

Active Multicasting for Heterogeneous Groups

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://ibr.cs.tu-bs.de/~wittmann/AMnet.html>

Abstract

Communication environments are becoming increasingly heterogeneous. This imposes new challenges on communication support for multimedia and collaborative applications. AMnet (Active Multicasting Network) provides multi-point communication support for large-scale groups with heterogeneous receivers. Active Multicasting nodes inside the network include 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 is presented with some first performance results.

Keywords

Group Communication, Active Multicasting, QoS Filter

1 INTRODUCTION

It is now well accepted that scalable and flexible multicasting support is increasingly needed for emerging communication systems. The demand is driven by multimedia applications, especially, by collaborative applications. The same holds for applications, such as scientific computing and distributed simulation. For many of these applications group communication is a natural paradigm. Therefore, proper support within the communication subsystem is required. 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. Thus, all participants in the group are provided with a homogeneous quality of service (QoS).

However, heterogeneity is increasing with respect to networks and end nodes. Consider a scenario where some receivers of the group use simple PDAs, which typically are poorly equipped and often connected over error prone slow-speed wireless links (cf., Figure 1). Others may use high-end workstations that are attached through




Figure 1 Multicast Tree with Active Multicasting Nodes.

a high-speed link. Due to limited processing power 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. Furthermore, group members do not only differ in terms of end-system capabilities, they may also be connected to networks with different characteristics. For example, an ISDN network offers much less bandwidth than an ATM network. Hence, different participants in a multi-point communication may have different requirements with respect to quality of service. The goal of AMnet is in the provision of user-tailored services for individual group members. Data streams with different characteristics will be offered, e.g., to a high-end workstation and to a wireless attached end node. 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 heterogeneous group communication. It is based on Active Networking (Tennenhouse & Wetherall 1996) in the sense that it uses active multicasting nodes inside the network in order to individually tailor data streams to the end users service requirements. Therefore, various service modules shall be dynamically implanted in network nodes. QoS filter and error control modules are typical examples.

An overview of the state-of-the-art in Active Networking can be found in (Tennenhouse et al. 1997).

The paper is structured as follows. Section 2 presents the AMnet approach. An

early prototype is discussed in section 3. It is based on RSVP as signalling protocol for QoS filter functions. Section 4 concludes the paper and gives an outlook on future work.

2 THE AMNET APPROACH

2.1 Basic Scenario

Figure 1 depicts 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, with $QoS A > QoS B > QoS C$. Therefore, the active multicasting nodes comprise service modules for QoS filtering and enhanced QoS signalling. Furthermore, error control capabilities are needed in the active multicasting nodes.

Thus, AMnet is based on active network nodes, that provide user-tailored services for heterogeneous group communication. Particularly, QoS filtering, enhanced signalling, and error control are considered by active multicasting nodes.

2.2 Active Multicasting Node

The structure of an active multicasting node is depicted in Figure 2. It consists of the following main components:

- service modules
- AMnet signalling
- AMnet manager
- QoS monitor
- packet filter

The service modules provide services, such as QoS filtering, error control and recovery, and the like. They can be dynamically activated and configured based on individual service requirements of participants in the group communication. These requirements are propagated through the AMnet signalling protocol. The signalling entity extracts the service requirements and forwards them to the AMnet manager. The AMnet manager is responsible for allocating resources and configuring the service modules accordingly. Information on the current load is provided by the QoS monitor. It observes the service modules and collects information about ongoing communication. The QoS monitor operates similar to the one described in (Zitterbart 1996).

AMnet signalling and the AMnet manager are not involved in the regular data path. The packet filter demultiplexes the incoming data stream and forwards signalling

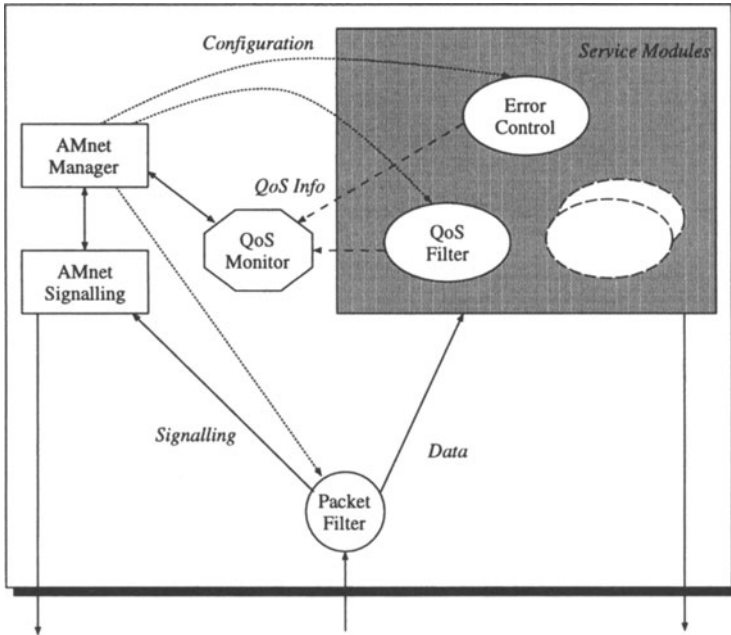


Figure 2 Structure of an Active Multicasting Node.

messages to the AMnet signalling entity and all other data to their corresponding service modules. Filtering is implemented content based, i.e., the corresponding service modules are addressed directly by the filter. Therefore, several demultiplexing steps are integrated in the packet filter.

2.3 Service Modules

The active multicasting nodes support service modules that are needed in order to support quality-based heterogeneous group communication. The service modules comprise, among others, QoS filtering, and error control and recovery.

QoS filtering

The service module QoS filtering reduces the amount of data forwarded to the next node in the multicasting tree towards the destination. QoS filtering is typically used to reduce the data rate of audio and video streams for low-end receivers. Several types of QoS filters may be active on an Active Multicasting Node (e.g., MPEG-1-filter, H.261-filter). QoS filter may perform different techniques to adapt a multimedia stream to specific QoS. As an example an MPEG-1 filter may support the following filter functions:

- Frame dropping filter
- Re-quantization filter
- Monochrome filter
- 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 quantizers can result in strange artefacts (Rao & Hwang 1996).

The *Monochrome filter* removes color information. As MPEG encodes chrominance and luminance information independently, discarding of luminance components is a relatively simple operation.

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, a slicing filter can drop slices easily and, thus, reduce the bandwidth consumed on the output link.

Several QoS filters have been investigated in (Yeaton 1996). However, signalling of QoS filter has not been addressed. Filters for hierarchical encoded streams are presented in (Wolf et al. 1994). They adapt continuous media streams to a specified QoS. Hierarchically encoded streams are also used in (McCanne et al. 1996). It is designed for networks like today's 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 (Wittig et al. 1994). Heterogeneity is not explicitly addressed. Application level gateways, that operate on top of RTP (Schulzrinne et al. 1996) are addressed in (Amir et al. 1995) for filtering conversion operations. Reliable multicast is not addressed in this work.

For the placement of QoS filters in the dissemination tree, several aspects have to be considered. Generally, minimal resources should be allocated with respect to bandwidth and processing power. Moreover, the end-to-end delay of user data should be minimized. Considering bandwidth, an optimal placement allocates a QoS filter as close to the root (i.e., sender) of the multicast tree as possible. The AMnet signalling protocol and the AMnet manager need to address this issue. Furthermore, dynamic join and leave operations must be supported. Therefore, dynamic allocation and re-configuration of QoS filters needs to be supported.

Error Control and Recovery

Active multicasting nodes provide also a sound bases for the implementation of error control and recovery. Thus, local group based concepts as they are discussed in the literature can be easily supported via dedicated service modules.

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 (Holbrook et al. 1995) and LGC (Hofmann 1996) 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 has to be aware of QoS filters.

Typically, mechanisms for error control and recovery are not applied in multimedia applications, such as video conferencing. However, they can still be advantageous for such applications. 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.

In AMnet, QoS filtering is supported by error control and recovery mechanisms. A negative acknowledgement scheme is applied. The active multicasting node collects the NACKs of a local group of receivers, processes them and, if necessary, passes a NACK to an active network node that is located closer to the sender of the multicast tree. The sender only receives NACKs from active network nodes of the highest hierarchy level.

Active multicasting nodes 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 releaves in many cases the burden of retransmissions from the source.

3 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 (Braden et al. 1994). RSVP (Zhang et al. 1993) 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. Currently the prototype consists of QoS filter modules, which are presented in the next section and an AMnet manager, which performs signaling tasks.

3.1 QoS filter

With respect to QoS filter, the focus was on compressed video streams that are MPEG-1 coded. The implementation and integration of an H.261 video is currently

under way. Other filters can be found in (Yeadon et al. 1996). The filter types presented in section 2.3 have been implemented in AMnet (Kupka 1997).

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 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 RSVP Extensions

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. Due to its flexible concept, RSVP can be 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 3(a). To assure compatibility with RSVP entities that 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 3(b) depicts a QF object for the implemented MPEG filter (cf., above).

The RSVP daemon has been extended with an interface to the AMnet manager to exchange QF objects (Krasnodembski 1997). In case a QF object is included in the received reservation (RESV) message, the QF object is extracted and forwarded to the AMnet manager.

3.3 Filter Merging

The AMnet manager is responsible for allocation and configuration of the QoS filter according to the requirements stated in the QF object. If the AMnet manager 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 AMnet manager keeps some state information concerning filter requirements of group members. The following cases can be distinguished:

1. no previous request
2. previous request for same QoS filter type
3. previous request for other QoS filter type

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/BW-Mode: 0	

(b) Example: MPEG Filter Class.

Figure 3 RSVP Filter Class.

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 AMnet manager 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 AMnet manager, two cases need to be distinguished.

Firstly, the new request can be merged with the current request and the resulting request does not impose higher requirements than the existing request. In this case, the corresponding filter functionality is activated by the AMnet manager. 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 newly 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 QoS level as the requirements of the joining group member, the filter may need to be moved to an active network node that is not in the dissemination path of the new member.




Figure 4 RSVP and QoS Filter.

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

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. This is first case mentioned in section 3.3. Thus, the RESV message is forwarded to Node E. Node E has already seen a RESV message without bandwidth constraints from receiver F. This situation corresponds to the third case in section 3.3. As a result, node E has to install a new 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.

In a subsequent step, receiver B joins the session. B requests a video stream with a moderate quality factor of 3 (cf., Figure 4(a)). 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. Once again, this situation corresponds to the third case in section 3.3. As a result, an RESV message with a requested video quality factor of 3




Figure 5 Filter vs. no filter.

is forwarded to node E. Moreover, a new filter is activated at Node D that performs filtering according to quality factor 5. Due to the new request at node F, the filter is re-configured by the AMnet manager to quality factor 3.

If B would have requested quality factor 5, no re-configuration would have been necessary. This corresponds to the second case listed in section 3.3.

3.4 Performance

In this section some performance results of the AMnet prototype are discussed. The end-to-end transfer delay of an MPEG-1 encoded video stream was measured with 160×120 pixels/frame. The group of pictures (GOP) was IBBPBB. Figure 5 depicts the results obtained when sender, receiver, and the active multicasting node were running on Sun Ultra 1 workstations. The workstations were connected through a lightly loaded 100 Mbit/s Fast-Ethernet. In the first experiment the video stream was transferred without filtering. In the second experiment a high quality filter (quality factor 1) was requested. In the third experiment a very low quality level (quality factor 5, dropping of B frames) was requested. Although the transfer delay depends on the requested quality level, QoS filtering introduces a substantial processing delay. Therefore, the video stream without filters leads to the best end-to-end delay. Due to the lightly loaded network with practically no data losses, QoS filters are not needed.

Figure 6 depicts results obtained in a more realistic scenario. The receiver was running on a 200MHz 586 PC, attached via a 2 Mbit/s wireless LAN (WaveLAN). The results demonstrate that QoS filtering can considerably reduce the end-to-end delay. Moreover, we obtained a reduced loss-rate of frames when QoS filter are used




Figure 6 Wireless vs.wired transfer.

(cf., Figure 6(a)). This is especially important when 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.

4 CONCLUSION

The AMnet approach presented in this paper provides communication support for heterogeneous group communication. Typical applications that require such services are collaborative applications and distributed simulations. AMnet is based on active networking, i.e., network nodes provide specific services to individual users. Active multicasting nodes implement QoS filters, signalling for QoS filters, error control and synchronization. The existing prototype basically considers QoS filters (MPEG-1) and RSVP-based signalling of QoS filters. The AMnet prototype is available at <http://ibr.cs.tu-bs.de/~wittmann/AMnet.html>.

Future work will focus on the impact of QoS filtering on synchronization and error control. Furthermore, merging and dynamic re-location of QoS filters will be studied in greater detail. Moreover, the prototype will be enhanced to allow for dynamically loadable service modules.

REFERENCES

Amir, E., McCanne, S. & Zhang, H. (1995), An application level video gateway, in 'Proc. ACM Multimedia '95', ACM, San Francisco, CA.

- Braden, R., Clark, D. & Shenker, S. (1994), Integrated services in the internet architecture: an overview, RFC 1633, ISI.
- Hofmann, M. (1996), A generic concept for large scale multicast, in 'Proceedings of International Zurich Seminar on Digital Communication (IZS'96)', Springer Verlag.
- Holbrook, H., Singhal, S. & Cheriton, D. (1995), Log-based receiver-reliable multicast for distributed interactive simulation, in 'Proceedings of SIGCOMM '95', ACM SIGCOMM, Cambridge, MA.
- Krasnodembski, K. (1997), Konzeption und Realisierung eines Signalisierungsprotokolls für Filterfunktionen, Diploma thesis, TU Braunschweig.
- Kupka, T. (1997), Skalierung von MPEG-I Videostreamen, Diploma thesis, TU Braunschweig.
- McCanne, S., Jacobson, V. & Vetterli, M. (1996), Receiver-driven layered multicast, in 'ACM SIGCOMM'96', San Francisco, USA.
- Rao, K. & Hwang, J. (1996), *Techniques and Standards for Image Video and Audio Coding*, Prentice Hall PTR.
- Schulzrinne, H., Casner, S., Frederick, R. & Jacobson, V. (1996), RTP: A transport protocol for real-time applications, Technical report, GMD Fokus.
- Tennenhouse, D. L., Smith, J. M., Sincoskie, W. D., Wetherall, D. J. & Minden., G. J. (1997), 'A survey of active network research', *IEEE Communications Magazine* 35(1), 80–86.
- Tennenhouse, D. L. & Wetherall, D. J. (1996), 'Towards an active network architecture', *Computer Communication Review* 26(2).
- Wittig, H., Winckler, J. & Sandvoss, J. (1994), Network layer scaling: Congestion control in multimedia communication with heterogeneous networks and receivers, in 'International COST 237 Workshop on Multimedia Transport and Teleservices', Vienna, Austria.
- Wolf, L. C., Herrtwich, R. G. & Delgrossi, L. (1994), Filtering Multimedia Data in Reservation-Based Internetworks, Technical Report 43.9608, IBM European Networking Center, 69115 Heidelberg, Germany.
- Yeadon, N. (1996), Quality of Service Filtering for Multimedia Communications, Phd thesis, Lancaster University.
- Yeadon, N., Garcia, F., Shepherd, D. & Hutchinson, D. (1996), Continuous media filters for heterogeneous internetworking, in 'Proceedings of the Conference in Multimedia Computing and Networking 1996', San Jose, California.
- Zhang, L., Deering, S. & Estrin, D. (1993), 'RSVP: A new reavailable ReSerVation Protocol. novel design features lead to an Internet protocol that is flexible and scalable', *IEEE network* 7(5).
- Zitterbart, M. (1996), User-to-User QoS management and monitoring, in 'Proceedings IFIP Workshop on Protocols for High Speed Networks', Sophia Antipolis, France.