime of the RSVP flow. If the requirements in terms of message processing and PATH state storage represent a system bottleneck, the second option is to establish ATM shortcuts for the transport of RSVP messages. In this case, either best-effort (UBR) VCs could be used, or QoS VCs with minimal requirements, since the bandwidth needed to transport RSVP PATH messages is very low. If several RSVP flows between two routers are active, these flows can share the same *RSVP control* ATM VC.

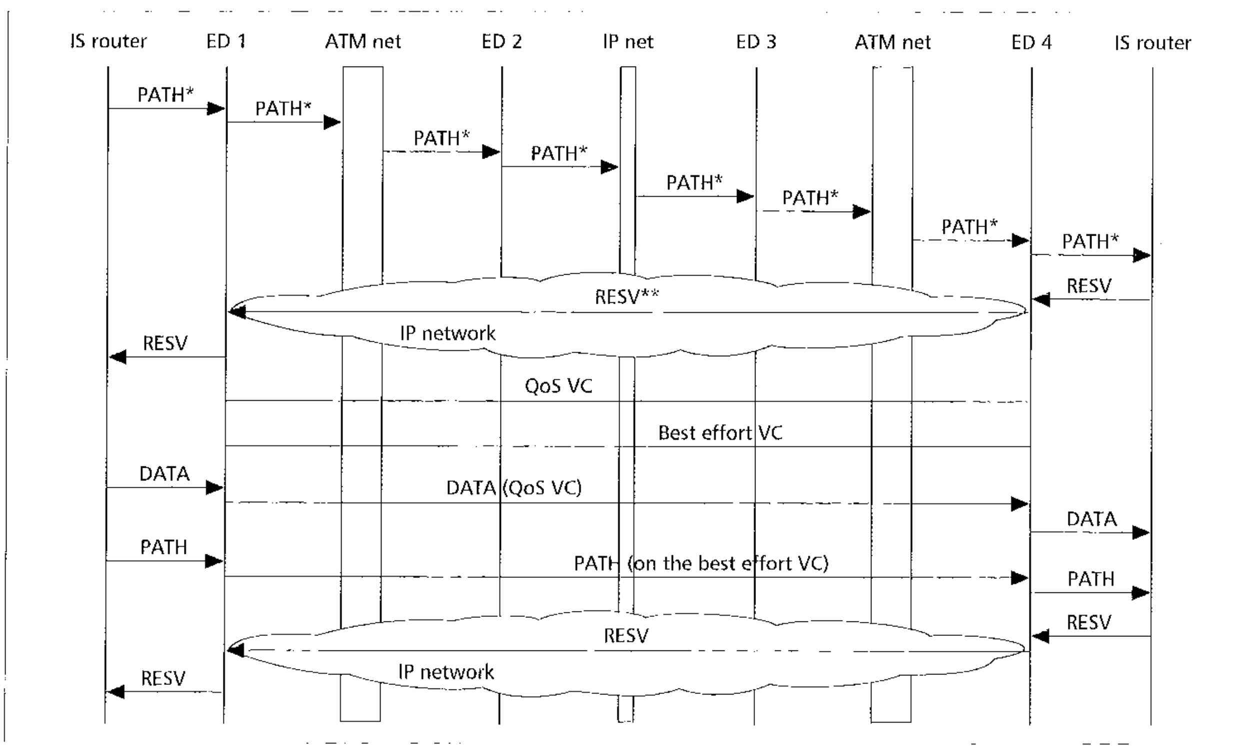


**Figure 3.** *Event trace of the proposed solution.*

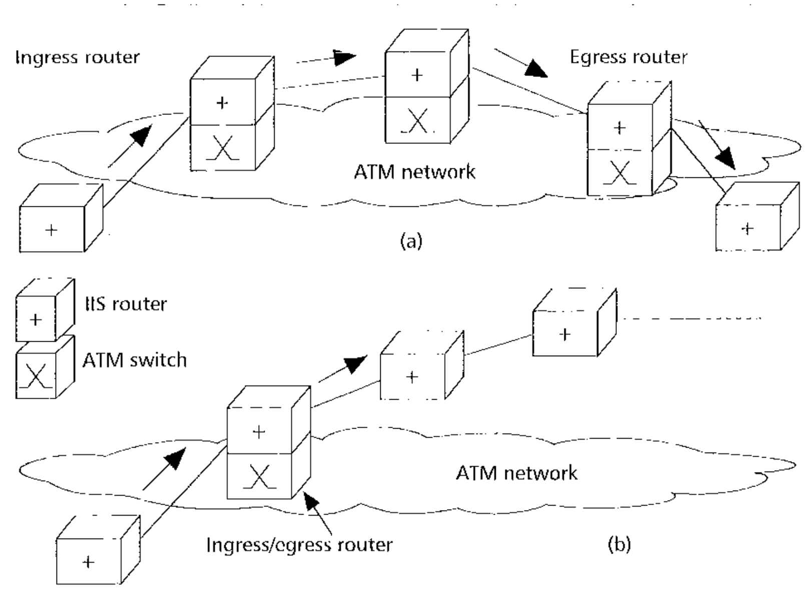

**Figure 4.** *The roles of the routers.*

It is worth noting that the ATM VCs for the transport of RSVP messages cannot be torn down because of low traffic volume, as in the classic implementations of shortcuts, but you must keep them alive as long as you need to route RSVP messages through them.

In the described scenario, a single core ATM network is represented (Fig. 1), but this solution is also able to support more complex network scenarios, where multiple independent ATM networks can be crossed. In such a situation it is possible to get an end-to-end path constituted by several shortcuts (each crossing a single ATM network), interconnected by means of IP sections. In fact, the egress router of the first ATM network takes care of removing the ATM_FHOP_IP_ADDRESS when forwarding the RSVP PATH messages; therefore, the proposed procedure can be applied again independently in the second network.

Under particular conditions (i.e., network topologies) the sequence of shortcuts and IP sections obtained by the described procedure is suboptimal. This happens when the IP route followed by RSVP messages goes out of an ATM network and then enters it again, maybe after having crossed a distinct ATM network. Figure 2 provides an example, showing the optimal ATM shortcut and the suboptimal sequence of shortcuts (dotted lines) and IP sections (thin lines).

It is possible to enhance the solution in order to find the optimal shortcut in the most general case. The needed enhancement to the procedure is that more than one IP address must be allowed in the ATM_FHOP_IP_ADDRESS class. The first ingress router in each ATM network will add its IP address in the ATM_FHOP_IP_ADDRESS class. In this case, the ATM_FHOP_IP_ADDRESS class will also be sent toward IP routers without ATM functionality (as described previously, they just forward the unknown classes).

An event trace of the application of the enhanced solution to the network depicted in Fig. 2 is shown in Fig. 3, where:
- The PATH* message is an RSVP PATH message enriched with an ATM_FHOP_IP_ADDR class to convey the ingress ED IP address
- The RESV** message is an RSVP RESV message enriched with an ATM_LHOP_ATM_ADDR class to convey the egress ED ATM address

A PATH message without the ATM_FHOP_IP_ADDR class reveals that the QoS VC will be torn down; then the ingress router should reroute the flow on a hop-by-hop path.

The price to be paid is the need to execute more checks; especially the egress router must be able to distinguish the IP address (if present) of the routers that belong to its own ATM network in order to store the farthest one. Each egress router can select the destination to which to send the RSVP RESV that allows it to set up the longest shortcut.

## New RSVP Class Definition and Processing Rules

### Defining New RSVP Classes

The definitions of the RSVP classes are reported hereafter. Two new classes are proposed:
- ATM_FHOP_IP_ADDRESS
- ATM_LHOP_ATM_ADDRESS

The maximum size of ATM_FHOP_IP_ADDRESS in the IPv4 case has been fixed to 20 bytes, which allows it to carry up to five different IP addresses. This is needed to enhance the solution, according to the concluding remarks of the previous section, while typically only one IP address is carried and the object content length is 4 bytes. The ATM_LHOP_ATM_ADDRESS has a length of 20 bytes as needed to carry the AESA ATM address format.

The choice of class number has an important impact on backward compatibility. Choosing to have the two highest order bits set to 1 implies that if an RSVP router receives the message and does not recognize the class number, it will forward the object unmodified, without generating any error message. The class number has been temporarily fixed to the following values, but could be modified in the future, according to Internet Assigned Numbers Authority (IANA) considerations [10].

ATM_FHOP_IP_ADDR Class = 208 (binary expression: 11010000)
- IPv4 ATM_FHOP_IP_ADDR Object: Class_Num= 208, Class_Type = 1

ATM_LHOP_ATM_ADDR Class = 209 (binary expression: 11010001)
- Ipv4 AESA ATM_LHOP_ATM_ADDR Object: Class_Num= 209, Class_Type = 1

With respect to performance, the proposed enrichments seem to have very limited impact on the size of the RSVP message, and on the processing and storage requirements in the RSVP routers. On the other hand, the proposed approach seems to provide a straightforward mechanism to set up the longest ATM shortcut without the need for additional address resolution mechanisms like NHRP.

### Processing Rules in the Routers

In this section the processing rules in the routers will be described. In the following, *IIS router* stands for Internet integrated services router, an RSVP-capable router. An IIS over ATM (IISoA) router is an IIS router that further relies on ATM.

The processing rules followed by a router also depend on its role; in Fig. 4a an ingress and an egress router are depicted, while an ingress/egress router is represented in Fig. 4b.

The algorithm performed in the routers when they get the first PATH message of each session will be described by means of a flow diagram as shown in Fig. 5, where the dashed boxes

| [illegible] | [illegible] | [illegible] | [illegible] | [illegible] | [illegible] | [illegible] |
|---|---|---|---|---|---|---|
| 1 | Yes | D | Yes/yes | R | Yes | S |
| 2 | Yes | S | Yes/yes | R | Yes | D |
| 3 | No | D | No/no | R | No | S |
| 4 | No | S | No/no | R | No | D |
| 5 | No | D | No/yes | R | Yes | S |
| 6 | No | S | No/yes | R | Yes | D |
| 7 | Yes | D | Yes/no | R | No | S |
| 8 | Yes | S | Yes/no | R | No | D |

**Table 1.** *A report of test configurations.*

border steps to be taken exclusively in the enhanced solution. Clarifying comments to the processing rules are given below.

An IISoA router must first "realize" whether it is an ingress router for the specified flow. If the PATH message has already entered the ATM network (i.e., passed through at least one ATM-capable router), it carries the ATM_FHOP_IP_ADDR class containing the ATM network ingress router IP address for the specified flow.

In the enhanced solution a router can be considered an ingress router for that flow only if in the PATH message it processes there are no IP addresses of IISoA routers belonging to the same ATM subnet.

In the case of branch a (egress router), only in the enhanced solution can the ATM_FHOP_IP_ADDR class contain several IP addresses belonging to different ATM subnets. In an intermediate router branch b will be followed. Branches c and d are related to the ingress router and ingress/egress routers, respectively.

The IIS routers should not receive messages containing new classes, but in the enhanced solution it is unavoidable. An egress router will forward the PATH message without the ATM_FHOP_IP_ADDR class since such information is unnecessary to subsequent routers in the path.

Further forwarding the PATH message without the ATM_FHOP_IP_ADDR class will allow the same process as above to be repeated in each subsequent ATM subnet.

The algorithm performed in the routers when they get the first RESV message of each session is described by means of a flow diagram in Fig. 6.

According to classical RSVP (branch a), a message not addressed to the considered node will simply be forwarded toward the destination station, with no processing.

An IISoA router that receives an RSVP RESV message can only be an ingress (branch d) or egress router (branch b); in fact, according to the conceived algorithm, the ATM network intermediate routers will never get RESV messages addressed to them.

Branch c represents an ingress/egress router which is the only crossed node belonging to a specific ATM network (the PATH message did not enter the same ATM network again after having crossed that IISoA router) and requires no actions to be performed.

Refresh RESV messages will follow the same path of the first RESV message for that session, obviously without creating any further RESV state.

## Trial Implementation

The enrichments to RSVP have been implemented in a testbed running on a LAN. This section gives an overview of the architecture of the testbed, describing the equipment and software tools used.

The testbed (Fig. 7) consists of three Pentium PCs running Linux (kernel 2.1.125) couple-interconnected via 10 Mb/s Ethernet links. Obviously the two external routers belong to different IP subnets.

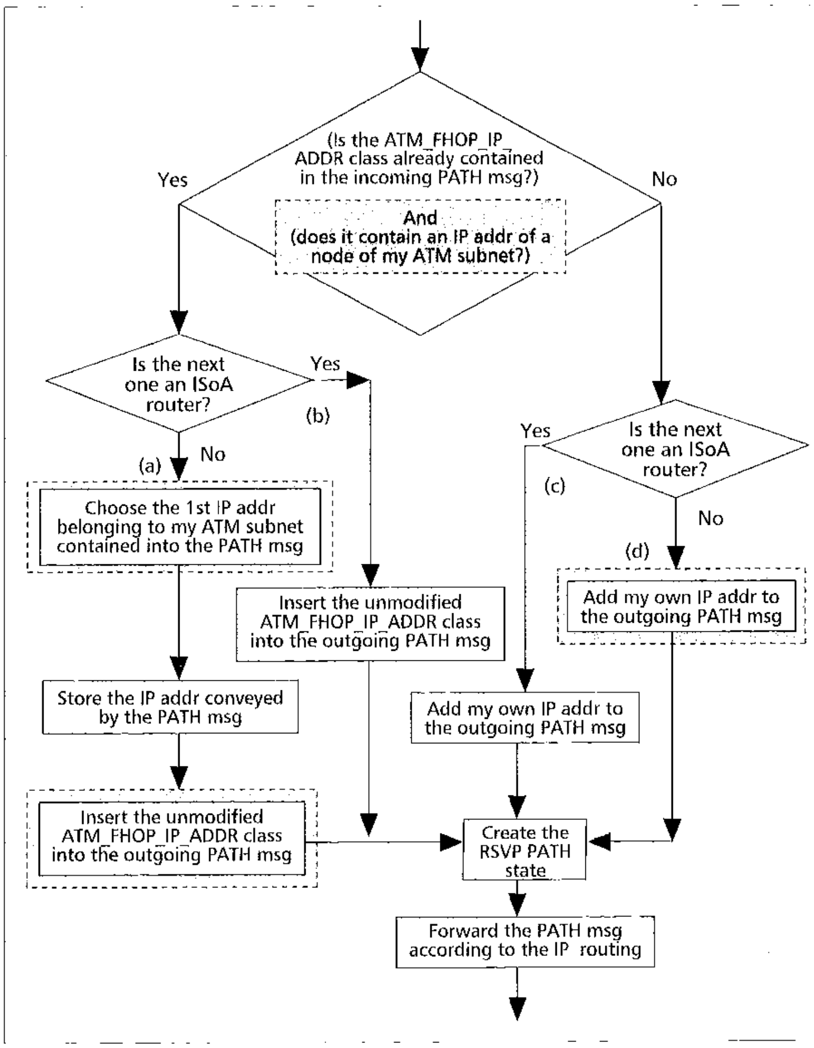


**Figure 5.** *The PATH message processing rules in IISoA routers.*

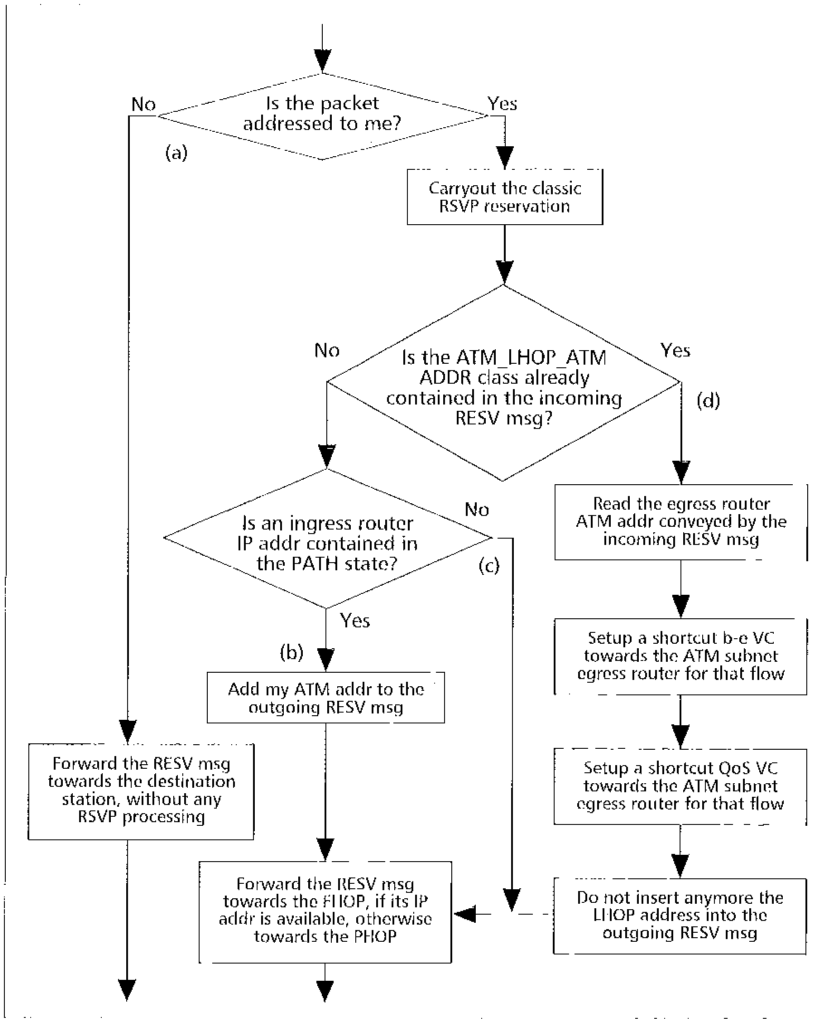


**Figure 6.** *The RESV message processing rules in the IISoA routers.*

The demonstrator has been realized modifying the source code of the *Rel.4.2a4* RSVP daemon [11]. The RSVP daemon is complemented in Linux by the iproute2 package, which tends to traffic control.

During the trial a simulation of both ATM-capable interfaces and IIS-only interfaces was needed: in each node a configuration file contains such information for each interface and eventually its own ATM address.

The trials are listed in Table 1, where S, R, and D stand for source, router, and destination, respectively.

Different topology configurations have been tested, changing the sequence of devices (ingress, intermediate, and egress IISoA routers; legacy IIS routers) The experimentation allowed confirmation of the feasibility of our proposal, functional verification of the processing rules in the routers, and estimation of the software implementation complexity. In fact, the software realization turned out to be very simple.

## Open Issues for the Multicast Case

Let us consider the extension of this approach to the multicast case. The goal is to use a similar mechanism to let the ingress ED obtain the ATM addresses of the set of egress EDs for a given IP multicast address. Then a shortcut ATM multicast VC can be established to support the QoS flow. This approach is basically compatible with the unicast case thanks to the fact that the ATM shortcut VC is established starting from the ingress ED, which can act as the root of the ATM point-to-multipoint connection. Additional leaves (i.e., egress EDs) can be successively added or removed according to RSVP requests. There are several open issues to be solved to define a working solution for the multicast case. A general problem in interworking between RSVP and ATM multicast is that different receivers can specify different QoS for the same flow, or even no QoS at all when they want to receive the flow on a best-effort basis. The interworking solution must specify the algorithm to map the IP multicast flow in one or more ATM point-to-multipoint VCs. Another problem to be considered is the interaction between the IP multicast routing and the ATM shortcuts that can be established modifying the topology. We are currently investigating on a solution where the IP multicast tree is replaced by ATM multicast subtrees whenever possible. This means that the ATM multicast can reach the end terminals if they are ATM-capable, or be confined to ATM-capable routers in the backbone, just as in the unicast solution.

## Conclusions

This article presents an extension of the RSVP protocol to support IP/ATM address resolution and the establishment of ATM QoS virtual channels. The goal is to allow the IIS architecture to fully exploit the capability of an underlying ATM network. Two new classes have been added to the RSVP PATH and RESV messages, allowing IIS over ATM routers to signal their capability (and will) to set up an ATM QoS VC and to exchange ATM addresses.

Compatibility with traditional RSVP hosts and routers has been maintained. No modifications are needed in the host or in routers that do not take part in the shortcut procedure. With respect to performance, the proposed enrichments seem to have very limited impact on the size of the RSVP message, and on processing and storage requirements in the RSVP routers. On the other hand, the proposed approach seems to provide a straightforward mechanism to set up the longest ATM shortcut without the need for additional address resolu-

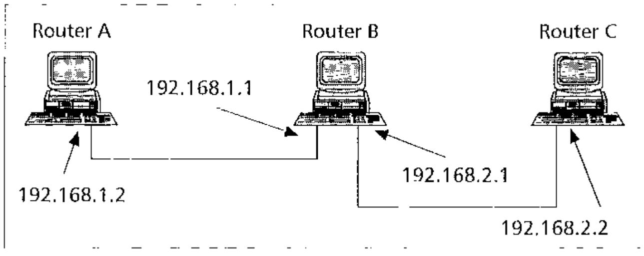


**Figure 7.** *The testbed.*

tion mechanisms like NHRP. The IP unicast case has been considered in detail. Although detailed extension to the multicast case is for further study, there are no conceptual obstacles.

## Reference


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

[2] R. Braden et al., "Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification," RFC 2205, Sept. 1997.

[3] ATM Forum Technical Committee, "LAN Emulation over ATM Version 2 - LUNI Specification," Dec. 1996.

[4] M. Laubach and J. Halpern, "Classical IP and ARP over ATM," RFC 2325, Apr. 1998.

[5] "Multi-Protocol Over ATM Specification V1.0," ATM Forum, July 1997.

[6] E. Crawley et al., "A Framework for Integrated Services and RSVP over ATM," RFC 2382, Aug. 1998.

[7] M. Garrett and M. Borden, "Interoperation of Controlled-Load Service and Guaranteed Service with ATM," RFC 2381, Aug. 1998.

[8] J. Luciani et al., "NBMA Next Hop Resolution Protocol (NHRP)," RFC 2332, Apr. 1998.

[9] D. H. Cancever, "NHRP Protocol Applicability Statement," RFC 2333, Apr. 1998.

[10] R. Braden and L. Zhang, "IANA Considerations for RSVP Version 1," Internet Draft, Apr. 1999.

[11] Rel. 4.2a4 RSVP daemon source code, http://www.isi.edu/div7/rsvp/release.html


## Biographies


ROBERTO COCCA (cocca@coritel.it) received his degree in computer science engineering from the University of Rome "La Sapienza" in 1998, discussing a thesis about the interaction of RSVP with ATM for supporting shortcut QoS VCs, developed in CoRiTeL. Since December 1998 he has been working in CoRiTeL, where he is taking part in the research activity of the ACTS ELISA (European Experiment on the Linkage between Internet Integrated Services and ATM) project group. His main research interests are IP and ATM interworking to support QoS and IP differentiated services architecture.

MARCO LISTANTI [M] received his Dr. Eng. degree in electronics engineering from the University of Rome "La Sapienza" in 1980. He joined the Fondazione Ugo Bordoni in 1981, where he led the TLC network architecture group until 1991. In November 1991 he joined the INFOCOM Dept. of "La Sapienza," where he is associate professor in switching systems. Since 1994 he has also collaborated with the Electronic Department of the University of Rome "Tor Vergata," where he holds courses in telecommunication networks. He has participated in several international research project sponsored by EEC and ESA, and is author of several papers published in the most important technical journals and conferences in the area of telecommunication networks. His current research interests are in traffic control in IP networks and the evolution of techniques for optical networking. He has represented the Italian PTT administration in international standardization organizations (ITU, ETSI).

STEFANO SALSANO [M] (salsano@coritel.it) received his degree with honors in electronic engineering from University of Rome "Tor Vergata" in 1994, and his Ph.D. from University of Rome "La Sapienza" in 1998. At the end of 1997 he joined CoRiTeL, a research institute on telecommunications, where he is now coordinating the research projects in IP related areas. He participated in the INSIGNIA ACTS project (on the integration of B-ISDN with intelligent network) in the areas of architecture and protocol specification. He is currently participating in the ELISA ACTS project (on the support of QoS in IP/ATM networks). His research interests include architectures for broadband networks, integration of IP and ATM, and QoS support in IP networks. He is co-author of several papers on these topics.