Comparison of Concepts for IP Multicast over ATM

Torsten Braun, Stefan Gumbrich, and Heinrich J. Stüttgen

IBM European Networking Center, Vangerowstr. 18, D-69115 Heidelberg
E-mail: braun@heidelbg.ibm.com, Phone: +49 6221/59-4352, Fax: +49 6221 3300

1 Introduction

There are several possible approaches to support multicast communication within an ATM subnetwork. A simple one is ATM Forum's LAN emulation, where a special server broadcasts all multicast packets to all nodes within a subnetwork. This paper compares two approaches published as IETF draft documents. These are the so-called multicast router approach where an IP multicast router is used as ATM multicast server and the Multicast Address Resolution Server (MARS) concept. MARS is currently being specified within IETF IP over ATM working group.

2 IP Multicast

IP multicast is based on IGMP (Internet group management protocol) [3][4] and DVMRP (distance vector multicast routing protocol) [10]. IGMP allows IP hosts to join and leave multicast groups. DVMRP is used as multicast routing protocol in multicast router systems.

The *IP multicast address*es of IP multicast messages are mapped to *MAC multicast addresses* if they are supported by the MAC protocol and the network adapters. Otherwise, multicast packets are sent to the MAC layer broadcast address.

IGMP is mainly designed to run over shared-medium LANs such as Ethernet, Token Ring or FDDI. A router sends periodically IGMP Host Membership Query messages over its interfaces to the ***all hosts IP multicast address***. End systems respond with an IGMP Host Membership Report Message in order to join or to stay member of a group. The IGMP Host Membership Report contains the corresponding group addresses. End systems listen to IGMP Host Membership Reports of other local end systems. If an end system detects that another local end system responds with an IGMP Host Membership Report for the target group, it doesn't send a second IGMP Host Membership Report. In shared medium LANs it is sufficient that the router knows the existence of any member of a certain goup on one of its subnetworks.

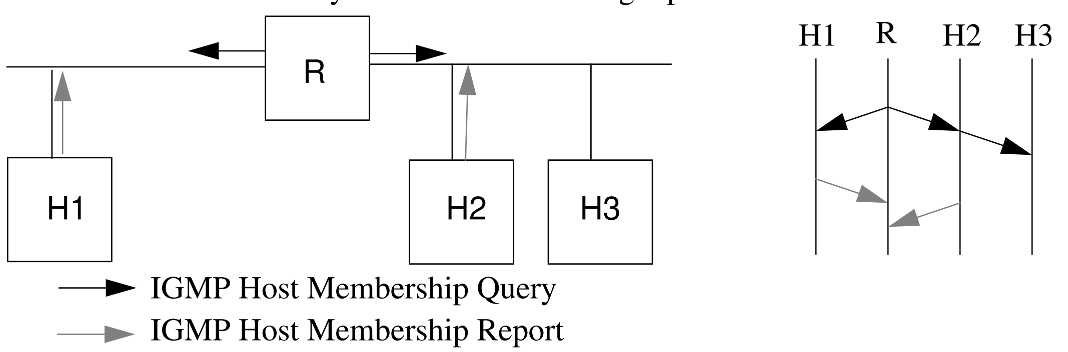

Figure 1: IGMP

3 IP-Multicast over ATM

To provide an IP multicast service over ATM there are several possible approaches. In addition to LAN emulation developed by the ATM Forum, there is ongoing work within the IETF working groups. The different approaches, in particular the IETF approaches, are described and compared with each other in the following sections.

3.1 LAN Emulation

IP multicast over ATM may be supported by ATM Forum's LAN emulation, which emulates the functionality of shared medium LANs such as Ethernet and Token Ring networks. The LAN emulation concept also emulates the broadcast facility of shared medium LANs. A dedicated server (broadcast and unknown server, BUS) broadcasts multicast packets to all ATM nodes within the local ATM subnetwork [9]. The maximum packet size is limited to the maximum packet sizes of the emulated Ethernet or Token Rings LANs. These packet sizes are much smaller than AAL5 packet sizes. The major drawback is that all multicast packets are sent to all nodes of a subnet and the nodes receive multicast packets for groups they do not belong to. Therefore, they have to filter multicast packets depending on their group membership.

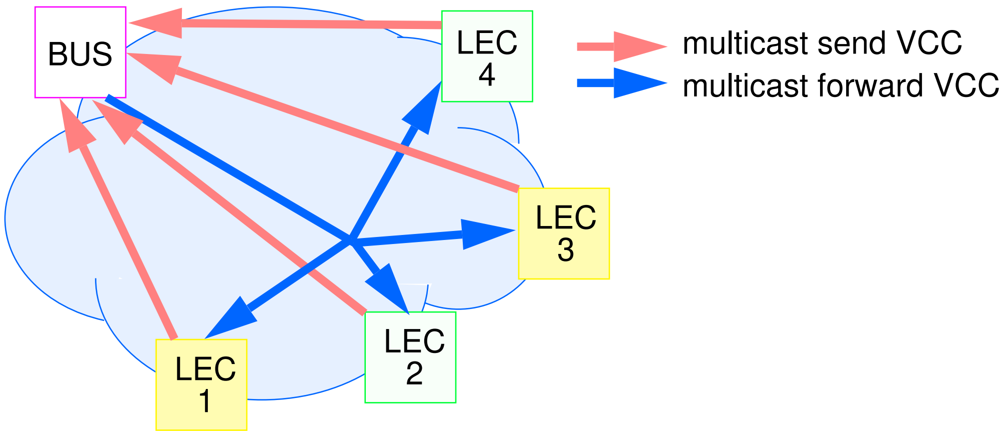

Figure 2: LAN Emulation

3.2 Multicast Router

Another approach is the extension of an IP multicast router by multicast server functionality for an ATM subnetwork. The server/router provides multipoint VCs for each multicast group of the local ATM subnetwork.

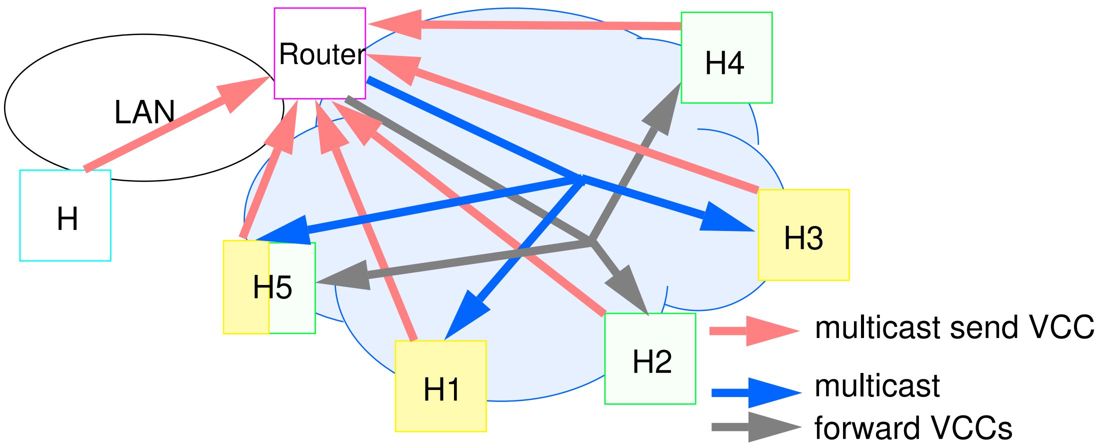

Figure 3: IP multicast router as multicast server

Multicast packets received from other subnetworks or from local senders are distributed over the multipoint VC associated with the particular IP multicast group.

Management of multipoint-VCs is based on IGMP message analysis. End systems send IGMP messages via a point-to-point VC to the multicast router in order to join a group. Each end system willing to join a group has to send IGMP Host Membership Reports, since the multicast router does not forward IGMP messages. Due to this, the router knows explicitely the IP address of each local member of any group. The ATM address of an end system has to be resolved by an ATM ARP server [5]. Based on this information, the router establishes and maintains one multipoint VC per group. Multicast packets received on point-to-point VCs from a sender are forwarded via the multipoint VC of the multicast group.

3.3 Multicast Address Resolution Server (MARS)

While the approach described above is limited to IP, the multicast address resolution server (MARS) concept is theoretically independent of the supported network (layer 3) protocol [1]. It is basically an extension of the ATM ARP server defined in RFC1577 [5]. However, new message formats and a new protocol are defined for multicast address resolution.

A 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 ATM addresses {ATM address 1, ATM address 2, ... , ATM address N} for the requested group. The sender stores the mappings in its local cache and establishes a multipoint-VC to all ATM end systems specified in the address list.

For joining and leaving a multicast group, an end system 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 to leave. The MARS forwards information about joining and leaving end systems over a control VC. Senders analyse the information and update their mappings in the local cache.

The MARS concept supports both

- an approach based on VC meshes and
- a server based approach.

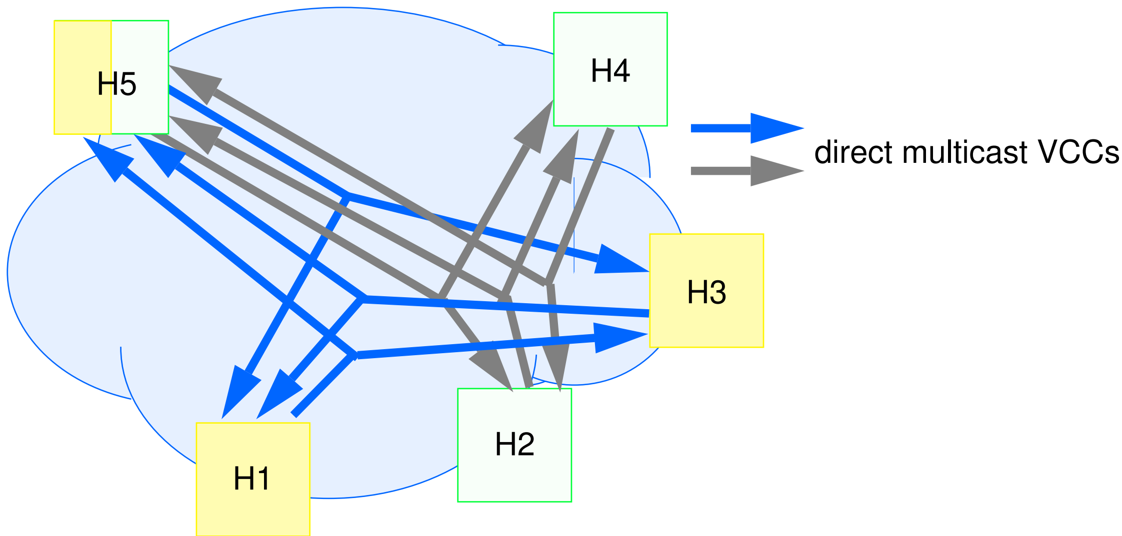

Figure 4: MARS VC Mesh

In the ***VC mesh*** approach (Figure 4), a multicast sender establishes direct multipoint VCs to the local ATM end systems. Multicast packets within the ATM subnetwork are sent directly from the sender to all receivers. Depending on group membership information distributed by the MARS protocol each sender maintains its own multipoint VCs.

In the ***server based approach*** (Figure 5), multicast packets are sent to one or more multicast server (MCS) nodes. In contrast to a list of ATM end systems addresses the MARS_MULTI message of the MARS contains a list of multicast server addresses. A sender establishes a multipoint VC to the specified MCSs and sends the multicast packets to the MCSs via this multipoint VC. Then, the MCSs distribute the multicast packets to the receivers of the local subnetwork. Several issues of the current specification are still open such as the existence of more than one multicast server for a certain multicast group and the synchronization between several MCSs.

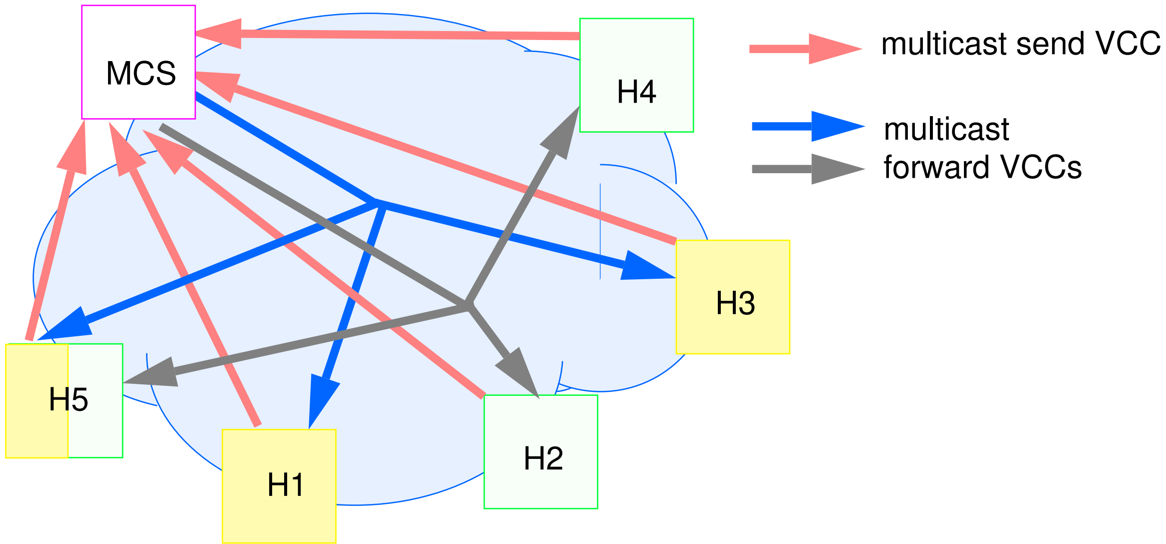

Figure 5: MARS MCS

4 Comparison

In the following, we compare the MARS and the multicast router approach considering the complexity, performance, ressource requirements and specific problems.

4.1 Complexity

The ***MARS*** approach introduces a new protocol with several new packet formats. The protocol is rather complex, in particular to support multicast servers.

In contrast, the ***multicast router*** approach requires relatively small changes of the multicast router software in order to receive and forward multicast packets. No additional protocol is required to implement the concept. The ATM end systems only have to know the ATM ARP server address and the multicast router address. An optimization is to perform address resolution for multicast packets using a RFC1577 ARP server. The ARP server returns the ATM address of the multicast router when it receives an ARP request for a multicast group. In that case there is no difference for the end system between unicast and multicast address resolution.

4.2 Performance

A disadvantage of the multicast router approach compared to the MARS VC mesh approach is that the multicast traffic is concentrated at the router. Local multicast traffic is delayed due to the additional hop. The same problem occurs with MARS MCSs.

If the sender is located on an other subnetwork than the local receivers, multicast traffic passes always the multicast router. With the multicast router approach and the MARS VC mesh concept, the multicast router forwards the data directly to the receivers. The amount of hops is equal in both cases. The MARS MCS approach requires an additional hop. The multicast router sends a multicast packet to the MCS, which finally sends the packets to the receivers.

4.3 Ressources

The main problem of the MARS VC mesh concept is the extensive usage of VCs [5]. For an ATM network with n senders and m multicast groups $n \cdot m$ multipoint VCs must be established and maintained to realize a MARS VC mesh, while only n + m multipoint VCs are necessary for a server based approach (either multicast router or MARS MCS).

The signalling overhead of the ATM signalling protocol and the MARS protocol grows with the number of VCs. A joining or leaving group member affects the modification n multipoint VCs for MARS VC meshes while only 1 multipoint VC must be updated for multicast server based approaches.

4.4 Specific Problems

A problem of the MARS MCS approach are possible routing loops, when routers receive a forwarded packet from a MCS. Furthermore, group management is performed twice, namely by IGMP at IP level and by MARS at ATM level. However, there are possible optimizations to reduce IGMP traffic using MARS group management information for IGMP.

5 Conclusions

The multicast router approach performs better than the MARS MCS concept and has only one disadvantage compared to MARS VC meshes. The throughput and delay for local ATM multicast traffic is expected to be worse than for MARS VC meshes. In scenarios with sender and receiver located in different subnetworks, the multicast router has no performance drawbacks. It has significant advantages when parameters like resource utilization and signalling overhead are investigated.

All approaches discussed above support best-effort traffic only and do not provide any quality-of-service (QoS) guarantees. Therefore, it seems reasonable to apply a server based approach for traffic without throughput and delay constraints.

Traffic with certain QoS constraints should be supported by resource reservation protocols such as RSVP. RSVP signalling can also be used to establish short-cuts between ATM end systems connected to different IP subnetworks but connected via several ATM switches [2]. The establishment of short-cuts is currently under study for the MARS approach and the multicast router concept.

6 References

[1] G. Armitage: Support for Multicast over UNI 3.1 based ATM Networks, Internet Draft, February 1996

[2] A. Birman, R. Guerin, D. Kandlur: Support for RSVP-based Service over an ATM Network, Internet Draft, February 1996

[3] S. Deering: Host Extensions for IP Multicasting, RFC 1112, August 1989.

[4] W. Fenner: Internet Group Management Protocol, Version 2, Internet Draft, August 1995.

[5] M. Laubach: Classical IP and ARP over ATM, RFC 1577, January 1994

[6] W. Milliken: Integrated Services IP Multicasting over ATM, Internet Draft, July 1995

[7] W. Milliken: IP Multicasting over ATM: System Architecture Issues, Internet Draft, July 1995

[8] R.R. Talpade, M.H. Ammar: Multicast Server Architectures for MARS-based ATM multicasting, Internet Draft, February 1996

[9] H.L. Truong, W.W. Ellington, J.Y. LeBoudec, A.X. Meier, J.W. Pace: LAN Emulation on an ATM Network, IEEE Communications Magazine, Vol. 33, No. 5, May 1995, pp 70-85

[10] D. Waitzman, C. Partridge, S. Deering: Distance Vector Multicast Routing Protocol, RFC 1075, November 1988.

Biographies

Torsten Braun received his diploma degree and his doctoral degree in computer science from the University of Karslruhe, Germany, in 1990 and 1993, respectively. From 1990 to 1994 he was a research assistant at the Institute of Telematics, University of Karlsruhe (Germany). From 1994 to 1995 he was a visiting scientist at the Institut National de Recherche en Informatique et en Automatique (INRIA) in Sophia-Antipolis (France). Since 1995, he has been a guest scientist at IBM European Networking Center in Heidelberg (Germany). His main research interests are the design and the implementation of high performance transport systems for multimedia communication. Dr. Braun is a member of the ACM and IEEE.

Stefan Gumbrich received his Diploma in Electrical and Computer Engineering from the Fachhochschule Bingen (Germany) and the University of Tennessee, Knoxville in 1991. In 1994 he received his Diplom in Computer Science from the University of Paderborn (Germany). Since 1994 he has been a guest scientist at IBM European Networking Center in Heidelberg. His main research interests are the aspects of video transmission over high speed networks and the design and implementation of high speed network protocols.

Heiner Stüttgen studied Computer Science at the University of Dortmund (Germany) and SUNY Buffalo, from where he received Diploma and M.Sc. degrees in computer science in 1979. In 1984 he obtained a Doctor of Science degree from the University of Dortmund with a thesis on hierarchical associative memories. In 1985 he joined the IBM development lab at Böblingen (Germany) to work on a mainframe UNIX development project. In 1987 he moved on to the IBM European Networking Center at Heidelberg, where he has been active in different research, development and marketing projects in the areas of heterogeneous networks, protocols and implementation techniques for high-speed networks and multimedia communications. Currently he manages the Broadband Multimedia Department focusing on multmedia communication and ATM networks.