A QoS Guaranteeing Framework for the Integration of IP and ATM in DIANA

Martin Lorang
University of Stuttgart, IND, Pfaffenwaldring 47, 70569 Stuttgart, Germany
E-mail lorang@ind.uni-stuttgart.de, Phone: +49 711 685 7991

Abstract

Recently, the Internet community has been working on extensions to the current Internet protocol suite in order to enable service guarantees or differentiation. However, standardised solutions to integrate IP and ATM still are restricted to best-effort service. Therefore, this paper deals with a framework for QoS based interworking between IP and ATM. Network element state with QoS attributes is thereby controlled by RSVP and ATM and an appropriate translation from and to each other. Signalling and traffic control issues related with that translation are of particular concern for this paper. Moreover, the building blocks of an interworking unit prototype architecture as planned by the ACTS DIANA project are presented.

1 Introduction

For many years, public network operators regarded ATM technology as *the* solution for a service integrating broadband network. Conceived as a logical extension of narrowband ISDN, its standardisation was influenced by the connection-oriented paradigm, signalling protocols and addressing scheme known from ISDN. While ATM research, development and standardisation was concerned with guaranteed QoS for typical broadband ISDN applications such as video-on-demand and multimedia conferencing as well as with an efficient rate control for data applications (Available Bitrate), the World Wide Web helped to establish IP networks as *the* carrier for data networking.

The current Internet architecture offers a flexible, but simple connectionless best effort service and for sure this is inadequate for applications sensitive to the QoS provided by the network. For this reason, the IETF has been working on extensions to the current Internet protocol suite in order to enable service guarantees (Resource Reservation Protocol RSVP together with Integrated Services, IIS) [1, 2] or at least differentiation (Differentiated Services, DS) [3, 4, 5].

In Integrated Services, RSVP allows applications to request either Guaranteed or Controlled Load Service for individual flows in an IP network. However, the wide scale deployment of RSVP must be approached with care because the processing of (periodically refreshed) reservation and control messages, the identification of each packet based on the IP header and the handling of per-flow reservation state becomes challenging in backbone routers passed by a huge number of individual flows.

Conversely, the Differentiated Services architecture achieves scalability by classifying and marking packets by means of the so-called DS field in the IP header at the ingress to a DS capable IP network. The goal is to receive a particular per-hop forwarding behaviour on DS routers along their path. A DS admission control service (DACS) [6] may ensure that the resources needed for the aggregate stream are available inside the DS network.

The standardised solutions to integrate IP and ATM that are available so far, namely

Classical IP over ATM (CLIP) [7], Next Hop Resolution Protocol (NHRP) [8], Multicast Address Resolution Server (MARS) [9], LAN Emulation (LANE) [10] and Multiprotocol over ATM (MPOA) [11] focus on IP to ATM address resolution to set-up switched best-effort ATM VCs. Applications with tight QoS constraints cannot be supported.

IP Switching and Tag Switching [12] combine network layer routing with switching. They may or may not use ATM as a link layer technology providing a high switching capacity but, in any case, they do not involve ATM signalling. Both solutions aim at replacing the relatively expensive look-up for IP prefixes by binding a fixed label or tag to a flow on a per-hop basis. Whereas IP switches create a label to flow binding based on a data flow analysis, Tag switches assign tags driven by control messages exchanged by routing protocols or RSVP. Both technologies are compatible with traditional IP routers, but similar to those, they may have to be enhanced by Integrated or Differentiated Services.

Since it is likely that various solutions will co-exist in the near future, a key issue for achieving convergence between Integrated Services, Differentiated Services and ATM to support QoS end-to-end across domain boundaries is the integration of the IIS model and RSVP signalling on the one hand, DS marking and per-hop behaviour as well as ATM service categories and signalling on the other hand.

In this area, DIANA started in March 1998 as a new project in the European Union 4th Framework Programme ACTS. As its main goal, the DIANA consortium will develop and evaluate resource reservation and traffic control functionality which seamlessly interoperates between ATM and IIS networks in order to provide guaranteed QoS end-to-end.

A so-called Integration Unit (IU) will be placed at the boundary between ATM and IP domains to provide the functionality needed for the translation between RSVP and ATM UNI [13] signalling, for the mapping of QoS specifications given by the flow descriptors objects with Integrated Services and ATM traffic descriptor information elements respectively, and for the allocation of ATM virtual connections for IP flows.

This paper is organised as follows: Section 2 starts with a review of previous work on a baseline RSVP and ATM networking framework and some issues for a more advanced implementation. In section 3, it is shown how signalling translation between RSVP and ATM can be solved in a RSVP over ATM and RSVP peering with ATM scenario. Section 4 discusses the QoS and traffic control parameter mapping and is mainly based on previous work. Finally, the building blocks of DIANA's Integration Unit prototype are presented.

2 RSVP and ATM Networking

A number of IETF drafts deal with a framework for Integrated Services and RSVP over ATM. This section points out in which way the outcome of that work could be applied to a baseline implementation in DIANA and discusses advanced issues related with address resolution, multicast support and flow aggregation strategies.

2.1 Baseline RSVP over ATM Implementation

A baseline implementation for RSVP over ATM as specified in [14] only requires an implementation to establish RSVP-initiated VCs to RSVP capable end points. An ATM network supporting switched VCs (SVCs) lies somewhere on the path between those RSVP sender(s) and receiver(s). This scenario, depicted in Fig. 1, is referred to as RSVP over ATM scenario throughout this paper.

The SVCs are set up, added to (in the case of multipoint trees), torn down, and controlled by the edge devices at the ingress to and the egress from the network, which act as both IP

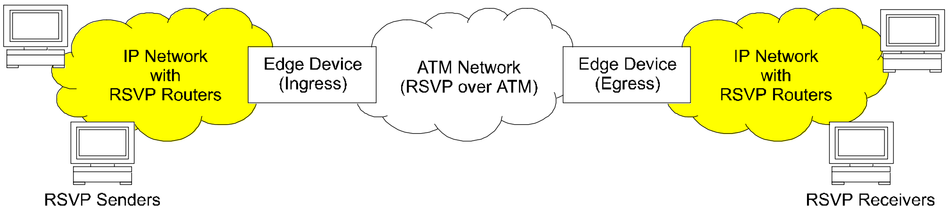

Fig. 1: RSVP over ATM networking scenario

routers and ATM access nodes, capable of initiating and managing VCs at the ATM UNI [15].

It is important to note that, in the RSVP over ATM scenario, the set-up, modification and release of ATM connections is always triggered by RSVP control messages.

A baseline implementation of an edge device [14] sends RSVP messages over the best-effort paths set up by one of the conventional IP over ATM protocols, e.g. CLIP, and uses independent VCs for each RSVP reservation thus doing without flow aggregation. As an alternative, a separate RSVP control VC with a reserved capacity between RSVP demons could be established in order to ensure a normally loss-free and hence more robust operation of RSVP (Fig. 2).

An open issue is how best-effort behaviour for non-conforming packets can be retained inside the ATM network. In [15] cell tagging is proposed, however, since priorities and per-VC queues have been introduced to ATM, tagged traffic in a stream with guaranteed QoS might still affect lower priority traffic in an unfair manner. A solution which is both fair and conformant to the requirements of the Integrated Services is not in sight.

2.2 Advanced Issues with Address Resolution, Multicast Support, Flow Aggregation

CLIP only resolves IP addresses that belong to the same Logical IP Subnet (LIS) based on ATM technology. For unicast address resolution, NHRP could complement (or replace) CLIP without major problems: Destinations outside the LIS are identified during normal routing operation, and thus, (at least) those resolution requests can be forwarded to NHRP.

Conversely, multicast address resolution allowing ATM shortcuts which span multiple subnetworks is problematic. MARS co-ordinates and distributes IP multicast group address to ATM endpoint address resolution information but its scope of responsibility is a so-called cluster which is likely to coincide with a LIS. Whereas IP group membership and routing protocols benefit from aggregation in each node of the multicast tree [16], a point-to-multipoint ATM shortcut VC set-up so far requires each individual end-point to be known to the root (as long as the network Leaf Initiated Join capability [13], see below, is not widely implemented). For this reason, bypassing of IP multicast routers with point-to-multipoint shortcut VCs is not recommended.

Furthermore, RSVP provides many features, such as receiver oriented reservations, heterogeneity within a multicast session, dynamic change of reservations, and multiple reservations styles that cannot easily be supported with ATM. In ATM, the first branch of a point-to-multipoint VC determines the QoS (statically) for the whole tree. Hence, ATM does not allow for heterogeneity in a single point-to-multipoint VC, and separate point-to-

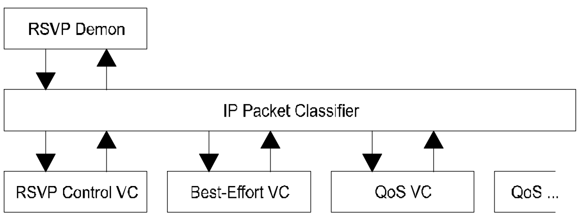

Fig. 2: RSVP over ATM networking scenario

multipoint VCs would have to be set up if ATM is supposed to support heterogeneous QoS.

The Leaf Initiated Join (LIJ) capability offered by UNI4.0 [13] enables ATM receivers to join a point-to-multipoint VC but without the option of specifying an individual QoS. Thus LIJ does not resolve the mismatch between RSVP and ATM with respect to receiver heterogeneity.

For this reason, several authors [14, 17, 18, 19] abandon the goal of full heterogeneity and introduce a limited heterogeneity or modified homogeneous model [14, 15, 19]. With limited heterogeneity, a best-effort and a single alternate QoS are offered, whereas only one VC with the maximum requested QoS is established in the modified homogeneous model.

Potentially, some of the scaling issues of RSVP can be addressed by aggregating several RSVP flows into a single VC if the path through the ATM network is the same for all the flows being aggregated. Similar to multicast, aggregation complicates VC management, and even worse, requires scheduling mechanisms at the originating point of the VC. On the other hand, RSVP would become invisible inside the ATM network and normal connection classification based on the VCI and VPI would apply. If the VC capacity is dimensioned appropriately, end-to-end QoS is feasible. However, since only flows with an identical path through the ATM network can be aggregated, the positive impact on the amount of reservation state and processing may be limited.

In contrast, RSVP or Differentiated Services [6, 20, 21, 22] based aggregation aims at reducing reservation state in each IP hop independently. In spite of some subtle differences, their common paradigm is that routers inside the aggregation region (the backbone) classify packets at the ingress with a limited number of different classes and do no maintain per-flow reservation state in the interior. The admission control will be done at the ingress too, but the decision will be based on congestion information within the aggregation region.

3 Signalling Message Mapping between RSVP and ATM

The integration of Integrated Services and ATM requires the interaction of the respective signalling procedures and the related traffic control as well as the mapping of service classes and their parameters from RSVP to ATM and vice versa. In this section, a networking scenario is developed based on IP as the common user plane overlay layer. Nonetheless, still a distinction can be made between a scenario where all applications use the same type of signalling - RSVP - and a scenario where some applications use - maybe by extending the standardised IP over ATM protocol implementations of CLIP, LANE or MPOA similar to Arequipa [23] - UNI signalling without involving RSVP.

3.1 Brief Overview of RSVP and ATM Design Paradigms

RSVP allows applications to request an Integrated Services based state with QoS attributes associated with a flow in all IP routers along a path from the source to the destination. As opposed to ATM, RSVP is receiver oriented and, by using soft state, can thus support dynamic group membership and receivers with heterogeneous QoS demands within a multicast session. In RSVP, a receiver signals its reservation request as part of a *Resv* message based on a traffic specification contained in *Path* messages previously sent by the source along the IP (multicast) route. Since the routers establish path state while processing *Path* messages, *Resv* messages can return along the reverse routes. Reservation set-up is uni-directional. Table 1 gives a survey of this and other differences between RSVP and ATM.

Table 1: Comparison of RSVP and ATM Signalling

RSVP	UNI4.0
Receiver requests reservation in response to a *Path* message, uni-directional	Sender sets up a bi-directional VC
Default best-effort service	No connectivity if set-up fails
Heterogeneous QoS within a multicast session	Point-to-multipoint VCs with a homogeneous QoS
Dynamic QoS: RESV can alter the reservation at any time	Static QoS, negotiated at set-up (Q.2963.x now specifies sender controlled modification procedures)
Multiple reservation filter styles to select different senders in a multipoint-to-multipoint scenario	Point-to-multipoint
Soft state: Messages are resent periodically	Hard state: Connections have to be explicitly released
Guaranteed, controlled load and best-effort service	CBR, rt/nrt-VBR, ABR, UBR

3.2 RSVP over ATM Scenario

DIANA's Integration Unit can be regarded as an implementation of an edge device as considered in [15] and section 2. With such an edge device a RSVP over ATM scenario can be realised, see Fig. 1 above, in which both connection end-points are controlled by RSVP.

With RSVP over ATM, illustrated in Fig. 3, a RSVP sender starts sending *Path* messages downstream towards the receiver. Provided that a RSVP control VC has been established before, those *Path* messages can pass the ATM network without triggering the exchange of signalling messages and finally arrive at the receiver. In normal operation, the receiver returns a *Resv* message that is to be processed by the IU devices which act as RSVP capable routers. The IU at the downstream ingress to the ATM network originates a call to set up an „unidirectional" ATM VC, i.e. a VC with no reservation in the upstream direction. A failure in the set-up of this ATM VC should trigger the same RSVP error messages as if the local reservation had failed.

The traditional method to account for changes in RSVP reservations, as described in detail by Berger [14], is to attempt to replace an existing VC with a new, appropriately sized VC. During the set-up of the replacement VC, the old VC must be left in place unmodified. If the set-up of the replacement VC fails, then the old QoS VC must continue to be used. If the new reservation is greater than the old reservation, the reservation request must be answered with an error in this case. When the new reservation is less than the old reservation, the request must be treated as if the modification was successful.

Latest ATM signalling standards permit modifications of traffic parameters by the connection initiator. In [24] and [25] modification procedures for the Peak Cell Rate and

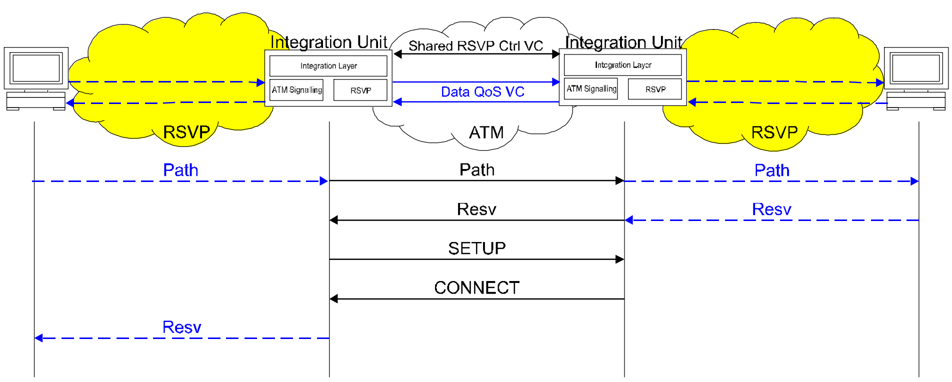

Fig. 3: Signalling message flow for RSVP over ATM (Simplified)

Sustainable Cell Rate have been defined, and this might ease to retain RSVP reservation dynamics in ATM. However, only parameters specified at connection set-up can be altered, so a change from one ATM service category to another one is not possible.
RSVP can identify from either explicit messages or time-outs when a data VC is no longer needed. Therefore, data VCs set up to support RSVP controlled flows should be released at the direction of RSVP. A release at the direction of ATM should only happen when ATM is the end-point of the QoS path, as is the case in RSVP peering with ATM scenario, see Fig. 4.

3.3 RSVP peering with ATM Scenario

In the RSVP peering with ATM scenario, an user application that communicates with an ATM API without involving RSVP interworks with an application that uses RSVP to signal its resource demands, as shown in Fig. 4. The IU is now the terminating entity for ATM signalling and the originating entity for RSVP *Path* messages or vice versa.

When setting up a RSVP reservation for a call originating from ATM, the IU may use the information contained in the incoming SETUP message to send a *Path* message towards the RSVP controlled destination, as illustrated in Fig. 5, provided that the relevant IP addresses have been communicated by non-standard means (see below). At the same time, an ALERT message is returned to avoid a premature expiry of signalling timers (T303, T310) [26].

Only on the receipt of a *Resv* message, it should answer to the SETUP with a CONNECT message. If the original SETUP message contained one or more Alternative ATM Traffic Descriptor or a Minimum Acceptable ATM Traffic Descriptor Information Element [26], a RSVP reservation request different from what was specified in the Path message may be taken into account in ATM parameter negotiation. In any case, an ATM connection is only established if the reservation was successful on the entire RSVP path. It is important to note that the ATM application should refrain from requesting bi-directional ATM VCs and instead rely on the RSVP destination to specify the QoS of the reverse path if needed.

Alternatively, the VC set-up may be finished before *Path* messages are sent towards the IP receiver. As before, there is no danger of a false expiry of signalling time-outs since the UNI signalling demon completes the ATM set-up without interaction with RSVP modules. However, if the receiver does not accept the traffic specification as advertised by the Path message, a *Resv* message arrives at the IU which does not match the original reservation the ATM connection was based upon. Since it is not the initiator of the VC, the IU does not dispose of the standard mechanism [24] to alter the parameters specified in previous ATM Traffic Descriptor Information Elements of the already active connection. But this applies to Fig. 5 too after the set-up has finished.

Even worse, since the ATM connection set-up is finished before the reservations in the IP domain are established, the ATM sender and originator of the QoS path in the upper half of Fig. 6 may already start sending before a QoS path has been prepared throughout the network. The opposite direction (see bottom half of Fig. 6) is less problematic because an ATM destination has no means to downgrade a reservation (By specifying an alternative or minimum traffic descriptor information element in the SETUP message, the call originator could give the receiver a chance of making a selection among different traffic descriptors, however, the IU,

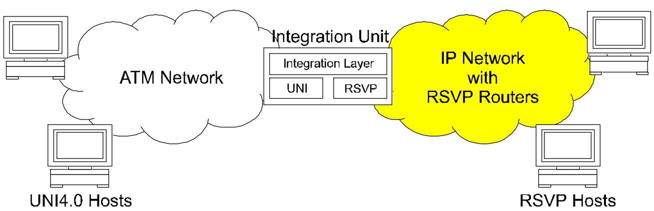

Fig. 4: RSVP peering with ATM scenario

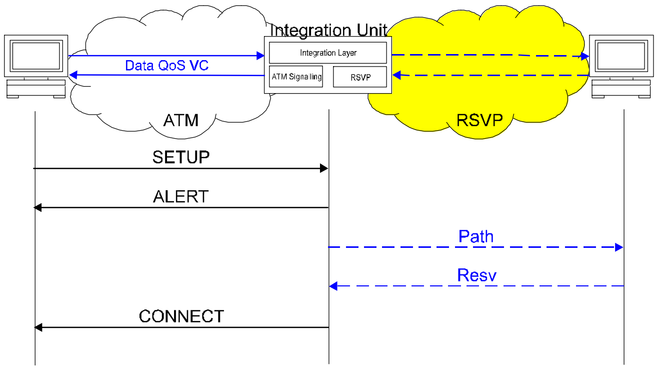

Fig. 5: RSVP peering with ATM: Signalling message flow for an ATM initiated QoS path, waiting for RSVP reservation request (Simplified)

being the call originator, can simplify the procedure by not offering this option).

If both the RSVP over ATM and the RSVP peering with ATM scenario was implemented, the IU would have to run the different signalling procedures shown in Fig. 3, Fig. 5 and Fig. 6. A classification of incoming SETUP messages based on the ATM address and Service Access Point would be required to select the right procedure. Since specifications of IP and ATM [15, 27, 28] as well as some examples in [13] already use the Broadband Lower Layer Information Element (BLLI-IE), the Broadband Higher Layer Information Element (BHLI-IE) could be used for making this distinction and for communicating the IP addressing information (in a non-standardised way). Nevertheless, the IP and ATM address mapping and signalling issues of RSVP peering with ATM suggest to refrain from implementing this scenario.

4 QoS and Traffic Control Parameter Mapping between RSVP and ATM

Prior to initiating the actual traffic and QoS parameter transformation, the IIS Guaranteed Service and Controlled Service respectively have to be mapped to appropriate ATM service categories [29].

Controlled Load Service (CLS) is intended to support applications that are sensitive to congestion but that can exactly characterise their traffic profile. CLS is expected to be equivalent with a best-effort service „under unloaded conditions“. The traffic descriptor consists of token bucket parameters, i.e. a bucket rate r and a bucket depth b, a peak rate p, a minimum policed unit m and a maximum packet size M. This is a variable rate specification,

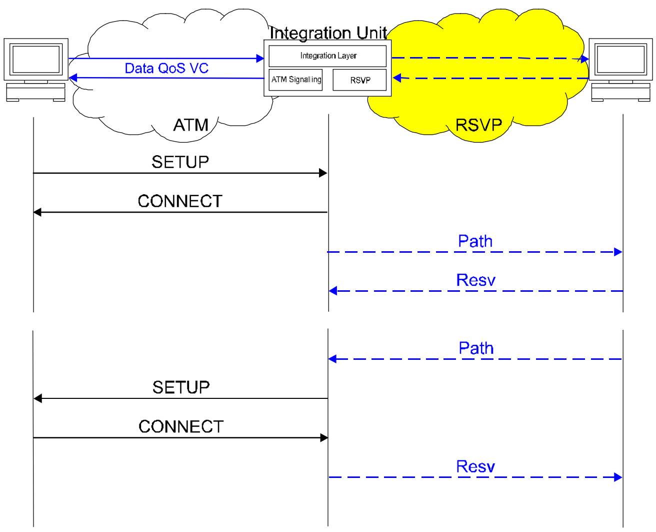

Fig. 6: Signalling message flow for RSVP peering with ATM when ATM set-up completes immediately (Simplified)

and, since CLS does not provide delay guarantees, non-real-time Variable Bit Rate (VBR*)* matches best. With a positive Minimum Cell Rate (MCR), Available Bit Rate (ABR) offers a fixed minimum throughput plus a dynamic component which matches as well. Even Constant Bit Rate (CBR) will work fine if one accepts an over-allocation of resources.

Guaranteed Service (GS) is foreseen for applications with stringent real-time requirements. The end-to-end behaviour of a series of network elements implementing GS is an assured level of bandwidth that, when used by a token bucket constrained flow, yields a delay-bounded service without queueing loss. The traffic descriptor includes the same parameters as for GS. Of course, only ATM service categories with delay guarantees can be used to relay GS traffic, i.e. either CBR or real-time VBR.

A convincing mapping of Guaranteed and Controlled and an in-depth description of related issues is given in [15]. This document also points out that the ingress edge device (or IU) should keep a table of QoS information, for the set of egress edge devices that it may establish VCs with in order to be able to check whether the requested IP service delay bounds may be supported. Similarly, an ATM request for a certain Cell Loss Ratio has to take IP packet size distribution and potential frame discard into account.

5 Building Blocks of an Integration Unit Prototype Architecture

DIANA's Integration Unit has to be fully functional in both RSVP/IIS as well as ATM. On the RSVP side, the IU adopts the role of a RSVP capable router, i.e. processes RSVP messages, reserves resources, and maintains soft state (in the control path), and classifies, polices and schedules packets (in the data path) before finally forwarding them. As an ATM end system, the IU sets up ATM connections and accepts or refuses incoming connections.

Due to the problems with the RSVP peering with ATM scenario, the IU prototype as described here only includes the functionality required for the RSVP over ATM scenario. The key issue is that admission control for Integrated Services is now dependent on the successful set-up or modification of a connection across the ATM network. In RSVP terminology, the Link Layer Dependent Adaptation Layer (LLDAL) is supposed to provide this functionality. As a first step, LLDAL translates between IIS to ATM traffic descriptors. Then the set-up of a new VC, the modification of an existing VC or, if aggregation is taking place, a simple look-up for a VC already active and the respective admission decisions will be performed.

In the IU, the ATM specific LLDAL is called Interworking Control module. The RSVP demon invokes this module whenever RSVP operation and message processing cause a reservation state change. Both the RSVP demon and the Interworking Control module as well as the UNI signalling demon manipulate kernel data structures that represent all attributes assigned to a connection and will communicate connection handles appropriately.

As QoS and traffic parameter mapping from RSVP to ATM as well as flow-to-VC management functions can be highly interdependent, especially when aggregation is taking place and a management function may decide to allocate extra resources in anticipation of further reservations, e.g. when a certain percentage of the available VC bandwidth is consumed, they deserve further studies treating QoS mapping, VC management, and CAC in an integrated manner. The outcome of this research will help to incrementally improve the Interworking Control module.

Although not shown in Fig. 7, CLIP will be used for the IP to ATM address resolution. Both the signalling demon and CLIP's ATMARP demon can either be called implicitly via the kernel as a result of a routing decision or, as an alternative, explicitly from the Interworking Control module. The latter is feasible because RSVP path state includes the IP address of the previous hop RSVP node, e.g. the ingress IU. With this information, the address resolution

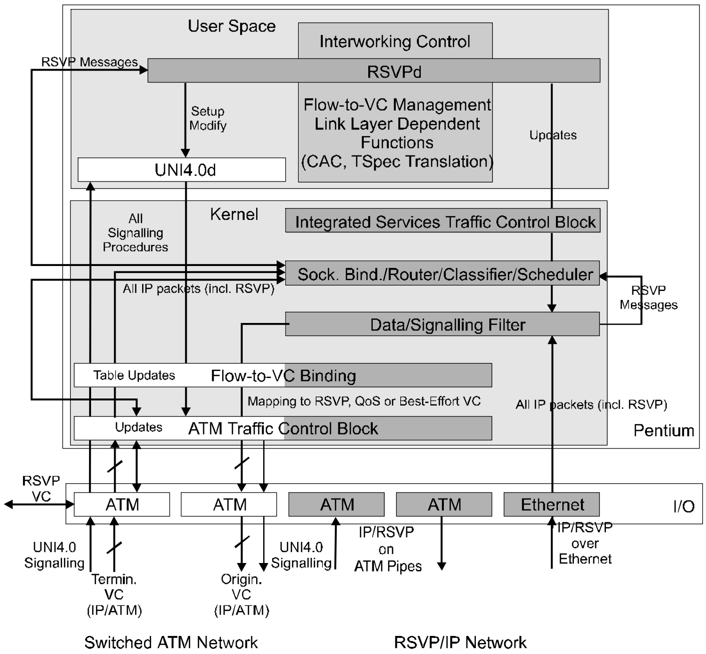

Fig. 7: Block Diagram of the Integration Unit

procedure can be started without involving the kernel routing tables provided that the incoming interface stored in the path state can be identified as an interface to the switched ATM network.

Signalling messages arriving or leaving on the standard signalling VCs are handled by the UNI signalling demon whereas other switched VCs, namely the RSVP control VCs and the RSVP over ATM VCs, terminate in the IP module in the same way as IP paths (on the bottom right of Fig. 7). From there, IP packets carrying the well-known RSVP port numbers are redirected to the RSVP demon while data packets are subject to routing and classification (to identify the flow they belong to) before they are scheduled to be sent. The settings of the classifier and scheduler are controlled by the RSVP demon by a set of Integrated Services specific data structures grouped in the Integrated Services Traffic Control Block.

6 Conclusions

This paper has presented an architecture which integrates IP and ATM on the basis of IETF Integrated Services. In order to accomplish QoS end-to-end, a comprehensive interworking architecture encompassing address resolution, routing, signalling and traffic control is required. Work in DIANA mainly focuses on the specification, implementation and evaluation of signalling translation, and related with that, traffic and QoS parameter mapping in a so-called Integration Unit which is to be placed at the boundaries between IP and ATM networks.

DIANA's networking model is based on IP as a common network layer and the assumption that end systems are connected to different link layers but use RSVP to provide the application with the appropriate control capabilities. However, also scenarios with end points using native ATM signalling have been studied.

Aggregation of flows to reduce processing in core network elements is supposed to be an important issue for the Internet of the future. As has been pointed out, the Integration Unit may aggregate several flows to one VC when using a preventive scheme to trigger re-negotiation procedures. For sure, the integration of those and further aggregation mechanisms into existing networks and their evaluation is a challenging task for the future.

References

[1] R. Braden (Ed.), L. Zhang, S. Berson, S. Herzog and S. Jamin, *Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification.* IETF Request for Comments, RFC 2205, September 1997.

[2] P. White, „RSVP and Integrated Services in the Internet: A Tutorial“, *in IEEE Communications Magazine,* Vol. 35, No. 5, pp. 100-106, May 1997.

[3] D. Black, S. Blake, M. Carlson, E. Davies, Z. Wang and W. Weiss, *An Architecture for Differentiated Services.* Work in Progress, IETF, <draft-ietf-diffserv-arch-00.txt>, May 1998.

[4] Y. Bernet, J. Binder, S. Blake, M. Carlson, E. Davies, B. Ohlman, D. Verma, Z. Wang and W. Weiss, *A Framework for Differentiated Services.* Work in Progress, IETF, <draft-ietf-diffserv-framework-00.txt>, May 1998.

[5] K. Nichols and S. Blake, *Definition of the Differentiated Services Field (DS Byte) in the IPv4 and IPv6 Headers.* Work in Progress, IETF, <draft-ietf-diffserv-header-00.txt>, May 1998.

[6] Y. Bernet, R. Yavatkar, P. Ford, F. Baker, L. Zhang, K. Nichols and M. Speer, *A Framework for Use of RSVP with Diff-serv Networks.* Work in Progress, IETF, <draft-ietf-diffserv-rsvp-00.txt>, June 1998.

[7] M. Laubach, *Classical IP and ARP over ATM.* IETF Request for Comments, RFC 1577, January 1994.

[8] J. Luciani, D. Katz, D. Piscitello, B. Cole and N. Doraswamy: *NBMA Next Hop Resolution Protocol (NHRP).* IETF Request for Comments, RFC 2332, April 1998.

[9] G. Armitage, *Support for Multicast over UNI 3.0/3.1 Based ATM Networks.* IETF Request for Comments, RFC 2022, November 1996.

[10] ATM Forum, *LAN Emulation over ATM - Version 1.0.* ATM Forum, AF-LANE 0021.000, January 1995.

[11] ATM Forum, *Multiprotocol over ATM Version 1.0.* ATM Forum, AF-MPOA-0087.000, July 1997.

[12] Y. Rekhter, B. Davie, G. Swallow, E. Rosen, D. Farinacci and D. Katz, „Tag Switching Architecture Overview“, in *Proceedings of the IEEE,* Vol. 85, No. 12, December 1997.

[13] ATM Forum, *ATM User-Network Interface (UNI) Signalling Specification - Version 4.0.* ATM Forum, AF-SIG 0061.000, July 1996.

[14] L. Berger, *RSVP over ATM Implementation Requirements.* Work in Progress, IETF Integrated Services Working Group, <draft-ietf-issll-atm-imp-req-03.txt>, April 1998.

[15] M. W. Garrett and M. Borden, *Interoperation of Controlled-Load Service and Guaranteed Service with ATM.* Work in Progress, IETF Integrated Services Working Group, <draft-ietf-issll-atm-mapping-05.txt>, March 1998.

[16] G. Armitage: *VENUS - Very Extensive Non-Unicast Service.* IETF Request for Comments, RFC 2191, September 1997.

[17] A. Demirtjis et al., *RSVP and ATM Signalling.* ATM Forum Contribution 96-0258, January 1996.

[18] E. Crawley, *A Framework for Integrated Services and RSVP over ATM.* Work in Progress, IETF, <draft-ietf-issll-atm-framework-01.txt>, November 1997.

[19] L. Berger *RSVP over ATM Implementation Guidelines.* Work in Progress, IETF Integrated Services Working Group, <draft-ietf-issll-atm-imp-guide-04.txt>, April 1998.

[20] S. Berson and S. Vincent, *Aggregation of Internet Integrated Services State.* Work in Progress, IETF Integrated Services Working Group, <draft-berson-classy-approach-01.ps >, November 1997.

[21] R. Guerin, S. Blake and S. Herzog, *Aggregating RSVP-based QoS Requests.* Work in Progress, IETF Integrated Services Working Group, <draft-guerin-aggreg-rsvp-00.txt>, November 1997.

[22] T. Li and Y. Rekhter, *Provider Architecture for Differentiated Services and Traffic Enginieering (PASTE).* Work in Progress, IETF Network Working Group, <draft-li-paste-00.txt>, January 1998.

[23] W. Almesberger, J-Y. Le Boudec and Ph. Oechslin, *Application REQuested IP over ATM (AREQUIPA).* IETF Request for Comments, RFC 2170, July 1997.

[24] ITU-T, *Peak cell rate modification by the connection owner.* ITU-T Recommendation Q.2963.1, July 1996.

[25] ITU-T, *Digital Subscriber Signalling System No. 2 - Connection modification: Modification procedures for sustainable cell rate parameters.* ITU-T Recommendation Q.2963.2, September 1997.

[26] R.O. Onvural and R. Cherukuri, *Signaling in ATM Networks.* Norwood, MA: Artech House, 1997.

[27] M. Perez et al., *ATM Signalling Support for IP over ATM.* IETF Request for Comments, RFC 1755, February 1995.

[28] R. Cole, D. Shur and C. Villamizar, *IP over ATM: A Framework Document.* IETF Request for Comments, RFC 1932, April 1996.

[29] ATM Forum, *Traffic Management Specification - Version 4.0.* ATM Forum, AF-TM-0056.000, April 1996.