AMnet: Active Multicasting Network

Ralph Wittmann, Martina Zitterbart
Institute of Operating Systems and Computer Networks
Technical University of Braunschweig
38106 Braunschweig, Germany
{wittmann|zit}@ibr.cs.tu-bs.de
http://www.ibr.cs.tu-bs.de/~wittmann/AMnet.html

Abstract

Today we are faced with an increasing variety of networks and end systems. The resulting heterogeneous environment imposes new challenges on communication support for multimedia and collaborative applications. AMnet is an approach that provides multipoint communication support for large-scale groups with heterogeneous receivers. Active network nodes are used to deliver data streams with user-tailored QoS. They provide so-called QoS filters that remove information from continuous media streams in order to reduce data rate for low-end receivers without affecting high-end receivers. A prototype implementation of AMnet based on RSVP Filer is presented.

1 Introduction

In the past decades a tremendous growth in the use of computer networks in general and the Internet in particular for communication, computation, and collaboration can be noticed. Especially, an increasing demand for multimedia and collaborative applications, like teleworking and teleconferencing systems can be observed. For many of these applications group communication is a natural paradigm. Therefore, proper support within the communication subsystem is required. This includes flexible and scalable multicast support, enhanced group management, and provision of integrated services by the network.

Multicast support is inherently needed in order to facilitate development and to increase communication efficiency. This demand is acknowledged by various recent approaches that address transport protocols for reliable multicast. The focus of most current proposals is on the support of reliability. Mostly, all group members experience the same level of service, independent of their network attachment and end system equipment.

However, heterogeneity is growing with respect to networks and end nodes. Consider a scenario where some receivers of the group use simple PDAs (Personal Digital Assistants), which typically are poorly equipped and often attached via wireless links. Others may use high performance workstations. Due to limited processing power or network capabilities low-end receivers may be unable to handle the same data stream as receivers using high-end workstations. Moreover, these devices may have small displays offering low resolution and color depth. Hence, different participants in a multi-point communication may have different requirements with respect to quality of service (QoS). The goal of the proposed research work is in providing user-tailored QoS support for individual group members. Different data streams will be delivered, e.g., to the high performance workstation and to the wireless attached PDA. In most current approaches the service provided to individual group members is penetrated by the group member with the lowest service capabilities. Such an approach is not acceptable for multimedia and collaborative applications in heterogeneous networking environments.

AMnet especially addresses the aspect of heterogenous group communication. It is based on Active Networking [13] in the sense that it uses active network nodes in order to individually tailor data streams to the end users service requirements. Therefore, in a first step so-called QoS-filters are implanted in network nodes. They adapt continuous media streams (e.g., MPEG-video) to specific QoS requirements.

The paper is structured as follows. Section 2 presents the AMnet approach and gives an overview of related work. A prototype implementation is presented in Section 3. It is based on RSVP as signalling protocol for QoS filters. Section 4 concludes the paper and gives an outlook on future work.

2 Heterogeneous Multicasting

A key aspect of AMnet can be seen in the usage of active networking in order to provide user-tailored services for heterogeneous groups.

2.1 The AMnet Approach

In Figure 1 a multicast tree is depicted that provides heterogeneous multicasting capabilities through active network nodes, called *Active Multicasting Nodes*. Some of the network nodes included in the dissemination tree are active, others are passive with respect to enhanced multicasting tasks. Data streams with different QoS are delivered to the end users of the group, where QoS B is lower than QoS A and QoS C is lower than QoS B. Therefore, the active network nodes comprise service modules for QoS filtering and enhanced QoS signalling. Furthermore, error control capabilities are implanted in the active network nodes.

The service modules can be selected and configured according to requirements of individual group members. Besides filtering and error control, ordering schemes, such as partial and causal ordered delivery, are other typical examples for selectable modules. The focus of this paper is on QoS filters and the corresponding filter signalling.

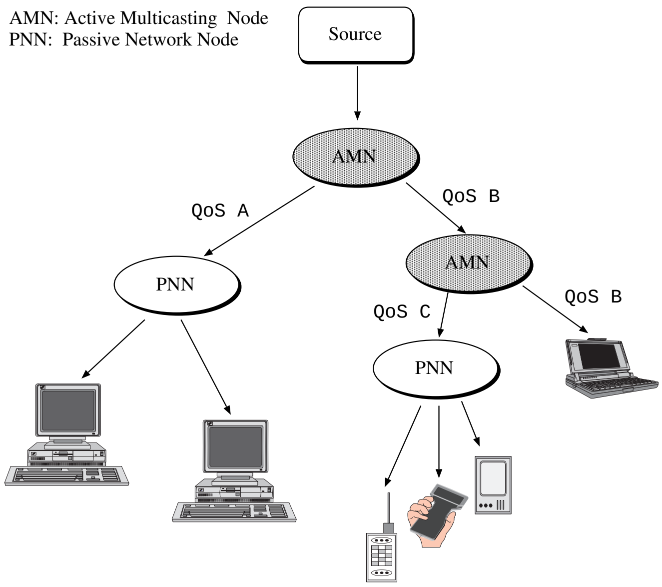

Figure 1: Multicast tree with Active Multicasting Nodes

In order to configure and activate the service modules inside the active network nodes, a signalling protocol is needed. With respect to signalling of QoS filters, the following requirements have to be met by the signalling protocol. A QoS filter should be placed as close as possible to the source of the multicast tree. All nodes above the selected node should have experienced higher QoS filter requests. Furthermore, dynamic join and leave operations must be supported. Therefore, dynamic allocation and re-configuration of QoS filters needs to be supported.

Active network nodes may also perform error recovery. If – in case of a NACK – the requested data are still cached inside the respective network node, it retransmits the data to the receiver. If not, the NACK is passed to active network nodes at higher hierarchy levels. This reliefs in many cases the burden of retransmissions from the source.

Typically, retransmissions are not used in video conferencing applications. However, they can still be advantageous for such applications [4]. Due to severe time constrains retransmission of continuous media from the source to the receiver is often impossible, because of the latency incurred. Since active network nodes can provide retransmissions, the experienced retransmission delay can be significantly reduced compared to the end-to-end delay.

Filtering of continuous data streams is closely related to synchronization. Many multimedia and collaborative applications need support for inter-stream synchronization. For example, a video-conferencing tool which synchronizes audio and video streams of a speaker. Support through the transport system is particularly needed if QoS filtering is used. Like error control, synchronization support has to be aware of QoS filters. This preserves the possibility to synchronize different streams in the end-systems. Moreover, the same argument holds for inter-stream synchronization.

In summary, AMnet is based on active network nodes, that provide user-tailored services for heterogeneous group communication. Particularly, QoS filtering, enhanced signalling, error control, and synchronization are supported by active nodes inside the network.

2.2 Related Work

AMnet addresses various aspects that have been subject of current research projects. This holds particularly for scalable reliable multicasting and QoS filtering.

Several QoS filters have been investigated in [17]. However, signalling of filter functions has not been addressed. Filters for hierarchical encoded streams are presented in [16]. They adapt continuous media streams to a specified QoS. Hierarchically encoded streams are also used in the RLM approach [10]. It is designed for networks like todays Internet, i.e., networks with best effort multi-point packet delivery. Each layer is associated with a multicast group. Receivers can join and leave a group in order to change the experienced QoS. This approach is mainly targeted towards network congestion. The usage of filters to control network congestion is presented in [15]. Heterogeneity is not explicitly addressed. Application level gateways, that operate on top of RTP [12] are addressed in [1] for filtering and conversion operations. Reliable multicast is not addressed in this work.

The problem of providing multicast communication to a large group of receivers is a vital research area in the networking community. Especially, approaches, such as LBRM [6] and LGC [5] are similar to the concept followed by AMnet. However, these protocols do not support heterogeneous receivers. Moreover, QoS filtering of data streams may violate data integrity. Therefore, error control and recovery and mechanisms for congestion control has to be aware of QoS filters.

An overview of the state-of-the-art in active networking can be found in [14]. Bhattacharjee et. al. [2] use active networking technologies for congestion avoidance. Aspects, such as multicasting or advanced QoS filtering are not addressed.

3 The RSVP-based AMnet Prototype

A prototype of AMnet has been implemented. In the context of the Internet, signalling protocols are currently investigated with the goal of providing an integrated services architecture [3]. RSVP [19] seems to be the main candidate for deployment in the Internet. Since its design is flexible with respect to future extensions, the prototype AMnet implementation is based on RSVP.

3.1 QoS filter

The focus was on compressed video streams that are MPEG-1 [7] coded. The implementation and integration of an H.261 video is currently under way. Other filters can be found in [18]. The following types of MPEG-1 filter are implemented [9]: frame dropping filter, re-quantization filter, monochrome filter, and slicing filter.

The *frame dropping* filter simply drops P and B frames of the MPEG-1 stream to reduce the frame rate for less powerful receivers.

The *re-quantization filter* operates on DCT-coefficients and, thus, requires semi-decompression of the video stream. With re-quantization many near-zero coefficients may become zero which leads to a better compression ratio in the subsequent entropy encoding step. With moderate quantization steps a good trade-off between bandwidth reduction and loss of play-out quality can be achieved. However, large quantization steps can result in strange artefacts [11].

The *monochrome filter* removes color information. As MPEG encodes chrominance and luminance information independently, discarding of luminance components can be efficiently performed.

Slicing filters exploit a property of MPEG encoded streams. The macro blocks of MPEG can be structured into *slices*. A slice contains information to synchronize the codec. If a slice gets lost, the codec resynchronizes on the following slice and may continue decoding. Since slices are not always included in MPEG coded data, the filter have to incorporate them. This may lead to a slightly higher data volume. However, slicing filter can drop slices easily and, thus, reduce the bandwidth consumed on the output link.

These types of QoS filters are configured in the active network nodes. In addition, a set of parameters is defined which specifies the requested filter function. The parameters are *quality factor*, *skip-frames*, *B/W-mode* and *number of slices*. The quality factor corresponds to a quantization level. The filter supports the quality factors 1 to 5, with 1 representing the best play-out quality, i.e., the least bandwidth reduction. The parameter skip-frames determines the number of frames to be dropped. The parameter B/W-mode activates the color reduction. Finally, the slice parameter specifies the number of slices per frame that the filter has to insert into the MPEG stream.

3.2 Signalling QoS Filter

Since QoS filters are located in the network, signalling is needed for configuration and allocation of QoS filters and for negotiation of filter parameters. As QoS filter reside on intermediate systems, they can be viewed as network resources. RSVP is a signalling protocol for network resources which is intended for the Internet. Originally, RSVP is not desinged to signal QoS filter functions. However, due to its flexible concept, RSVP can be easily extended to configure and control QoS filters. In the current AMnet prototype RSVP serves as signalling protocol.

RSVP provides the concept of classes that can be defined in order to extend RSVP to new resources. The format of the new RSVP class *QoS filter* (QF) is depicted in Figure 2(a). To assure compatibility with RSVP entities which are not extended to handle QF objects, the QF class carries the class number 208. With this class number a QF object is simply forwarded by an RSVP entity if the QF class is unknown. To select a certain type of filter class the type field is used. As an example, Figure 2(b) depicts a QF object for the MPEG filter.

0 … 31

Length	Class-Num: 208	C-Type
Filterparameter		

(a) Generic QF Class.

0 … 31

Length: 12	Class-Num: 208	C-Type: MPEG
Picture-Quality: 3	Skip Frames: 0	
Slices per Frame: 2	Color/B\W-Mode: 0	

(b) Example: MPEG Filter Class.

Figure 2: RSVP filter class.

3.3 Basic implementation architecture

The basic implementation architecture is depicted in Figure 3. It consists of the following main components:

- RSVP Daemon
- QF Daemon
- QoS Filter (QF)

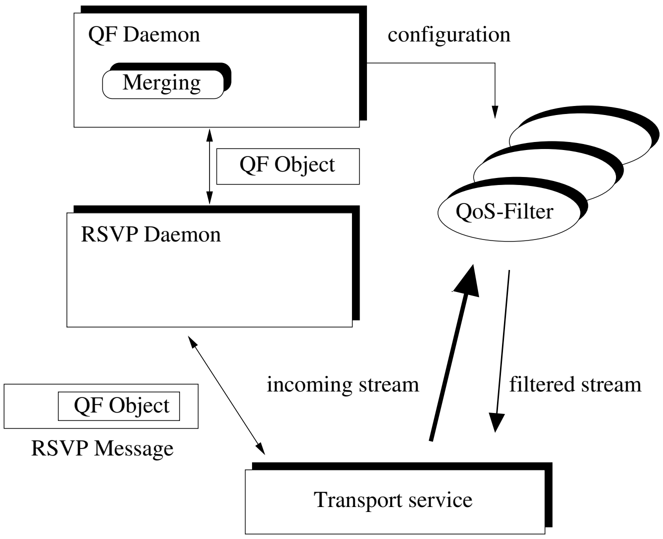

Figure 3: Architecture of an active node

The RSVP daemon has been extended with an interface to the QF daemon to exchange QF objects [8]. In case a QF object is included in the received RESV message, the QF object is extracted and forwarded to the QF daemon.

The QF daemon is responsible for allocation and configuration of the QoS filter according to the requirements stated in the QF object. If it receives a QF object from the RSVP daemon, it needs to check whether other QoS requirements for that group have been received previously. Therefore, the QF daemon keeps some state information concerning filter requirements of group members. The following cases can be distinguished:

- no previous request
- previous request for same QoS filter type
- previous request for other QoS filter type

If no previous request is available, the RESV message is forwarded to the next node in the multicast tree.

If another filtering request for the same group and the same QoS filter is available, merging of the filter parameters is needed. The QF daemon updates the QF object accordingly and forwards it to the RSVP daemon for further propagation towards the sender using RESV messages and standard RSVP procedures.

If a previous request for a different QoS filter type is registered at the QF daemon, two cases need to be distinguished.

First, the new request can be merged with the current request and the resulting request is not higher than the existing request. In this case, the corresponding filter functionality is activated by the QF daemon. The QF object that is periodically retransmitted with RSVP messages needs not to be updated. The node requires still the same amount of input data from the hierarchically higher nodes in the multicast tree.

Secondly, the new requested QoS filter can not be merged with the active filter. In this case, different filters serve different group members in the same active network node. To receive a sufficient data stream the highest requirements for both requests have to be identified. They form the new QF object to be forwarded in the RSVP messages.

A joining group member with new requirements can also lead to the re-location of already installed QoS filters. If the data is already filtered to a lower rate as the requirements of the joining group member, the filter may need to be moved to an active network node that is closer to the group members.

RSVP takes a soft state approach to manage the reservation state in intermediate systems and hosts. A soft state is created and periodically refreshed by PATH and RESV messages. QF objects are included in these messages.

QoS filters are released if group members leave the group and filtering at that node is no longer required. The same time out mechanism used by RSVP is applied for that purpose. Thus, AMnet currently only provides soft-states.

RSVP daemon and QF daemon are involved in the control flow only. User data are directed directly to the corresponding QoS filter. The output of the QoS filter is then forwarded to the packet scheduler of the output link.

3.3.1 Example

In order to illustrate configuration and merging of filters, the example depicted in Figure 4 is considered. There are three receivers: A, B and F. They participate in a multicasted MPEG-1 video stream. Receiver A has severe bandwidth constraints due to a wireless link and, thus, requests a very low quality factor of 5. Receiver A issues a RESV message that includes the corresponding QF object. No other RESV messages have passed nodes C and D. Thus, the RESV message is forwarded to Node E. Node E has already seen a RESV message without bandwidth constraints from receiver F. As a result, node E has to install a new QoS filter that is configured to the quality factor of 5 (cf., Figure 4(a)), i.e., receiver F continues to receive the full stream and node D receives a reduced data rate.

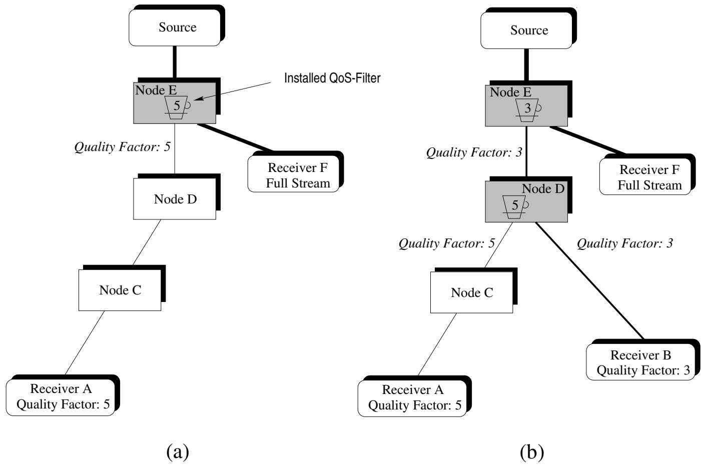

Figure 4: RSVP and QoS filter

In a subsequent step, receiver B joins the session. B requests a video stream with a moderate quality factor of 3 (cf., Figure 4(b)). The RESV message of receiver B is forwarded to node D. Node D needs to merge the different filter requests from receivers A and B. As a result, a RESV message with a requested video quality factor of 3 is forwarded to node E. Moreover, a new filter is activated at Node D, which performs filtering according to quality factor 5. Due to the new request at node E, the filter is re-configured by the QF daemon to quality factor 3.

Such a simple merging style is possible, since the discussed filter types are not additive, i.e., applying the filter twice does not change the result.

3.4 Performance

In an experiment a sender running on a Sun Ultra-1 workstation sent a MPEG-1 encoded video stream with 160×120 pixels/frame, the group of pictures was IBBPBB. The average sizes for I, P, and B frames were 7908 bytes, 7100 bytes, and 3037 bytes. The receiver was running on a 200 MHz 586 PC attached via a 2 Mbit/s wireless LAN (WaveLAN). In the first measurement no filter was used, in the second measurement a filter with quality factor 5 was installed. In both cases the end-to-end delay to transfer a video frame from the sender to the receiver was measured.

Figure 5 depicts the results. I and P frames of the original video stream were much larges than B frames, which results in a relatively long end-to-end delay for these frame types. QoS filter reduce the size of the frames and, thus, lead to a lower end-to-end delay. For I and P frames the savings were approximately 43% (180 ms to 120 ms). Note, that the end-to-end delay includes the processing time in the active node.

Moreover, no losses of frames on the wireless links were observed if a QoS filter was used. In contrast without QoS filters some I and P frames were lost (cf., Figure 5(a)). This is especially important in cases where the coding scheme is very sensitive to loss as in the case of MPEG. It should be noticed, that the unfiltered stream could not be decoded at the receiver. Thus, QoS filtering is highly required in such environments.

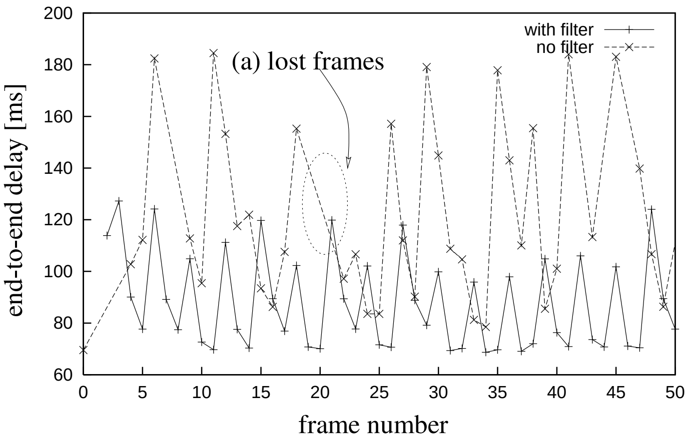

Figure 5: Transfer delay with an without filters

4 Conclusion

The paper addresses the issue of multicasting communication support for heterogeneous group communication. Typical applications that require such services are multimedia and collaborative applications. AMnet is an approach that is based on active networking, i.e., network nodes provide specific services to individual users. Active network nodes implement QoS filters, signalling for QoS filters, error control and synchronization. The existing prototype basically considers QoS filters and RSVP-based signalling of QoS filters. Several types of MPEG-1 filters have been implemented and tested.

Future work will focus on the impact of QoS filtering on synchronization and error control.

References

[1] Elan Amir, Steve McCanne, and Hui Zhang. An application level video gateway. In *Proc. ACM Multimedia '95*, San Francisco, CA, November 1995. ACM.

[2] Samrat Bhattacharjee, Ken Calvert, and Ellen W. Zegura. An architecture for active networking. In *High Performance Networking (HPN'97)*, White Plains, NY, April 1997.

[3] R. Braden, D. Clark, and S. Shenker. Integrated services in the internet architecture: an overview. RFC 1633, ISI, June 1994.

[4] A. Fieger and M. Zitterbart. Transport protocols over wireless links. In *Proceedings of 2nd IEEE Symposium on Computer and Communications (ISCC'97)*, Alexandria, Egypt, July 1997. IEEE.

[5] M. Hofmann. A generic concept for large scale multicast. In *Proceedings of International Zurich Seminar on Digital Communication (IZS'96)*. Springer Verlag, February 1996.

[6] H.W. Holbrook, S.K. Singhal, and D.R. Cheriton. Log-based receiver-reliable multicast for distributed interactive simulation. In *Proceedings of SIGCOMM '95*, Cambridge, MA, August 1995. ACM SIGCOMM.

[7] ISO IEC JTC 1. Information technology - coding of moving pictures and associated audio for digital media up to about 1.5 Mbit/s. International Standard ISO/IEC IS 11172, 1993.

[8] Kai Krasnodembski. Design and Implementation of a Signalling Protocol for Filter Functions (in german). Diploma thesis, TU Braunschweig, August 1997.

[9] Thomas Kupka. Scaling of MPEG-I encoded Videostreams (in german). Diploma thesis, TU Braunschweig, August 1997.

[10] Steven McCanne, Van Jacobson, and Martin Vetterli. Receiver-driven layered multicast. In *ACM SIGCOMM'96*, San Francisco, USA, August 1996.

[11] K.R. Rao and J.J. Hwang. *Techniques and Standards for Image Video and Audio Coding*. Prentice Hall PTR, 1996.

[12] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A transport protocol for real-time applications. RFC 1889, GMD Fokus, January 1996.

[13] D. L. Tennenhouse and D. J. Wetherall. Towards an active network architecture. *Computer Communication Review*, 26(2), April 1996.

[14] David L. Tennenhouse, Jonathan M. Smith, W. David Sincoskie, David J. Wetherall, and Gary J. Minden. A Survey of Active Network Research. *IEEE Communications Magazine*, 35(1):80–86, January 1997.

[15] Hartmut Wittig, Jörg Winckler, and Jochen Sandvoss. Network layer scaling: Congestion control in multimedia communication with heterogenous networks and receivers. In *International COST 237 Workshop on Multimedia Transport and Teleservices*, Vienna, Austria, November 1994.

[16] Lars C. Wolf, Ralf Guido Herrtwich, and Luca Delgrossi. Filtering Multimedia Data in Reservation-Based Internetworks. Technical Report 43.9608, IBM European Networking Center, 69115 Heidelberg, Germany, August 1994.

[17] N. Yeadon. *Quality of Service Filtring for Multimedia Communications*. Phd thesis, Lancaster University, May 1996.

[18] N. Yeadon, F. Garcia, D. Shepherd, and D. Hutchinson. Continuous media filters for heterogeneous internetworking. In *Proceedings of the Conference in Multimedia Computing and Networking 1996*, San Jose, California, January 1996.

[19] Lixia Zhang, Stephen Deering, and Deborah Estrin. RSVP: A New Resource ReSerVation Protocol. Novel design features lead to an Internet protocol that is flexible and scalable. *IEEE network*, 7(5), September 1993.