Implementation of IP Multicast over xDSL/ATM-based Access Networks

A. Banchs, M. Gabrysch, S. Vannay, T. Dietz, B. Lange, H. J. Stüttgen

NEC Europe Ltd, Computer and Communication Research Laboratories Heidelberg

adsl@ccrle.nec.de

Abstract: *ADSL technology provides a new platform for delivering broadband services to homes and small businesses, supporting a wide variety of high bandwidth applications. Since Multicast is essential in many of these high bandwidth applications, Multicast support will be a key requirement for the success of ADSL networks. In this paper we present two IP Multicast architectures for ADSL: the Multicast Gateway and Multicast over PVC. We evaluate architectural strengths and weaknesses, provide performance measurements and discuss accounting considerations. Both architectures have been implemented and will be used in an ADSL field trial at RWTH Aachen.*

1 Introduction

The classical technologies to access the global Internet or a corporate network have become insufficient to satisfy the growing demands of the end-users due to their limitation in bandwidth. With 56 kbit/s, the analog modems have reached their physical limits. Even with ISDN the bandwidth can only be increased by using multiple channels. xDSL technology is a solution that overcomes this bandwidth limitation by using the existing telephone-access network infrastructure. Of all the members of the xDSL family, ADSL is currently the most deployed[1]. With ADSL, bandwidths of up to 8 Mbits/s downstream (towards the customer) and up to 768 kbits/s upstream (towards the Central Office) can be achieved.

Our research laboratory has been actively involved in an ADSL field trial in Aachen, Germany, performed in a collaboration of Deutsche Telekom AG, RWTH Aachen and NEC. This field trial started in February 98, and has recently been extended to two additional years in the framework of a DFN proposal. The main goals of the field trial were to gain experience with ADSL technology and to study the behavior of the users in such a high-speed scenario. From this experience, we have been able to identify further work that needs to be done in this area.

---

[1] Note that even though this paper is focused on ADSL, all the concepts here discussed can also be applied to any other xDSL technology.

One important requirement detected during the field trial has been the need for support of group applications based on point-to-multipoint and multipoint-to-multipoint. From the feedback of the field trial users, it can be seen that a high-speed scenario such as ADSL encourages the use of applications like video on demand, video conferencing and distributed games, and it provides a suitable scenario for teleworking and distance learning. Most of these applications require point-to-multipoint or multipoint-to-multipoint communication, which is much more efficiently handled by IP Multicast than by IP Unicast. As a consequence, IP Multicast is a very important requirement for ADSL.

In this paper we present two architectures that have been implemented for providing IP Multicast over ADSL. In section 2 a basic background about ADSL networks is given. In section 3 the two Multicast architectures for ADSL are described and performance results are presented for each one. Section 4 discusses how accounting could be performed for these Multicast architectures in a real scenario. Finally, section 5 concludes the paper.

2 ADSL Networks

This section provides some background information about ADSL networks. For further information the reader is referred to [1].

2.1 ADSL-based Broadband Network Architecture

The ADSL-based broadband network architecture [2] consists of the following subnetworks shown in Figure 1: the customer premise network, the ADSL access network, the Regional Broadband Network (RBN) and the Network Service Provider (NSP) network.

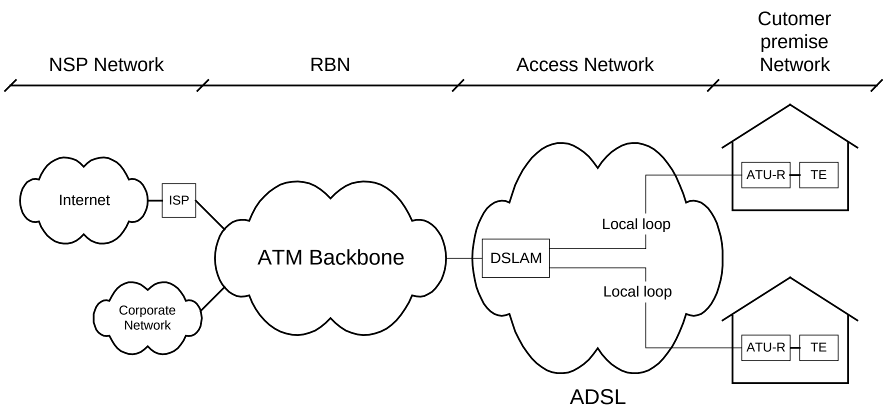

Figure 1 ADSL-based Broadband Network Architecture

2.2 PPP over ATM over ADSL

The ADSL end-to-end service interoperability model [2] is based on an end-to-end ATM network between the end-user and the NSP (Figure 2). This ATM over ADSL architecture preserves high-speed characteristics and provides QoS guarantee in the ADSL environment without changing protocols.

Once ATM layer connectivity is established between the end-user and the NSP, the session setup and release phases at the link level and network level are established using PPP [3]. PPP over ATM provides, among other features, AAA (Authentication, Authorization and Accounting) and interaction with RADIUS servers. The model used for PPP over ATM in ADSL networks is the IETF proposed standard for PPP over AAL5 [4].

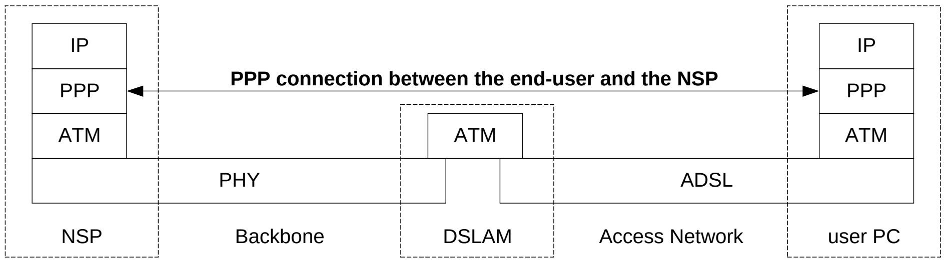

Figure 2 PPP over ATM over ADSL

2.3 SVC Support

For the connection between the end-user and the NSP there are two possibilities: ATM PVCs and ATM SVCs. The advantage of using ATM SVCs is that it reduces per-user provisioning and it supports dynamic access to multiple NSPs. However, ubiquitous ATM SVC service is currently not available and the operation infrastructure for it is incomplete. In addition, many NSPs are currently not prepared to terminate large numbers of ATM SVCs.

For Unicast the gateway architectures (namely, L2TP Access Aggregation and PPP Terminated Aggregation [5]) have been developed as near-term options that provide many of the benefits of SVCs in a PVC environment. These gateway architectures enable the customer to select the NSP via the PPP layer rather than by switching at the ATM layer. The goal of these architectures is to allow ADSL deployments to progress while the industry develops a mature set of ATM capabilities for a long-term solution based on SVCs. An overview of this evolution towards a SVC environment is given in [6]. A similar approach would also be desirable for the Multicast case: a short-term solution based on PVCs and a long-term solution based on SVCs.

3 Multicast Architectures for ADSL

In this section we present two Multicast architectures for ADSL: the *Multicast Gateway* and *Multicast over PVC*. The *Multicast Gateway* is an easy way to provide Multicast to ADSL users, but it does not scale well for a large number of users.

Another possibility for providing Multicast over ADSL networks is MARS/MCS [7], the proposed standard by IETF[2] for Multicast in ATM networks. This solution presents a good behavior regarding scalability but it relies on SVCs, which are currently not available. Hence a scalable solution for Multicast in the PVC based access networks of the near future is missing.

In order to cover this need we have developed our own proposal: *Multicast over PVC*. This architecture is based on MARS/MCS, adapting it to the PVC environment of current ADSL networks. It has been designed in order to enable a smooth transition to SVCs when they become available on a broader basis in the future.

A more detailed discussion of architectural strengths and weaknesses of the *Multicast Gateway*, *Multicast over PVC* and other possible architectures for providing Multicast to ADSL networks can be found in [9].

3.1 Multicast Gateway

The *Multicast Gateway* [10] is a Multicast client-server application that provides access to Multicast streams to hosts located in Non-Multicast networks through Unicast tunnels dynamically established between the application running at the server and at the client. This application was originally developed for dial-up ISDN clients at University of Mannheim. It has been adapted to the special needs of ADSL at NEC CCRLE Heidelberg.

The *Multicast Gateway* architecture (Figure 3) consists of two components: a server and a client part. The client mainly provides a user interface for the communication with the server. The server receives packets from a Multicast group, replicates them and forwards them to the clients that have requested packets from that group. The request of a session by a client works in the following way: The server listens to the session descriptions being announced with the session announcement protocol SAP [11], which is commonly used in the Mbone. When a client gets connected to the server, the server sends him the available sessions

[2] Even though MARS/MCS is the IETF proposed standard, other architectures have been proposed for supporting Multicast over ATM. These architectures could also be used in ADSL networks. In [8] a comparison between these architectures is provided.

received through the SAP. Then the client chooses a session to join, and from this moment on, the server will transmit this session to the client via Unicast.

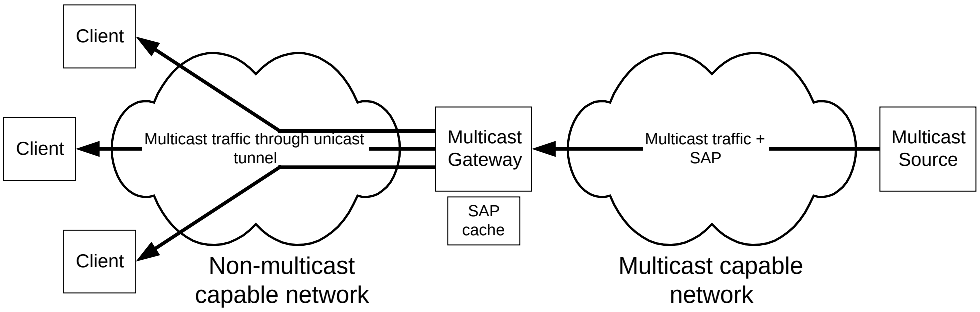

Figure 3: Multicast Gateway Architecture[3]

3.2 Multicast over PVC

The *Multicast over PVC* approach is a solution that, even though it requires a bigger implementation effort than the *Multicast Gateway*, is preferable since it presents a better behavior in terms of packet replication efficiency and scalability regarding number of clients.

The *Multicast over PVC* architecture is based on MARS/MCS: for address resolution a MARS server is used and for delivering Multicast data a MCS server is used. Note that even though in the explanation given here only one MCS is considered, the architecture could be easily extended to several MCS serving either different Multicast groups or different hosts each one [12].

When a client wants to join or leave a session, it issues a MARS_JOIN or MARS_LEAVE message to the MARS (using statically configured PVC between the MARS and the host). MARS uses this information to keep track of all the group membership information in the cluster.

When a client has a Multicast packet to transmit, it sends it to the MCS, which is responsible for delivering it to the members of the corresponding group. This communication with the MCS is also done using statically configured PVCs. The MCS obtains then the membership information from the MARS through a MARS_REQUEST message; the reply from the MARS, one or more MARS_MULTI messages, contains the mapping between the requested Multicast group and the ATM addresses of the hosts that have joined that group. This mapping is stored in the local cache of the MCS, and will be updated with new joining and leaving information, that is forwarded by the MARS to the MCS.

[3] Note that the Multicast Gateway can be located anywhere inside the Multicast capable network and not necessarily in the boundary between the Multicast and the Non-Multicast capable networks.

Once the MCS has obtained the group membership information, it has to take a decision about the way of delivering the Multicast data. There are two possible approaches for data forwarding:

- A PMP PVC from the MCS to all the hosts of the cluster
- Multiple PP PVCs from the MCS to each host

The PMP PVC approach (Figure 4) broadcasts the Multicast data to all hosts[4], regardless whether they have joined the group or not: the PMP PVC is statically set to reach all hosts, so it cannot be selected which hosts it has to be delivered to. This approach is very efficient in the data replication, which is done at the ATM level, but, on the other hand, user's accesses might get flooded with useless Multicast traffic.

The multiple PP PVCs approach (Figure 5) allows the MCS to selectively send a replicated copy of the packet to each of the PVCs corresponding to hosts that have joined the group. This approach avoids flooding user's accesses with Multicast data they do not require, but has the disadvantage of requiring a high cost for the replication, which is done by the MCS at the network level.

Note that in this approach, the PVCs used to transmit the outgoing Multicast traffic from the hosts to the MCS can be reused to transmit the incoming Multicast traffic in the opposite direction.

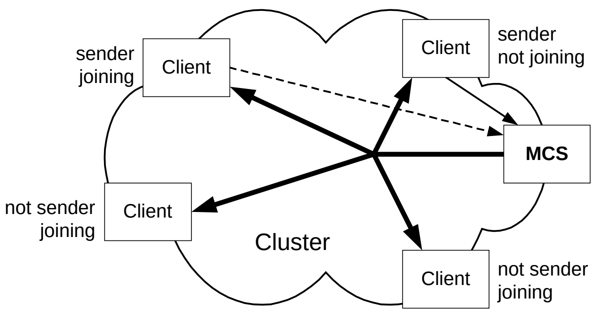

Figure 4: Delivery of Multicast data using PMP PVCs to all hosts.

[4] Note that this approach is based on the *broadcast and filter principle*: end hosts are responsible for filtering the data corresponding to Multicast groups they have not joined. Multicast on Ethernet networks is also based on this principle, but while in the Ethernet the filtering is performed by the adapter hardware, in the PMP PVC approach it is performed by the terminal software (i.e., only in the latter case CPU is consumed for the filtering operation).

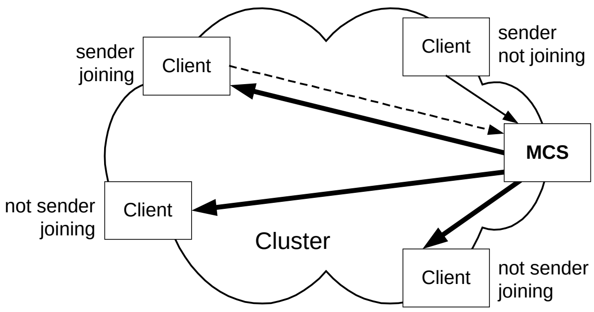

Figure 5: Delivery of Multicast data using multiple PP PVCs.

3.2.1 Multicast Delivery option: Decision criteria

The decision of which is the most appropriate way to replicate and forward the packets of a Multicast group (with a broadcast PMP PVC or with multiple selective PP PVCs), can be based on different criteria (or a combination of them). Possible criteria for such a decision are:

- *Number of hosts that have joined the group*

In the case of many hosts per group, the broadcast PMP PVC approach is the best option. Since a lot of hosts have joined the group, not too many users will have their access bandwidth wasted with useless traffic; moreover, the selective PP approach in this case would have a very high cost in replication.

On the other hand, in the case of only few hosts per group, the selective PP approach will be the best option: the replication won't have a very high cost, and the broadcast PMP option would imply waste of bandwidth.

- *Avoid flooding user's accesses*

Flooding a user access with Multicast streams the user has not joined should be avoided. Thus, another possible criterion could be to use always the more efficient approach (i.e. the PMP PVC) except when that means flooding a user's access. The drawback of this option, however, is that a bandwidth control mechanism is required for detecting when a user's access is being flooded.

- *Preassigned bandwidth to content providers*

Another criterion could be to divide the access bandwidth between a bandwidth assigned to the user and a bandwidth assigned to content providers, in such a way that each content provider has statically assigned a certain bandwidth of the user access, for which it is charged. Then, since there is no contention for this preassigned bandwidth, the PMP PVC approach can be used without risk of flooding the user access.

3.3 Performance Results

The performance of the Multicast Gateway and the Multicast over PVC architecture (both PP and PMP options) were evaluated in terms of throughput and losses as a function of the number of clients. The purpose of these tests was to gain insight into the impact of the number of clients and the replication level[5] to the performance.

Figure 6 shows the experimental setup of the measurements. For the machine forwarding the Multicast stream (i.e., the Multicast Gateway and the MCS, respectively), a Pentium 200 MHz MMX with Linux 2.1.90 and experimental Linux-ATM 0.35 software was used[6]. Note that the use of experimental ATM drivers might have some affect in the results, especially in the case of the MCS when high throughputs (i.e., near the 155 Mbps maximum capacity of the ATM NIC) are involved.

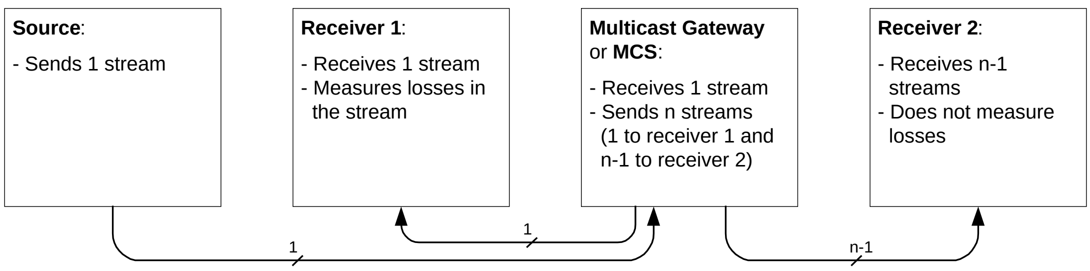

Figure 6 Experimental setup for the performance tests.

The results of the measurements are presented in Figure 7 and Figure 8. Figure 7 shows the average loss rate of the Multicast Gateway as a function of the bandwidth of the Multicast stream for 1, 2 and 5 clients.

Figure 8 shows the same results for the Multicast over PVC architecture. Note that in the 5 clients with 20 Mbps case, the MCS has to deal with a total of 120 Mbps. In the Point-to-Multipoint case the losses at the MCS do not depend on the number of clients, because the replication is not performed by the MCS but at the ATM level.

Table 1 shows the approximate bandwidth, which can be supported without losses in each of the architectures (Multicast Gateway, Multicast over PVC – PP approach and Multicast over PVC – PMP approach) with 1, 2 and 5 clients respectively. In these table it can be observed that the performance of replication at the MCS is about 10 times better than the

---

[5] The replication is done at the application level in the case of the Multicast Gateway, the network level in case of the MCS PP approach and the ATM level in case of the MCS PMP approach

[6] With a more powerful machine a better performance is expected. However, the primary goal of this test is to compare different architectures. The hardware dependence of the architectures will be object of further studies.

Multicast Gateway. This is a consequence of the fact that the *Multicast Gateway* replicates the packets at the application level (end-user space), and that has a much higher overhead than the network level replication at the MCS (kernel).

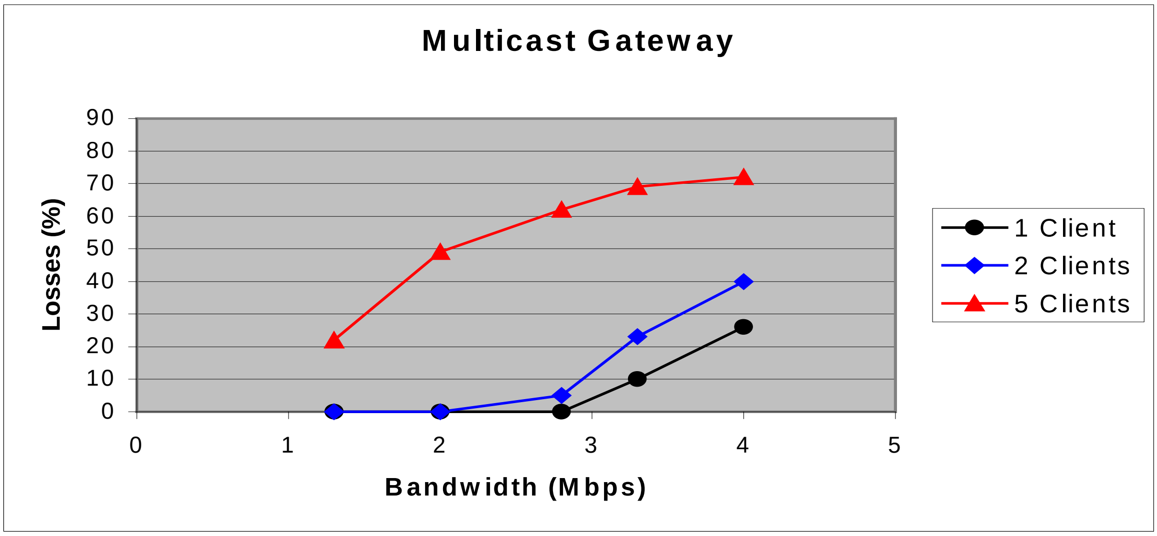

Figure 7: Multicast Gateway Performance Results

It can be observed that losses increase considerably with the number of clients for both the Multicast Gateway and the PP approach of the Multicast over PVC architecture. Hence the Multicast Gateway architecture is not scalable with the number of clients. In the Multicast over PVC architecture, scalability will only be achieved with a proper combination of the PP and the PMP approaches.

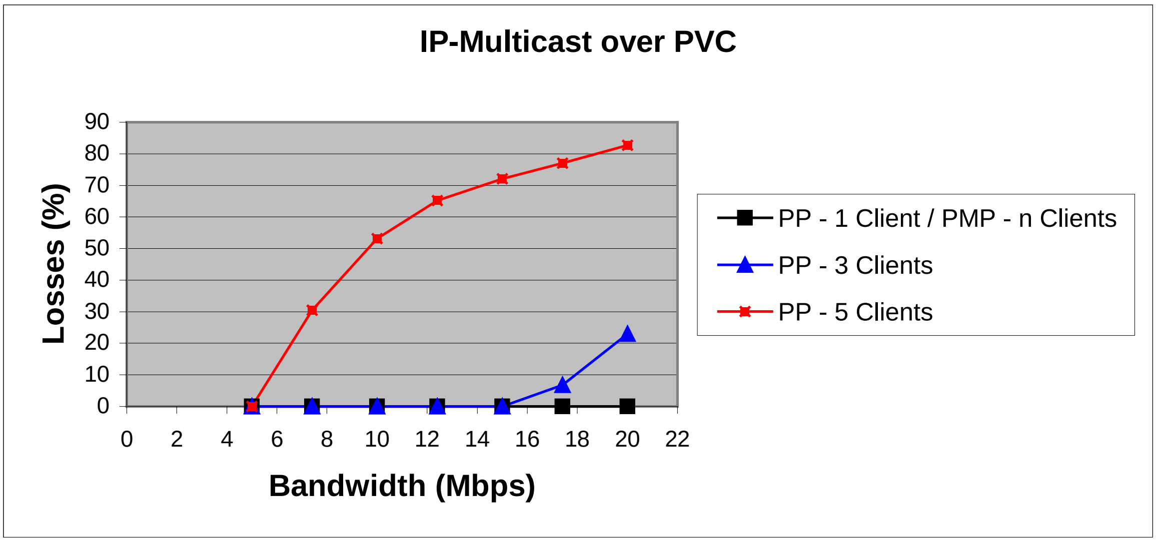

Figure 8: IP-Multicast over PVC Performance Results

Our main conclusion from these results is that even though the *Multicast Gateway* architecture is a fast and easy solution to provide Multicast to ADSL end-users, another architecture with a more efficient packet replication has to be used if a large number of end-users needs to be supported. Multicast over PVC can be such architecture if the PP and the PMP approaches are properly combined.

	Multicast Gateway	**Point-to-Point**	**Point-to-Multipoint**
1 Client	3 Mbps	29 Mbps	29 Mbps
2 Clients	2,4 Mbps	26 Mbps	29 Mbps
5 Clients	<1 Mbps	6,2 Mbps	29 Mbps

Table 1: Bandwidth supported without losses for the different architectures.

4 Accounting Considerations

In section 1 the need for IP Multicast support in an ADSL environment to efficiently support Group Communication based application was discussed. Theoretically IP Unicast could also be used for these applications[7]; in this case, however, the global utilization of the network resources would be much more inefficient, driving to a much lower QoS than the one achievable with IP Multicast. For the global benefit of everybody, thus, the users must be encouraged to use IP Multicast[8].

One way of encouraging users to use IP Multicast could be the pricing. For this purpose, though, a charging and accounting architecture is needed. Accounting for Unicast is rather straightforward, and PPP and RADIUS are widely used at the NSPs for this purpose [13]. Accounting for Multicast, however, is significantly more challenging [14]. In this section we discuss how charging and accounting could be provided to the Multicast architectures that have been described in the previous section.

---

[7] Radio distribution is a good example of a widely popular application in the Internet that has clearly a Multicast nature but currently is used in most cases with IP Unicast.

[8] It must be noted that one individual user decides to receive one specific service through Unicast instead of Multicast does not make a noticeable difference in the network usage efficiency. The QoS experienced by users, thus, is not an element that will encourage individual users to use IP Multicast.

As has been stated in section 3.1, the *Multicast Gateway* forwards the Multicast data to the clients through Unicast. If the *Multicast Gateway* is located at the NSP, then all the communication between the NSP and the clients is done through Unicast, and accounting can therefore be performed exactly in the same way as it is performed for Unicast. Thus, the *Multicast Gateway* architecture is not only an immediate solution from the point of view of implementation effort but also from the accounting point of view.

One possibility for performing accounting in the *Multicast over PVC* architecture is using the *MARS Proxy* [15]. In this approach, the MARS Proxy, located at the NSP, receives all Multicast related data from the proxy clients (the PPP clients) through the same PPP connection used for Unicast traffic. Data membership information is sent from the proxy client to the MARS Proxy via IGMP. The MARS Proxy then issues the corresponding MARS_JOIN and MARS_LEAVE messages to the MARS on behalf of that client. Multicast packets from the proxy client are also sent to the MARS Proxy, which forwards them to the MCS for their distribution. This is shown in Figure 9.

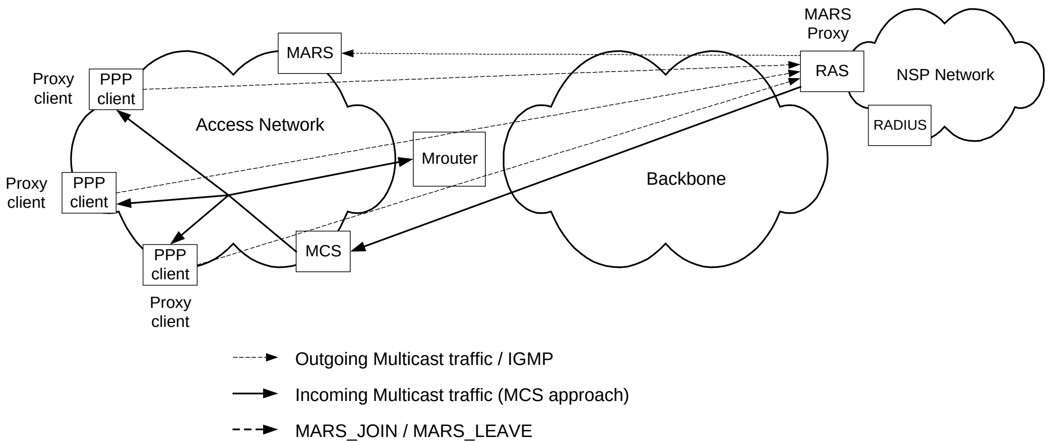

Figure 9 MARS Proxy

With the *MARS Proxy* approach, since outgoing Multicast data and group membership information are sent to the NSP, accounting can be performed at the NSP based on the following parameters: the Multicast traffic sent by the user, the Multicast sessions joined by the user and the duration of these sessions. This solution, however, assumes that a client always has to join a Multicast group in order to receive the data sent to this group. But this requirement is not met in the PMP PVC approach; in this approach, a client that is a leaf node of the PMP PVC will receive a Multicast stream sent to this PVC regardless whether this client has joined the corresponding group or not. This behavior makes it very difficult to perform accounting on the client for the Multicast sessions being distributed through the PMP PVC. This observation leads us to suggest an accounting architecture based on using the PMP option for those services that are only charged to the sender and the PP option for those services for which the receivers are also charged. Note that this solution fits very nicely the '*preassigned bandwidth to content providers*' decision criterion described in section 3.2.1

The accounting approach described above has some obvious limitations. One important limitation is the fact that in the PMP PVC approach the receivers can not be charged. Another limitation is that in the case of the PP option, the receiver can only be charged for the the sessions joined and the duration of these sessions. Other important parameters, such as the traffic being sent to these sessions, can not be accounted for. In order to overcome these limitations, more sophisticated architectures, such as the ones described in [16] or in [17], could be used. These architectures are based on the idea of performing the accounting in the ingress and egress routers. These architectures allow a more detailed evaluation of the resource consumption than the *MARS Proxy*. On the other hand they have two major drawbacks: first they require a much higher implementation effort (ingress and egress routers must be modified), and second they require collaboration between the access network operator and the NSP in the performing of the accounting.

The above discussion shows that a tradeoff has to be made between the level of detail in the accounting and implementation effort. The *MARS Proxy* approach requires a small implementation effort and provides an acceptable level of detail when the '*preassigned bandwidth to content providers*' decision criterion is used. These features make the *MARS Proxy* the solution that fits best the aim of the *Multicast over PVC* architecture of being a short-term solution for first generation ADSL networks.

5 Conclusions

In a high-speed network scenario such as ADSL, point-to-multipoint and multipoint-to-multipoint applications become more and more important. Thus, IP Multicast will be a key point in the success of ADSL networks. We feel that Multicast should be available in ADSL from the very beginning of its deployment.

At this moment on time an architecture that supports Multicast in current ADSL networks is missing. We have implemented two architectures for this purpose: the *Multicast Gateway* and *Multicast over PVC*. Even though both architectures provide the same functionality (i.e., Multicast in ADSL), they cover different needs.

The goal of the *Multicast Gateway* is to provide an easy way of studying the response of the users to a Multicast capable network. The study of this response can be used in order to determine whether it is worth or not investing a bigger effort for a more scalable architecture.

The goal of the *Multicast over PVC* architecture is to provide an efficient solution for IP Multicast in current ADSL networks, supporting a bigger number of users. We expect this architecture to be used in the first stage of ADSL deployment (extended field-trials) , and we expect it to be substituted by MARS/MCS in a further future. For this purpose, the *Multicast* *over PVC* architecture has been designed to enable a smooth migration to MARS/MCS when SVCs become available on a broader basis.

We plan to use both architectures in our ADSL field trials. In the first stage, the *Multicast Gateway* will be used in order to get a first experience with Multicast. In the second stage, the *Multicast over PVC* architecture will be used, with the goal of evaluating its scalability and performance in a real scenario with a large number of users.

Acknowledgments

We gratefully acknowledge Christoph Kuhmünch and Wolfgang Effelsberg from University of Mannheim for giving us their *Multicast Gateway* code. Without their help and collaboration the experiments described in section 3.1 would not have been possible.

References

[1] ADSL Forum Home Page, http://www.adsl.com

[2] ADSL Forum Technical Report TR-012, *Broadband Service Architecture for Access to Legacy Data Networks over ADSL Issue 1*, June 1998.

[3] W. Simpson, *The Point-to-Point Protocol (PPP)*, RFC1661, IETF, July 1994.

[4] G. Gross, M. Kaycee, A. Lin, A. Malis, J. Stephens, *PPP over AAL5,* Work in progress – Internet draft, IETF, July 1998.

[5] RedBack, Cisco, Westell, Alcatel, AG Communication Systems, Bellcore, Ameritech, BellSouth, Bell Atlantic, *Core Network Architectures for ADSL Access Systems*, ADSL Forum 98-017, March 1998.

[6] G. Young (BT), *ATM Architectures for ADSL: Comparison & Migration*, ADSL Forum 98-080, June 1998.

[7] G. Armitage, *Support for Multicast over UNI 3.0/3.1 based ATM Networks*, IETF RFC 2022, November 1996.

[8] T. Braun, S. Gumbrich, H. J. Stüttgen, *Comparison of Concepts for IP Multicast over ATM,* Proceedings of the 2nd IEEE ATM Workshop, San Francisco, August 1996.

[9] A. Banchs, M. Gabrysch, T. Dietz, B. Lange, H. J. Stüttgen, *Supporting Multicast in ADSL networks,* accepted to the IEEE ATM Workshop, Kochi, Japan, May 1998.

[10] C. Kuhmünch, *A Multicast Gateway for Dial-In Lines*, Technical report TR-98-0598, University of Mannheim, July 1998.

[11] M. Handley, *SAP: Session Announcement Protocol*, Work in progress – Internet draft, IETF, May 1997.

[12] R. Talpade, M. Ammar, *Multicast Server Architectures for MARS-based ATM multicasting*, IETF RFC 2149, May 1997.

[13] C. Rigney , *RADIUS Accounting*, RFC2139, IETF, July 1994.

[14] S. Herzog, *Accounting and Access Control for Multicast Distributions: Models and Mechanisms*, PhD dissertation, University of Southern California, USA, 1996.

[15] D. Allan, *MARS Proxy*, Work in progress – Internet draft, IETF, March 1998.

[16] G. Carle, M. Smirnov, T. Zseby, *Charging and Accounting Architecture for IP Multicast Integrated Services*, 4th International Symposium on Interworking, Ottawa, Canada, July 1998.

[17] M. Yamanouchi, N. Ishikawa, O. Takahashi, *RADIUS Extension for Multicast Router Authentication*, Work in progress – Internet draft, IETF, March 1998.