

Supporting Multicast in ADSL Networks

A. Banchs, M. Gabrysch, T. Dietz, B. Lange, H. J. Stüttgen
NEC Europe Ltd, Computer and Communication Research Laboratories Heidelberg
E-mail: adsl@ccrle.nec.de

Abstract: In this paper different approaches to support multicast in ADSL networks are discussed and compared in terms of scalability, performance and implementation effort. Existing approaches, such as the Multicast Gateway or RAS replication, have the advantage of easy implementation, but suffer from undesirable scaling properties and low performance. On the other hand, more scalable solutions, such as MARS/MCS, rely on SVCs, which are currently not available in ADSL-based broadband networks. Thus, so far there is no solution based on currently available PVC operation providing a good performance. In order to cover this need, we propose a new architecture: Multicast over PVC. This architecture, besides providing a good performance, gives the possibility of an easy migration towards future SVC operation.

1 Introduction

ADSL technologies are a new platform for delivering broadband services to homes and small businesses. ADSL can support a wide variety of high bandwidth applications, such as high-speed Internet access, telecommuting, virtual private networking, videoconferencing, audio and video distribution, distributed games and push technologies. These services were either not possible to support or were ineffectively supported by conventional dial-up modem data delivery technologies. Since multicast is essential in many of these high bandwidth applications, multicast support is a key requirement for the success of ADSL networks.

Although this paper is focused on ADSL, all concepts discussed here can also be applied to any other xDSL technology.

2 ADSL Networks

This section provides some background information about ADSL networks. For further information refer 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.

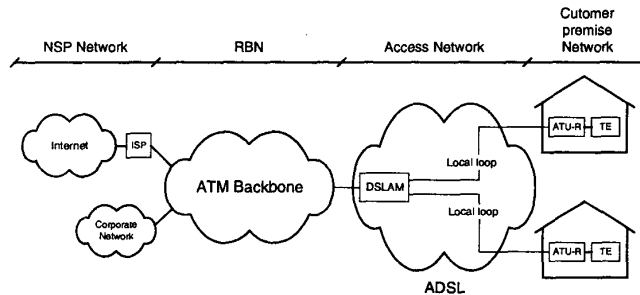


Figure 1 ADSL-based Broadband Network Architecture

2.2 PPP over ATM over ADSL

For the connection between the end-user and the NSP, the ADSL Forum recommends the 'PPP over ATM over ADSL' model [2]. The ATM over ADSL architecture preserves high-speed characteristics and provides QoS guarantee in the ADSL environment without changing protocols, and using PPP over ATM provides, among other features, AAA (Authentication, Authorization and Accounting) and interaction with RADIUS servers. The model used for PPP over ATM is the IETF proposed standard for PPP over AAL5 [3].

Since unicast in ADSL networks will be based on this 'PPP over ATM over ADSL' model, a Multicast Architecture for ADSL will have to be compatible with this approach for unicast.

2.3 SVC Support

In the 'PPP over ATM over ADSL' model described above, there are two possibilities for the ATM connection between the end-user and the NSP: 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 [4]) 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 [5]. 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 Requirements for a Multicast Architecture in ADSL

In this section the desirable features in a Multicast Architecture for ADSL networks are introduced:

- **Scalability:** The following dimensions of scalability can be distinguished:
 - *Scalability in terms of number of groups.* In order to support a large number of groups, the delivery of the packets corresponding to a multicast group has to be done selectively to the hosts that have joined the group.
 - *Scalability in terms of number of hosts.* In order to support a large number of hosts, the replication of the multicast packets has to be done in an efficient way.
 - *Scalability in terms of backbone bandwidth consumption.* The consumption of backbone bandwidth depends on the location where the replication takes place.
- **Avoid broadcasting multicast groups** (related to the scalability in terms of number of groups)
- **Efficiency of replication** (related to the scalability in terms of number of hosts)
- **Efficient use of the backbone bandwidth** (related to the scalability in terms of backbone bandwidth)
- **Support for dynamic joining and leaving**
- **Compatibility with PPP for unicast**
- **AAA support for Multicast**
- **Feasibility of implementation:** Architectures based on SVCs will only be feasible as long-term solutions.
- **Migration to a SVC scenario:** In the case that the architecture is based on PVCs, it should provide a low-impact migration to SVCs.
- **Implementation effort:** Only in long-term SVC architectures or in PVC architectures providing a low-impact migration to SVCs it makes sense to invest a high implementation effort.

4 Proposals for Multicast Support

In this section different proposals for multicast support in ADSL networks are introduced. Architectural strengths and weaknesses of each are discussed with respect to the requirements for a multicast solution introduced in the previous section.

4.1 Multicast Gateway

The *Multicast Gateway* [6] is a client-server application that allows the transmission of multicast sessions to hosts located in non-multicast capable networks. The multicast sessions are transmitted through tunnels that are dynamically established between the applications running at the server and at the client. The *Multicast Gateway* was originally developed for dial-up ISDN clients at the University of Mannheim and has been adapted to ADSL at NEC CCRLE Heidelberg.

This solution, which has proven to be easy to implement, has a major drawback: performance. The fact that the *Multicast Gateway* runs at the end-user space of the server makes the replication of packets highly inefficient (the replication is done at the application level). As a consequence, this approach provides a very low scalability in terms of number of hosts.

The performance of the gateway was evaluated in terms of throughput and losses as a function of the number of clients. The purpose of these tests was not to get the maximum number of clients supported by the gateway architecture (which strongly depends on the hardware) but to understand better how the number of clients impacts the gateway's performance. Figure 2 shows the experimental setup of the measurements.

The results of the measurements are shown in Figure 3: for different numbers of clients, the average relative loss rate is given as a function of the bandwidth of the multicast stream. It can be observed that losses increase considerably with the number of clients. Therefore, the *Multicast Gateway* architecture does not scale with the number of clients. This is a consequence of the low performance due to the inefficient replication at the application level. The conclusion that can be drawn from these results is that, if a large number of end-users needs to be supported, another architecture with a more efficient packet replication has to be used.




Figure 2 Experimental setup




Figure 3 Performance Results

4.2 RAS Replication

As it has been mentioned in section 2.2, the ADSL Forum recommends PPP for the connection to the NSP with unicast. A straight-forward approach in this scenario would be to provide multicast through the same PPP connection used for unicast. Thus, multicast could be easily provided via replication at the Remote Access Server (RAS), with the RAS functioning as a multicast router and performing replication of multicast traffic to each PPP attached destination. This approach has the main advantage that it provides the easiest migration from the existing NSP access infrastructure.

This approach, however, has also a major drawback: since the replication is done at the RAS, the backbone is flooded by as many streams of each multicast session as hosts have joined that session. Thus, this solution makes a very inefficient use of the backbone bandwidth and is therefore not scalable in terms of backbone bandwidth consumption. Figure 4a illustrates this undesirable behavior of *RAS Replication*, whereas Figure 4b shows the desirable behavior.

The *RAS Replication* solution, even though it provides a more efficient replication than the *Multicast Gateway* (RAS replication is done at the network level instead of the application level), does not use the inherent multicast facilities provided by ATM (i.e., PMP VC). Thus, another drawback of this approach is performance: replication is not very efficient, and, therefore, the scalability in terms of number of users will be low.




Figure 4 Undesirable behavior of RAS Replication

4.3 MARS/MCS

MARS/MCS¹ [7] is the proposal by the IETF for multicast in ATM networks. MARS/MCS is based on a Multicast Address Resolution Server (MARS) for address resolution, and for data forwarding two different approaches are supported:

- SVC meshes based approach
- Server based approach (MCS)

¹ 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.

In both approaches, the MARS keeps a cache of {layer 3 multicast address, ATM address 1, ATM address 2, ..., ATM address N} mappings. A sender that wants to send a multicast packet to a certain group sends a MARS_REQUEST message to the MARS in order to get the corresponding set of ATM addresses for the specified IP multicast address. The MARS responds with one or more MARS_MULTI messages containing the complete set of addresses {ATM address 1, ATM address 2, ..., ATM address N} for the requested group. In the *SVC meshes* approach this group of addresses corresponds to the hosts of the cluster that have joined the group, while in the *server based* approach it corresponds to the multicast servers that are serving this group. When receiving the MARS_MULTI, the sender stores the mappings in its local cache and establishes a PMP SVC to all hosts specified in the address list.

For joining and leaving a multicast group, a host has to send MARS_JOIN and MARS_LEAVE messages to the MARS indicating its ATM address and the multicast addresses of the groups to join or leave. The MARS forwards this information about joining and leaving hosts over a control VC, so senders can update their mappings in the local cache.

In the *SVC meshes* approach (Figure 5a), a multicast sender establishes direct PMP SVCs to the hosts of the cluster that have joined the group. Multicast packets are then sent directly from the sender to all receivers. Depending on group membership information distributed by the MARS protocol each sender maintains its own PMP SVCs.

In the *server based* approach (Figure 5b), multicast packets are sent to one or more multicast server (MCS) nodes. A sender establishes a PMP SVC to the specified MCSs and sends the multicast packets to the MCSs via this PMP SVC. Then, the MCSs distribute the multicast packets to the receivers inside the cluster.




Figure 5 MARS SVC Mesh (a) and MARS MCS (b)

One of the desirable features mentioned in section 3 is AAA support. The *MARS Proxy* [9] approach provides such support in a MARS/MCS scenario.

The key idea of this approach is that the MARS Proxy receives all multicast related data from the proxy clients (the PPP clients) through the same PPP connection used for unicast traffic. Group 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 either forwards them to the MCS for their distribution or directly distributes them via a SVC mesh.

With this approach, authentication, authorization and accounting can be performed by the MARS Proxy. Figure 6 shows the scenario of the MARS Proxy: the RAS server functions as a MARS Proxy, and the PPP clients as proxy clients.




Figure 6 MARS Proxy

The major drawback of the MARS/MCS approach is that it is based on SVC operation, and as a consequence, cannot be implemented as a short-term solution. Nevertheless, the MARS/MCS approach provides a very good behavior with respect to the requirements given in section 3, and is therefore a very good candidate for a long-term solution.

4.4 Multicast over PVC

Of all the approaches described so far, only MARS/MCS fulfills the requirements specified in section 3 in a reasonable way. However MARS/MCS relies on SVCs, which are currently not available. Hence, a short-term solution fulfilling reasonably the requirements given in section 3 is missing. In order to cover this need, we have developed a new proposal: *Multicast over PVC*.

The *Multicast over PVC* architecture is based on MARS/MCS [7]. For address resolution a MARS server (see section 4.3) is used. For delivering multicast data a MCS server (see also section 4.3) is always used.

For joining and leaving a multicast group, a host issues a MARS_JOIN or MARS_LEAVE message, respectively, to the MARS (using a preconfigured PVC between the MARS and the host). MARS uses this information to keep track of all the group membership information in the cluster.

For transmitting a multicast packet, a host 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 through a preconfigured PVC. The MCS then obtains the membership information from the MARS using 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.

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 (point-to-multipoint) PVC from the MCS to all the hosts of the cluster
- Multiple PP (point-to-point) PVCs from the MCS to each host

The PMP PVC approach (Figure 7a) broadcasts the multicast data to all hosts, 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 the packet 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 accesses might get flooded with useless multicast traffic.

The multiple PP PVCs approach (Figure 7b) 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 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.




Figure 7 PMP PVC to all hosts (a) and multiple PP PVCs (b)

The decision whether it is more appropriate to replicate and forward the packets of a multicast group using the broadcast PMP PVC or the multiple PP PVCs approach is taken dynamically by the MCS. This decision can be based on different criteria:

- Number of hosts that have joined the group.
- Avoid flooding user accesses.
- Preassigning parts of the user access bandwidth to content providers.
- A combination of the above criteria.

In order to provide AAA facilities to the *Multicast over PVC* architecture, a similar approach to the *MARS Proxy* based on PVCs instead of SVCs can be used.

5 Comparison of the Proposals

Table 1 summarizes the features of the proposals discussed in the previous section, with respect to the requirements given in section 3. The Multicast Gateway can be located either at the *NSP* or in the *boundary* between the access network and the backbone. Since its behavior is quite different in each case, different entries have been used in Table 1. Different entries (*PMP* and *PP*) have also been used for each of the two possible approaches of the *Multicast over PVC* architecture.

	Multicast Gateway (boundary)	Multicast Gateway (NSP)	RAS Replication	Multicast over PVC (PMP)	Multicast over PVC (PP)	MARS/ MCS
Broadcasts Multicast groups	NO		NO	YES	NO	NO
Replication efficiency	VERY LOW		LOW	HIGH	LOW	HIGH
Efficient use of the backbone	YES	NO	NO	YES		YES
Dynamic joining and leaving	YES		YES	YES		YES
Compatible with PPP	YES		YES	YES		YES
AAA	NO	YES	YES	YES		YES
Requires SVC support	NO		NO	NO		YES
Low-impact migration to SVC	NO		NO	YES		-
Implementation effort	LOW		LOW	HIGH		HIGH

Table 1 Comparison of the Multicast Proposals

6 Conclusions

In this paper different approaches for providing Multicast in ADSL networks have been proposed discussing strengths and weaknesses of each.

For immediate provisioning of multicast services to a small number of ADSL end users the *Multicast Gateway* architecture or the *RAS Replication* approach can be used, since they are easy to implement.

On a large scale deployment scenario these solutions cannot be used for scalability and performance reasons. If many end-users have access to ADSL services and the ADSL-based broadband network supports SVCs the *MARS/MCS* approach for multicast proposed by the IETF will have the best performance. However, current ADSL-based broadband networks do not support SVCs.

In order to provide an efficient and scalable solution for multicast in the PVC based access networks of the near future a new architecture, *Multicast over PVC*, has been designed. This architecture provides the best tradeoff between feasibility of implementation and performance, and gives the possibility of an easy migration towards future SVC-based solutions.

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 4.1 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] G. Gross, M. Kaycee, A. Lin, A. Malis, J. Stephens, *PPP over AAL5*, Internet draft, IETF, July 1998.
- [4] 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.
- [5] Gavin Young (BT), *ATM Architectures for ADSL: Comparison & Migration*, ADSL Forum 98-080, June 1998.
- [6] C. Kuhmünch, *A Multicast Gateway for Dial-In Lines*, Technical Report TR-98-98, University of Mannheim, July 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] D. Allan, *MARS Proxy*, Internet draft, IETF, March 1998.