AN EDGE DEVICE FOR SUPPORTING INTERNET INTEGRATED SERVICES OVER SWITCHED ATM NETWORKS

H. Hussmann
Dresden University of Technology
Germany
Heinrich.Hussmann@inf.tu-dresden.de

B.F. Koch
Siemens AG
Germany
Bert.Koch@oen.siemens.de

G. K. Mamais
National Technical University of Athens
Greece
gmamais@telecom.ntua.gr

I.S.Venieris
National Technical University of Athens
Greece
venieris@softlab.ece.ntua.gr

ABSTRACT

This paper presents an architecture for combining the Internet approach for the support of multimedia services with switched ATM network technology. As the idea is to avoid changes on the underlying technologies, we introduce an Edge Devise (ED) which provides all functionality required for the efficient support of Internet Integrated Services (IntServ) through switched ATM networks. Core functions of the ED are the interoperability between RSVP (Resource Reservation Protocol) and broadband signaling at the control plane, the synergy between RSVP and DiffServ (Differentiated Services) at the User plane, the support of multicast services and address resolution via VC management, as well as charging and billing. Compared to proposed Internet routers the ED presents increased functionality both at the user and Control plane as it fully exploits the advantages of IP to ATM and vice versa. The paper proposes a set of RSVP / signaling scenarios for efficiently using the control capabilities of both networks towards the support of IP based multimedia applications.

KEYWORDS:

ATM, INTERNET, GUARANTEED SERVICES, DIFFERENTIATED SERVICES, RSVP INTEGRATED SERVICES, SIGNALING.

1. INTRODUCTION

The Internet has brought a world-wide breakthrough in the acceptance of multimedia services. The critical mass of applications exists to make multimedia communication an important part of daily life for many people. Unfortunately, the current technological state of the Internet does not yet adequately support advanced multimedia applications which are based on real-time communications involving video and voice. Necessary features for such advanced multimedia applications like guaranteed Quality of Service (QoS) are not available in the Internet. This deficiency of the Internet results from its heterogeneous nature; Internet is currently formed by a mixture of different network types which are mainly based on a connection-less, 'shared-medium' environment. Even in those parts of the Internet where the underlying link layer technology would be able to give good support for multimedia applications, the Internet Protocol (IP) and the suite of related protocols hides such features from application programs.

There are several initiatives from the Internet community to resolve this problem, in principle the so-called *Integrated Services* [1, 2] and the so-called *Differentiated Services* [3, 4] approach. But, unfortunately, the first approach which uses explicit resource reservation is considered as rather difficult to scale up to a world-wide network, and the second does not actually solve the underlying technical problem but tries to split the Internet traffic according to different priorities. So, there is still good reason to investigate alternative ideas.

A typical example of a network infrastructure which provides optimal support for novel interactive multimedia applications is ATM. Application developers are interested in ATM because of its ability to set up on demand a point to point Virtual Channel (VC) with specified Quality of Service (QoS). For network operators, ATM has additional attractive features as individual charging and billing for network usage, and comparatively high security standards. The standards for ATM have been defined in such a way that it is suitable for wide-area networks (where it is called B-ISDN) as well as for Local Area Networks (LANs). However, the trend of the last few years has shown that the deployment of an end-to-end ATM infrastructure is too expensive compared to the competing technology (in particular Fast Ethernet and Gigabit Ethernet in Local Area Networks). As a result, the current main usage of ATM is as a lower-layer technology for wide-area networks which in most cases carries IP-related traffic. Even worse for the ATM technology, there is a severe lack of application programs which are constructed for the usage of pure ATM switched connections. The low market penetration of ATM LANs has the consequence that application developers, although basically interested in the features of ATM, concentrate on the IP protocol stack.

This situation can be summarized as a 'bottleneck' problem. Application developers have some interest in advanced network features. The required features are becoming available in wide-area ATM networks. But protocol issues stemming from Internet traditions prohibit the efficient usage of the network technology.

This paper promotes one particular strategy to remove the bottleneck mentioned above. The idea is here to combine existing approaches from the Internet community and from ATM wide-area network technology in order to relieve the current limitations. More specifically, the idea is to realize this without any changes to the underlying technologies (ATM, Internet protocols). An appropriate *edge device* will be specified and realized which achieves the integration between the two technologies. Let us assume an end system which is connected by a powerful, but not ATM-based, local area network to such an edge device. The edge device can rely on the behavior of the local network it is connected to (where for instance a resource reservation protocol in routers or switches may be in effect). For wide-area connectivity, the end system can rely on the *illusion of an Internet Integrated Services network on a global scale.* The end system uses the relevant protocols from the Integrated Services architecture, e.g. RSVP for reserving bandwidth towards another end system. However, the edge device translates such RSVP requests into appropriate control messages for a wide-area ATM network. Somehow this can be seen as a *Virtual Integrated Services Network*. It is worth noting that the edge device is completely free in its choice for the organization of transport of IP packets over the ATM network, so also the Differentiated Services approach can be used to a significant extent.

This paper describes the specific approach taken by the project ELISA (European Experiment on Linkage of Internet Integrated Services with

ATM) which is supported by the European Union under research grant AC310. The key properties which distinguish ELISA from other competing approaches are its consequent usage of existing technologies (both on Internet and ATM side) as well as its commitment to practical experimentation over real European ATM networks.

This paper is organized as follows. Section 2 gives an introduction into the network and software architecture of the proposed solution. Emphasis is put here on a clear separation between control plane and user plane aspects. Section 3 discusses the basic concepts which enable interworking of Internet-style resource reservation (Integrated Services), ATM signaling and Internet Differentiated services. Section 4 gives some more technical details, in particular ideas for address handling and optimization of the use of connections in the wide-area network. Finally, a brief account of performance issues for the proposed architecture is given in section 5. This paper has been produced at a rather early stage of the ELISA project; therefore some of the reported material is not yet a description of a final prototype but a discussion of a detailed plan.

2. ARCHITECTURE FOR IP OVER SWITCHED ATM FOR GUARANTEED QUALITY SERVICES

A network architecture that aspires to provide guaranteed quality services that are based on the Internet protocol stack, must meet a set of criteria which are often in contrast with others. Therefore, besides the basic requirement which is the provision of guaranteed quality services, the network architecture must satisfy the following demands:

- To provide full backwards compatibility with the existing Internet and Internet applications.
- To be scaleable.
- To be simple, economical and easy to be implemented.
- To induce inconsiderable changes in the users terminals
- To provide effective mechanisms for charging and billing.

A candidate network architecture that meets most of the above requirements is illustrated in Figure 1. The architecture consists of end-user terminals, Edge Devices (EDs) and ATM switches which compose the core network. Customer Premises Equipment (CPE) running Internet Integrated Services (IIS, IP/RSVP multimedia applications) are grouped into local access networks. Local access networks can be based in a variety of network technologies (e.g. Fast Ethernet, switched LANs, N-ISDN, ADSL and others) depending on the needs and on the attributes of the end-user. Each access network is connected to the public ATM network via an ED which acts as a gateway. Users belonging to the same or different access network are able to seamlessly interoperate by using their standard IP applications through the ATM network with the additional provision of guaranteed QoS.

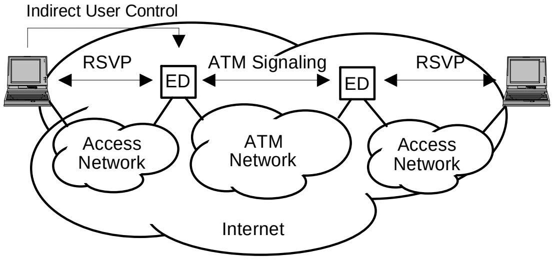

Figure 1: Network architecture

QoS is preserved end-to-end. To do so, both RSVP and ATM signaling mechanisms are employed in a synergetic manner. The RSVP, realized between the terminal and ED and between EDs, transfers user requests for network resources to the ED. The ED translates the request to ATM signaling messages and forwards them to the network. Doing so, the appropriate resources are reserved in the access network (RSVP) and in the ATM network (ATM signaling.) Therefore, the flow admission control realized at the EDs for the allocation of resources to IP requests will closely cooperate with the connection admission control algorithm used in the ATM switches for handling signaling requests. Furthermore, any traffic that is produced by existing, RSVP unaware Internet applications running on CPEs is handled in the ED as best effort; i.e., it is handled as a low priority typical Internet service. The core network of the proposed architecture realizes switched ATM connections and therefore it concentrates all advantages stemming from statistical multiplexing and dynamic resource allocation.

2.1 Control Plane

In the network architecture, illustrated in Figure 1, more than a single network protocol stack co-exist. As far as the control plane of the architecture is concerned, see Figure 2, we can distinguish between the ATM signaling stack (User Network Interface (UNI) 3.1 or 4.0 [10]) and Resource ReSerVation Protocol (RSVP) [11]. At the ATM side, both permanent and switched virtual channels (PVC and SVC) can be established. These can support traffic of any service category (i.e., CBR, VBR and UBR for UNI 3.x and VBR-rt and ABR for UNI 4.0) and traffic descriptor in point-to-point but also point-to-multipoint configurations. User terminals connected to the access network, employ RSVP for reserving network resources. RSVP differs from UNI signaling in many ways: First of all, RSVP is designed to operate with IP and therefore it can not be interpreted by ATM switches. Moreover, issues like dynamic QoS renegotiation, multicasting, call setup origination, call release are handled differently by RSVP compared to UNI signaling. Therefore, a major task of the ED is to act as signaling interworking unit between the ATM-based core network and the IP-based access network. To do so, the control plane of the ED must contain both UNI and RSVP protocol stacks. Furthermore, the appropriate interworking functions that map RSVP messages to UNI signaling messages and vice versa must be also present. As it is explained in section 4 below, the ED has to translate from IP to ATM addresses. For this purpose, the IETF's Next Hop Resolution Protocol (NHRP) is used shown in Figure 2. Finally, efficient interoperability between RSVP and ATM signaling cannot be achieved without the translation of RSVP traffic and QoS parameters to the corresponding ATM ones. Therefore a flow admission control mechanism must be present in the ED to decide whether a call can be serviced or not. For interoperability reasons the flow control algorithm should be in harmony with the Connection Admission Control (CAC) algorithm employed in the ATM part of the architecture.

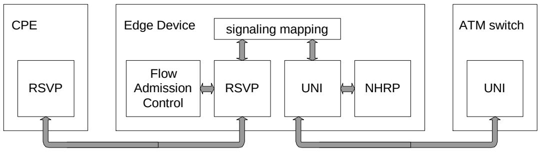

Figure 2: Control plane

2.2 User Plane

Figure 3 illustrates the user plane of the proposed architecture. As we can see, the user plane of the user terminal, and the ATM switch remain simple and unmodified. This fulfills the requirement for minimum changes on the already developed part of the network. The user plane of the ED is composed of the protocols required to map IP packets to ATM cells (e.g. IP, AAL5 etc.) as well as a set of processes executing a number of tasks required by the particular interworking environment. While a typical Internet router requires only the existence of the IP protocol and a look-up table that correlates destination IP addresses with physical interfaces, the ED considered here has additionally to reserve network resources per IP flow. To do so, the mechanisms required are those of policing, Differentiated Services and traffic shaping. The policing mechanism monitors each IP flow and checks whether the incoming traffic profile conforms with the traffic parameters agreed for the particular flow during the initial setup of the connection. The objective is to protect the network from intentional or unintentional deviations from the negotiated traffic parameters that can adversely effect the QoS of other users. In principle, policing will be based on the Generic Cell Rate Algorithm (GCRA) of I.371 [9], but in this case the incoming data will be rather variable length IP packets than fixed size ATM cells. There is an additional reason for using a policing mechanism in the ED that will become clear in the following example. Let us assume, that a user who has requested a CBR IP flow of 64Kbit/s for transferring voice, transmits IP packets which a rate higher than the negotiated. In the absence of policing in the ED, the UPC (Usage Parameter Control) function of the ATM switch would drop cells indiscriminately, that is cells belonging to different IP packets would be rejected. Assume now that one out of one hundred cells is dropped and that each IP packet is about 1000 bytes which is approximately equal to 20 cells. This means that one out of 100/20=5 IP packets would not be re-assembled in the receiving ED while policing in the ED would result in the loss of one out of one hundred IP packets in the sending ED.

A mechanism of equal importance implemented in the user plane of the ED is IDS (Internet Differentiated Services). IDS is used to assign priorities to the IP packets according to the service category they belong to. Priority differentiation is achieved by virtue of the value of the TOS (Type Of Service) byte of the IP header. IDS enabled routers handle IP packets according to their priority. The ED is able to recognize the particular flow an incoming IP packet belongs to and based on the previously received RSVP information for the particular flow sets the value of the TOS byte accordingly. If the incoming IP packet does not belong to any established IP flow it is handled as best effort. IDS is particularly important for existing Internet routers which are completely unaware of QoS issues and therefore are unlikely to be upgraded in the near future so that they can support resource reservation. IDS is very simple and easy to implement which makes us to believe that IDS enabled routers will be available soon. Therefore, the support of IDS in the ED offers not only backwards and but also future compatibility with Internet.

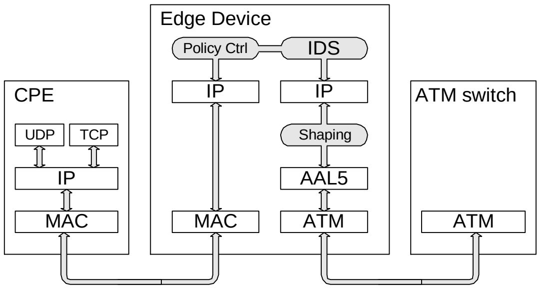

Figure 3: User plane

3. INTEROPERABILITY BETWEEN RSVP, IDS AND ATM SIGNALING

We describe the ED actions necessary when a reservation request (i.e. an RSVP reserve message), originated from a user within an access network, is received. At this point, there are a number of decisions that have to be taken in the ED. First of all, the ED must decide whether the user has the appropriate administrative permissions that allow him to request the particular reservation. This means that the ED must be capable to identify the user id. Therefore, secure cryptographic methods must be used for user identification. Supposing that the user authorization is successful, the ED must decide whether there are enough internal resources for the realization of the specific reservation. Usually this operation is executed by the flow admission control module of the ED. Assuming that there are sufficient internal resources in the ED, the request must be forwarded to ATM network. At this point there two major design decisions that must be made: the former is related to the way RSVP signaling messages pass through the ATM network towards the destination ED (see 3.1), while the later concerns the way IP traffic is concentrated at the ED so that signaling message processing is minimized in the ATM network (see 3.2).

3.1 Passing RSVP messages through the ATM network

We consider two approaches on how to transfer the RSVP messages from one ED to another ED. In the first approach, the RSVP is terminated in the EDs. This means that whenever an ED receives an RSVP message from the access network, it translates it to the appropriate UNI signaling message. The ED that receives the incoming UNI setup message recreates the original RSVP message and forwards it to the access network. Since there is a number of differences between RSVP and ATM signaling, recreation of the original RSVP message becomes difficult if not impossible. For example, several fields present in RSVP messages, (e.g. IP addresses, destination and source port etc.) can not be mapped directly to any of the fields of the UNI setup message. Additionally, RSVP requires hosts to periodically send RESV messages in order to keep an existing reservation alive (soft state of RSVP.) On the other hand, the resources in ATM are reserved for the duration of a connection, which must be explicitly and reliably deleted. Hence, there is no straightforward way to map RSVP “refreshing” messages to UNI messages. Moreover, there is no equivalent UNI message for the PATH message of RSVP which is sent periodically by the sender. For these reasons, we exclude the approach of terminating the RSVP in EDs.

Alternatively, RSVP messages can pass transparently through the ATM (RSVP tunneling) network. Whenever an ED receives an RSVP message it processes it and pass it to the destination ED through an open ATM channel similarly to the normal user traffic. The receiving ED recognizes that a particular IP packet carries an RSVP message by examining the protocol identifier field of the IP header. This approach is much more straightforward than the former one and therefore it is adopted in the design of the ED.

3.2 Processing RSVP messages

Any RSVP request that reaches an ED can be of the following type:

- A request for a new reservation to be established
- A request to tear down an established reservation

- A request to modify the QoS of an existing flow

The ED has many alternatives to fulfil any of these requests. For example, a request for a new IP flow reservation can lead the ED to request a VC from the ATM network or to aggregate this new flow with an existing one within the same ATM channel. The advantages and disadvantages of each approach are discussed in the next section.

In the following, we will present two possible scenarios that demonstrate the interoperability of the RSVP with the IDS and with ATM signaling. Both scenarios assume the network infrastructure and the protocol stacks presented in Section 2. Our approach is similar but not identical to the "RSVP over ATM" approach presented in [12].

3.2.1 Scenario 1

We assume that for any data interchange between two EDs switched virtual channels are always used. Moreover, the resource reservation at the ATM layer is done by the receiving ED. The scenario starts when the sending host, attached in an access network, generates an RSVP PATH message that characterizes the traffic and which is received by the ED (ED_s) that serves the sending host. At this point the ED_s does not call any function of the admission or policy control since the real reservation will be done by the other ED when it will receive the corresponding RESV message. In fact, the ED_s examines the contents of the PATH messages and checks whether there are sufficient free resources that could be reserved for that flow. If there are not sufficient resources, the ED_s modifies slightly the PATH message by adding in the AddSpec the amount of its available resources. This information will be used later by the receiving host that will make the real reservation, so that the QoS parameters of the reservation are less or equal to the available resources of the ED. After that, the ED_s must transfer the PATH message to the other ED (ED_r) that serves the receiving host. To do so, ED_s uses an existing bearer connection between the two EDs if there is at least one available. If a bearer connection between the two EDs already exists, the ED_r receives the PATH message and like in the ED_s case, ED_r does not call any function of the admission control module, but only calculates the amount of available local resources. If this value is less than the value of AddSpec field, the AddSpec is updated with the new value. The next step in this scenario is the transmission of the PATH to the receiving host through the access network. Normally, the receiving host responds by sending an RESV RSVP message to ED_r. The RESV message is actually a request for a new reservation and it contains the desired QoS parameters for that IP flow. Upon the reception of an RESV message, the ED_r calls the procedures of the local admission control module in order to decide whether the new request can be accepted or not. The next step for the ED_r is to reserve a path across the ATM network for the requested IP flow. At this point there are two options; the ED_r either tries to aggregate this IP flow with another one that exists between ED_r and ED_s into the same ATM virtual channel (multiple flows in one SVC), or requires a new connection from the ATM network by sending a UNI setup with QoS parameters equal to the ones requested by the receiving host (one flow - one SVC). The two approaches and their advantages and disadvantages are discussed in the next section. In either case, the ED_r transmits the RESV message and any future RSVP messages for that IP flow, are routed through the ATM channel used to convey user traffic. An alternative way would be to pass the RESV message through the ATM channel through which the PATH message was originally passed. However, in this case it would be more difficult for the ED_s to map the specific IP flow with the pair of VPI/VCI that would be used for sending the user traffic. At this point, the reservation in the ATM network has been realized and the RESV message has been received by the ED_s. The ED_s will call the flow admission control module to decide whether the new reservation can be accepted or not. In the normal case, the flow admission control of ED_s accepts the request and the RESV message is transmitted to the sending host indicating that the reservation is complete. Moreover, the ED_s creates a new IDS instance that is based on the QoS parameters of the new IP flow. The IDS instance monitors all the incoming packets that belong to this IP flow and puts the appropriate priority value in the TOS byte of the IP header. Figure 5 illustrates the messages exchanged in the various network components for a successful reservation for this scenario.

3.2.2 Scenario 2

The physical network infrastructure for the 2nd scenario is the same as the one of the 1st scenario. The differences between the two scenarios exist mainly in the handling of the permanent and switched connections of the network. In this scenario, permanent connections between neighboring EDs are required (see Figure 4).

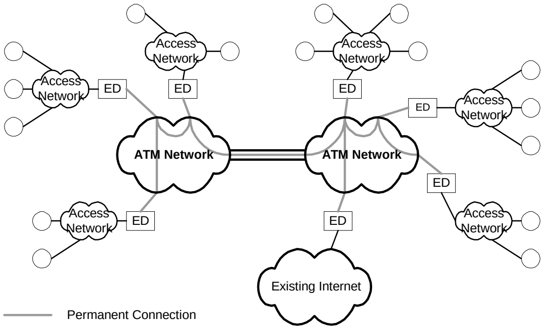

Figure 4: Network configuration for the second scenario

Since the EDs have routing capabilities, the permanent connections between neighboring EDs compose an IP network that allows the intercommunication of any hosts attached to any access network. From now on, we will refer to the IP network that is based on PVCs as "standard network". The standard network is used to carry the best effort traffic that is generated by the existing Internet applications (e.g. ftp, WEB browsers). This is in contrast to the assumptions made for the first scenario where for any data communication only switched connections are used. However switched connections are not excluded at least for QoS sensitive applications like video conference or video on demand. Let us assume now, that a user attached to an access network is the originator of the traffic of a guaranteed service. The user application will generate a PATH message which characterizes the traffic and which will be received by the ED_s. Like in the first scenario, the ED_s calculates the value of the AddSpec field of the PATH message and then it communicates the PATH message to the ED_r through the standard network. Communicating the PATH message through the standard network causes two problems. Firstly, the standard network which is used to carry the best effort traffic, may be overloaded and consequently there is the possibility that the PATH message will never reach at the ED_r. In order to overcome this problem, the ED_s increases the priority of the IP packet that carries the PATH message by putting a high priority value in the TOS byte of the IP header. Doing so, all intermediate EDs, which are IDS enabled, serve the PATH message immediately and therefore the probability of losing the PATH message decreases. Secondly, all the intermediate EDs, which are also RSVP enabled, will try to process the PATH message. In order to overcome this undesirable situation, we *demand* from the EDs to process only those RSVP messages that originate from the access network and are destined to the ATM network and vice versa. All the RSVP message that come from and must be forwarded to the ATM network are not processed by the local RSVP module but are simply routed towards their destination like normal user traffic. Continuing the description of the scenario, the ED_r receives the PATH message and after the value of the AddSpec field is updated, it transmits the message to the receiving host. The receiving host responds by sending an RESV message filled with the desired QoS parameters to the ED_r. At this point, the ED_r makes the required local reservations by calling the flow admission control module and creates a new IDS instance, but it does not make any reservation in the ATM cloud as it does in the first scenario. Moreover, ED_r communicates the RESV message to the ED_s through the standard network. The ED_s detects that the incoming RSVP does not belong to any active RSVP session and therefore a new reservation must be established. Again, the ED_s has two options; it either sets-up a new virtual channel, based on QoS parameters of the RESV message, or it aggregates the new IP flow with others that are destined to the ED_r. Furthermore, ED_s reserves the required local resources, creates a new IDS instance based on the QoS parameters of the RESV message and finally forwards it towards the sending host.

At this point, we observe another major difference between the two scenarios; in the first scenario, the reservation at the ATM layer is made by the receiver's ED while in the second scenario by the sender's ED. This fact has a positive impact in the multicast issue for the second scenario. Let us suppose that we have one sender sending the same data stream to multiple receivers. In the first scenario, each receiver's ED opens a point-to-point ATM connection with the sender's ED and the multiplication of the traffic is done at the IP layer in the sender's ED. It is obvious that in this case, the multicast capabilities of ATM are not exploited and that network resources are wasted. On the other hand, the second scenario allows us to utilize the ATM multicast capabilities. The sender's ED delivers the PATH message to all receivers. For the first arriving RESV message, the sender's ED responds by setting up a new ATM connection towards the ED of the user that firstly responded. For any other incoming RESV messages that belong to the same RSVP session, the sender's ED can utilize the multicast features of ATM by adding a new party to the existing session.

Figure **6** and Figure 7 illustrate the message flow diagrams for the unicast and multicast cases for the second scenario.

4. ADDRESS RESOLUTION, VCI MANAGEMENT, AND IP FLOW AGGREGATION

The architecture described in this paper aims to make optimal use of the features of a public ATM network, in particular Switched Virtual Circuits (SVCs), to avoid bottlenecks in a future Internet infrastructure with QoS support. In order to take benefit from SVCs, two more detailed questions have to be answered. The first issue is to find the addresses to which SVCs are established and the second is to determine the right number and attributes of these SVCs.

4.1 Address Resolution

Address resolution is required since the end systems involved in multimedia services are identified using the addressing scheme of the Internet (IP addresses). The VCs among the EDs, however, have to be established towards ATM end systems of a public ATM network. Therefore some translation between IP addresses and ATM addresses has to take place. There is already a well-established solution which was designed for the usage in local area ATM networks, that is the ATMARP (ATM Address Resolution Protocol) which is used in the Classical IP over ATM standard [5]. However, this approach just considers communication among hosts within a single Logical IP Subnetwork (LIS), such that communication to other IP sub-networks still has to take place over an IP router. Moreover, Classical IP over ATM requires all hosts to be attached directly to ATM which is not the case in the target configuration assumed here. Therefore, a more adequate candidate for address resolution is the Next-Hop Resolution Protocol (NHRP) [6] which was explicitly designed for the establishment of shortcut routes through Non-Broadcast Multiple Access (NBMA) networks like ATM. Given an IP address of an end system, NHRP finds the address of a router which is close to the given IP address. Since the EDs include a router functionality, the NHRP protocol can be employed in order to find the ATM address of an ED through which an end system IP address can be reached.

4.2 VC Management

An even more important issue is management of VCs. Individual VCs can be used to support a single IP traffic flow, or several different IP traffic flows can be multiplexed into the same VC. The ELISA project has not yet decided on a specific strategy; it may be very well possible that the trial platform will offer selectable VC management strategies. Below the main options for VC management are discussed briefly.

4.2.1 One Flow - One SVC

The most simple approach to VC management is a one-to-one correspondence between application-requested IP flows and ATM SVCs. In this case the implementation of the admission and policy control of the ED is relatively easy since the connection admission control algorithm of the ATM network can be used to deal with acceptance of reservation requests. However, there are also serious drawbacks. The user will experience rather high setup latency if she has to wait until a call setup in the public ATM network between EDs is completed. It is difficult to deal in this approach with a dynamic change of QoS attributes (which is a feature of RSVP). Advanced ATM signaling features like dynamic bandwidth modification may be required here.

4.2.2 Multiple Flows in One SVC

There are significant advantages of multiplexing together several IP flows in an ATM VC. First of all, this can be done without any additional overhead since the IP protocol supports the multiplexing and demultiplexing functions. Moreover, signaling latency may be reduced by pre-establishment of VCs or of spare capacity within a VC. Dynamic QoS changes can be realized in many cases by local resource assignments in the EDs. Nevertheless, at some points also the aggregated VCs may require a dynamic modification. However, an intelligent VC management function can remove or minimize the negative effects of such changes on the users (e.g. latency). One idea is for instance an overlapping setup of a new SVC as long as an old SVC is still in use such that the change of the VC attributes is achieved by a simple routing update. The main drawback of this strategy is of course the higher complexity of the EDs.

An interesting aspect is here that several parallel VCs may be required carrying aggregated traffic of different classes (e.g. constant vs. variable bit rate). This concept seems to be suitable to be combined with an implementation of the IDS framework on top of ATM.

4.2.3 Separated RSVP Control

Independent of the choice between the two alternatives from above, there are several different ways how to handle RSVP control messages (Path and Resv). Either the RSVP messages are multiplexed with the data traffic (as usual in IP networks) or there exist special VCs between EDs which are dedicated for QoS control. The second approach has the advantage that a stable and rapid transport of control messages is ensured independently of the user traffic.

5. PERFORMANCE ISSUES

As evident from the previous presentation the ED provides far more complicated functionality than typical IP routers. Therefore the ED performance becomes a critical issue as it may easily become the bottleneck entity of the entire architecture. Our approach to performance evaluation is based on the separation between User and Control Plane functions. In the User Plane provision should be taken so that the ED accomplishes its task of providing several service categories with the corresponding QoS guarantees to incoming flows.
To do so, the ED should at first be able to predict the statistics of the incoming flows which maps onto the accurate characterization of traffic profiles. It should then be able to pass this information to ATM through signaling which means that traffic should be described in terms of the well known ATM traffic parameters specified in [9] (i.e. PCR, SCR, MBS). The next step for the ED is to provide those mechanisms that guarantee conformance of the flows to the parameters agreed with the ATM side (e.g. traffic shaping, access priorities and buffering priorities per service class etc.). These mechanisms should rely on the principles applied in the ATM side for CAC and policing (i.e. the GCRA). It becomes apparent that traffic characterization on one hand and development of ATM compatible traffic handling mechanisms at the ED on the other are the main issues related to the performance of the User Plane.

The multiplicity of service classes foreseen for the future Internet (real-time streaming applications, real-time block transfer applications, non real-time applications), result in various traffic patterns and often conflicting quality requirements. Therefore it is natural to expect that the corresponding traffic models will appear even more complex than the already complex models developed using measurements from nowadays Internet applications [7, 8]. The above remark also applies to the traffic handling mechanism a situation that further strengthens the approach followed in this paper for combining RSVP at the Control Plane with Differentiated Services at the User Plane.

To ensure that the advantages of using RSVP at the Control Plane of the ED are not altered by excessive delay in resource reservation, modeling of the control plane operations is required. Using queuing models we can represent the RSVP and signaling traffic flows among the different entities comprising the architecture under consideration. Every request from the user side generates an RSVP message at the CPE which in turns generates an RSVP and ATM signaling scenario. An RSVP and signaling scenario is described by the numbers of messages exchanged between the participating physical entities. Messages are defined according to the protocols applied. Using these models we are able to define the bottleneck entities as well as to study alternatives of message flows towards performance enhancement. While user plane traffic needs to be described statistically by using the notion of self-similarity, request arrivals can follow a simple Poisson distribution.

6. CONCLUSIONS

This paper has shown that significant issues in the transition to a future multimedia-suitable Internet can be addressed without introducing completely new network concepts. Instead, a well-designed additional network component (the edge device described here) can provide a smooth integration and interworking of existing concepts from the Internet world (like Integrated Services, Differentiated Services) and from telecommunication networks (like ATM traffic classes, ATM switched connections).

The general approach taken here may be seen as not very exciting compared to the design of new innovative network concepts. Nevertheless, it is symptomatic for the current situation in the convergence process of information and telecommunication technology. The necessary technologies exist and provide all the required features for very powerful innovative services. But a huge amount of incompatibilities and overhead is introduced when all these technologies are just combined in a naive manner. Significant additional effort has to be spent to design intermediate devices like the edge device described in this paper, to analyze the behavior of large-scale networks based on such combined technologies and to carry out practical experiments with prototype implementations. The aim of this type of research is to make new network features available at minimal cost just by making optimal use of the infrastructure which already exists.

APPENDIX

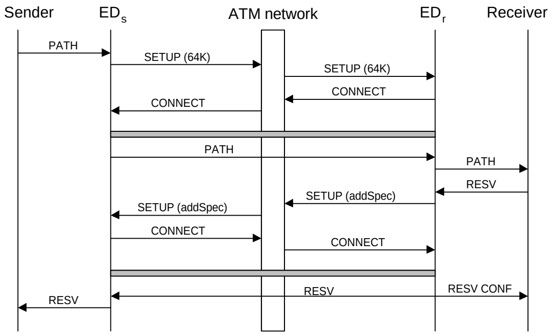

Figure 5: Flow message diagram for the first scenario

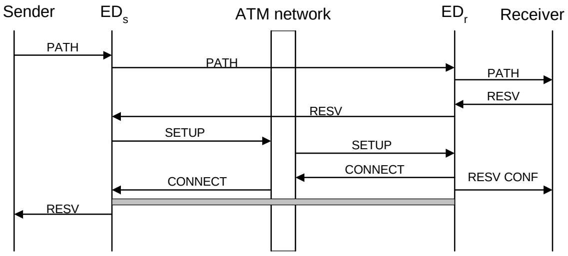

Figure 6: Flow message diagram for the second scenario

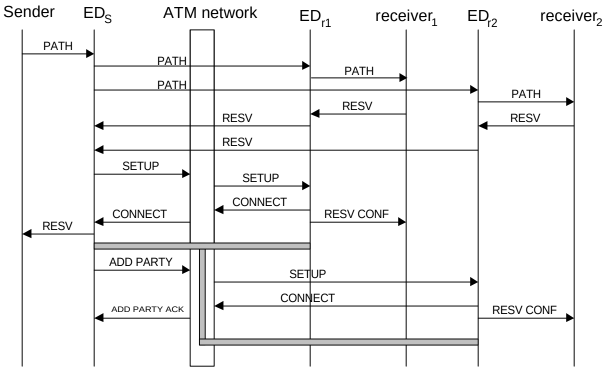

Figure 7: Flow message diagram for the second scenario (multicast)

REFERENCES

[1] D. Clark, S. Shenker, L. Zhang, “Supporting real-time applications in an integrated services packet network: architecture and mechanism”. Proc. ACM SIGCOMM, August 1992, pp. 14-26

[2] R. Braden, D. Clark, S. Shenker: “Integrated services in the Internet architecture: an overview”. RFC 1633, IETF, 1994.

[3] Y. Bernet, R. Yavatkar, P. Ford, F. Baker and L. Zhang, “A framework for end-to-end QoS combining RSVP / Intserv and differentiated services” , Internet Draft, IETF, March 1998.

[4] K. Nichols, S. Blake, “Differentiated services operational Model and definitions”, Internet Draft, IETF, February 1998.

[5] M. Laubach, “Classical IP and ARP over ATM. Request for Comments (Proposed Standard)” RFC1577, IETF, 1994.

[6] J. Luciani, D. Katz, D. Piscitello, B. Cole and N. Doraswamy, “NBMA Next Hop Resolution Protocol (NHRP)”, Internet Draft, IETF, 1998.

[7] T. C. Know, "Residential Broadband Internet Services and Applications Requirements", IEEE Commun. Mag., vol. 35, June 1997, pp. 76-83.

[8] M. E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", IEEE/ACM Trans. Networking, vol. 5, December 1997, pp. 835-846.

[9] ITU-T Recommendation I.371 - “Traffic control and congestion control in B-ISDN”, 8/96

[10] ATM Forum. “ATM User Network Interface (UNI) Specification Version 3.1”. Prentice Hall, June 1995.

[11] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin. “Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification”. RFC 2205, September 1997.

[12] R.Guerin, D.Kandlur, L.Li, V.Srinivasan, L.Berger. “Support of Shortcuts for RSVP Flows Over ATM.” Internet Draft, IETF, 1997