LOCATING SERVICES IN PROGRAMMABLE NETWORKS WITH AMNET-SIGNALLING

Anke Speer, Ralph Wittmann, Martina Zitterbart
Technical University of Braunschweig
[speer|wittmann|zit]@ibr.cs.tu-bs.de

Abstract AMnet provides a framework for flexible and rapid service creation. It is based on Active and Programmable Networking technologies and uses active nodes (AMnodes) within the network. These AMnodes are executing service modules for the provision of individual communication services. Using on-demand loadable service modules helps to enhance the functionality of intermediate systems without long global standardization processes.

Placing application-dedicated functionality within the network requires a flexible signalling protocol to locate and announce as well as to establish and maintain the corresponding services. The AMnet signalling protocol discussed within this paper uses active packets – evaluation packets – for service location. The AMnet service control architecture further comprises periodic service announcements to make available services known to receivers.

Keywords: Active Multicasting, Service Location, Active Signalling, Active Networking, Programmable Networking

1. INTRODUCTION

In the past decades a tremendous growth in the use of computer networks in general and the Internet in particular can be noticed. Besides the provided interpersonal communication, like e-mail, the networks advance in becoming the media for distributed computation, telecollaboration, distance learning, e-commerce and the like. Many of these applications are inherently based on group communication. Therefore, efficient and scalable group services are needed for proper communication support. The particular group service to be provided highly depends on the application itself and on the media transported, e.g., text, audio or video. Issues that vary within these group services are reliability and

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-0-387-35522-1_37

H. R. van As (ed.), *Telecommunication Network Intelligence*

quality of service support. It is important to note that no single group service will be able to serve the high variety of existing and emerging group applications.

With respect to group members, situations may appear in which different group members would like to experience a different group service leading to a *heterogeneous group service*. This can be motivated due to a low-bandwidth access link (e.g., a wireless link) or due to a lower monetary budget of a group member. Such group members should not be able to degrade the overall service quality of a collaborating group. However, 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.

The AMnet approach makes user-tailored data streams available to individual receivers and, thus, enables heterogeneous group services [15]. It therefore utilizes dedicated support within the network and applies the ideas of Active and Programmable Networking to avoid global standardization processes. *Service modules* are dynamically placed on active network nodes (AMnodes) dependent on user requirements. These service modules then individually tailor data streams according to the group members service requirements. Typical examples of such service modules are QoS-filter and modules for error control.

Figure 1 shows a tele-collaboration scenario – a typical example for possible AMnet support – where the participants differ in their Quality of Service (QoS) requirements. According to the AMnet approach the AMnodes are responsible for adapting the original data stream in accordance with the desired QoS-levels.

It should be highlighted that AMnet provides a general framework for heterogeneous group communication that is not limited to service modules like QoS-filter or error control. In contrast, AMnet provides an open framework where new services can easily be incorporated. An active signalling protocol is used for the on-demand placement of service modules. It is discussed in detail within this paper.

Current Heterogeneity Support. Currently, different approaches to signal and provide heterogeneous communication services are developed. The provision of heterogeneity can be distinguished in end-to-end support and network support.

To provide heterogeneity, a possible action of the sender is hierarchical coding of the output stream. Receivers can select different media qualities by joining appropriate multicast groups related to the coding levels (e.g., [13]). In [1] a communication model is described where hi-

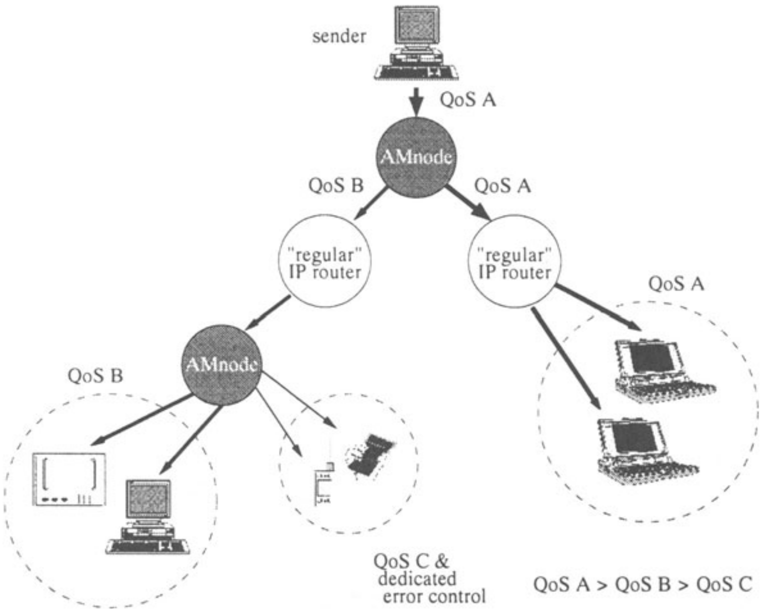

Figure 1 Basic AMnet scenario: Tele-collaboration

erarchical coding into one output stream allows for lossy decompression at the receiver. Dependent on the loss ratio the original signal can be restored on a corresponding quality level.

The model of integrated services within the Internet [5] is a step towards heterogeneous services. Using the RSVP protocol heterogeneous resource reservations are possible. However, current flow specifications limit heterogeneity support to network performance parameters.

A well-known way to place enhanced functionality within the network is the establishment of transport level [4] or application level gateways [3, 7]. Common to those solutions is the withdrawal of transparent end-to-end network operation. The gateway system which hosts the additional functionality is the peer entity of both sending and receiving system.

Active and Programmable Networking. Based on the idea of placing additional functionality inside the network the concepts of Active and Programmable Networking were developed. These concepts allow for application level processing within network nodes, such as IP routers. This is done with capsules [18] in the Active Networking approach and via programmable switches in the Programmable Networking approach

[22]. Capsules carry themselves the code to be executed on intermediate systems [8]. Whereas, programmable switches execute pre-loaded application code on packets flowing through them [22]. Both approaches help to enhance the functionality of the network without long and intractable standardization processes.

According to different research work the provision of group services with Active and Programmable Networking technology appears to be a promising idea. Not only capsules [20] but also programmable switches [2] and in particular, sophisticated acknowledgment [6] and congestion control [9] are proposed to provide basic multicast control mechanisms. However, these approaches do not consider the support of heterogeneous group communication.

To put it in a nutshell, it can be said that the potential of group communication with heterogeneous services has not been investigated in detail in the context of Active and Programmable Networking.

In contrast, our work focuses on service provision in heterogeneous group communication environments. The establishment and maintenance as well as the location and announcement of corresponding services requires a flexible signalling protocol which is presented throughout this paper. The paper is organized as follows. Section 2 introduces the AMnet approach for the provision of heterogeneous group communication. In section 3 the service concept of AMnet is described. Section 4 presents the general concepts of the active signalling protocol with basic implementation approaches. Section 5 comprises some results of **ns**-simulations performed to analyze the protocol behavior. Section 6 closes with a summary and an outlook on ongoing research.

2. AMNET

In the following, a brief introduction into the AMnet service model is given. A more detailed discussion of AMnet can be found in [15]. Basically, AMnet is based on IP-multicast and the group concept in general.

AMnet aims at the provision of scalable heterogeneous group communication with efficient and rapid service creation. It is based on the placement of additional functionality inside the network by *service modules.* Service modules are responsible for the adaptation of data streams to specific service demands. These modules are dynamically loaded by Active Multicasting Nodes (AMnodes) which form the core building blocks of AMnet [14] and operate on the communication path between sender and receivers. In this sense AMnet represents a Programmable Networking approach.

A goal of AMnet is to support service heterogeneity transparently to the origin of a data stream as well as to the receivers. Thereby, AMnet follows a receiver oriented approach. Loadable service modules with out-of-band signalling are used for service creation instead of capsules. The reasons are two-fold. Firstly, capsules have to be provided by the sender, i.e., the desired sender transparency and receiver orientation could not be achieved with capsules. Secondly, hardware support with dedicated service modules for performance improvement becomes possible for AMnet with out-of-band signalling whereas a combined HW/SW solution is envisioned. As the structure of the hardware dependent modules differs significantly from software-based modules, a capsule-based approach is not easily applicable.

Our approach for rapid service creation uses programmable switches, the active signalling protocol described in this paper, however, uses capsules to select appropriate service providers. These capsules are called *evaluation packets* and contain an evaluation program to be executed on the AMnodes in order to determine whether the AMnodes satisfy the qualifications for providing a service individually for a receiver.

Placing application-dedicated functionality within the network raises some questions: where should those services be located, how should they be established and maintained, how should a receiver be associated to a dedicated service and how should different services be managed within a *session*? In this context, a session refers to a communication scenario where a designated sender issues a data stream which can be received from several communication participants without or after adaptation, thus, different service levels are provided. Altogether the foregoing questions have to be solved by a flexible signalling protocol which provides a receiver-controllable localization of the service modules.

In a first prototype implementation RSVP was used as signalling protocol [21]. With this prototype QoS filters were requested by receivers via RSVP. To this end appropriate parameters are added to RSVP reservation messages. But RSVP does not offer the flexibility to support optimal allocation of service modules, because RSVP tends to locate the service as close to the sender as possible. While this is desirable for media translation, protocol adaptation can not be realized this way. In the following sections the model of a new active signalling protocol to be used within AMnet is explained.

3. AMNET SERVICE CONCEPT

This section presents first how AMnet is able to provide heterogeneous group communication. Then the mechanisms for service location and service announcement are described.

3.1. SERVICE LEVELS

Service heterogeneity within a session needs to be bound to a manageable degree of diversity. Therefore, one concept of AMnet service control is to logically group receivers with similar service demands into distinct multicast receiver groups – the *service level groups* (see Figure 2), distinguishable by their multicast-addresses. The receivers join the corresponding group on demand through IGMP [10].

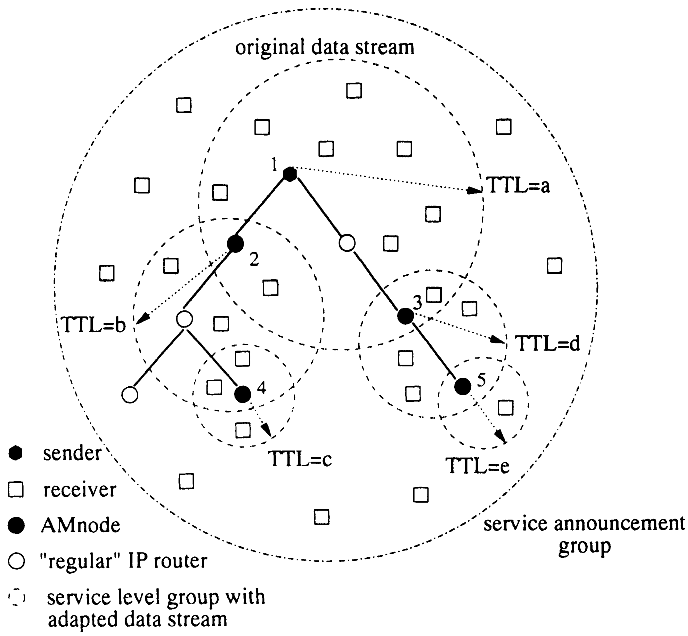

Figure 2 Example multicast tree

Each service level group within a communication session represents all receivers whose service demands can be resolved with a single group service. Therefore, each group represents a different view onto the same original data corresponding to the adaptation within the AMnodes. The communication service offered by AMnet can be described by a hierarchically ordered tree of service level groups (e.g., Figure 2). The service of a group is supported by an AMnode through the use of appropriate service modules. The actual service may be derived from the processing of the original data stream (see AMnode 1 in Figure 2) or from the service of another service level group (see AMnode 2 in Figure 2). The scope of an AMnet service level group is limited by the actual TTL value assigned to packets issued by the corresponding AMnode. The TTL value of every group must not exceed the scope of the service announcement group (see below).

Furthermore, the establishment of service level groups permits the provision of different service qualities within one region without the services interfering with each other. Two data streams with different media formats or individual error control must sometimes coexist on a communication link and have to be distinguishable by the appropriate receivers. Therefore, all packets are explicitly assigned to corresponding service level groups by their multicast destination address.

The hierarchical order of service level groups allows for an efficient establishment of different quality levels within a session. One distinct service quality might be easily derived from another already available quality level if, for instance, only a different (weaker) error control policy for network overload conditions has to be inserted into the already adapted data stream.

The service quality experienced at the receiver is a function of the service level of the group and the current network conditions between group source and each individual receiver. The service within a group can only differ in performance-oriented, packet-based service parameters such as delay or loss probability. Other parameters which define the content-based nature of the service are homogeneous (e.g., media format, acknowledgment strategy, error control) within a service level group. However, the distinction of services can be triggered by both: performance oriented parameters and content. Consider two receivers with very different loss probabilities. In that case different acknowledgment strategies and error control mechanisms might be necessary which require different service level groups.

The hierarchical ordering of services does not automatically imply hierarchical degradation of all service parameters. Some parameters can be provided unchanged, other parameters even improve. As an example, the insertion of a new service can improve media playback quality due to less jitter at the cost of higher, but uncritical delay. This could be useful for video distribution, for example.

3.2. SERVICE LOCATION AND ANNOUNCEMENT

In contrast to the capsule approach, AMnet's service creation is not based on executable code to be transmitted in data packets. Instead, the AMnet signalling provides procedures to locate and announce service modules and to establish and maintain appropriate services on the AMnodes. These signalling procedures utilize Active Networking technologies. Control of AMnet services is maintained completely out-of-band.

The following components are part of the service control. Figure 3 shows their correlation:

- Session announcement,
- service announcement,
- service module repository, and
- service access.

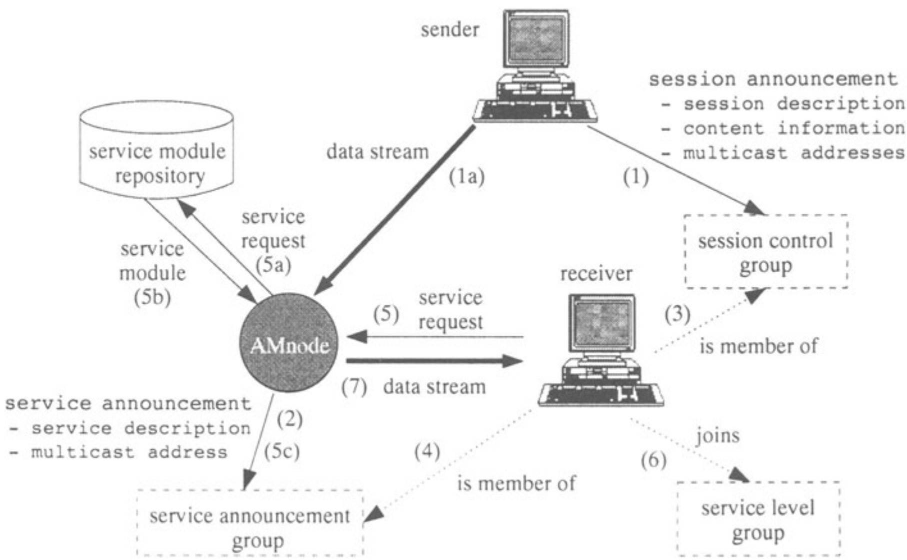

Figure 3 Service Control Architecture

Session Announcement. A session announcement is advertised by a sender on a separate multicast group – the *session control group.* In this group every AMnet session is announced similar to the well-known Session Announcement Protocol (SAP) [12] of the MBone. SAP itself was not utilized to avoid the use of its standardized port and because the overhead of introducing the needed new data-fields and functionalities was considered as too much.

The session announcement contains a description of the session, including information such as bandwidth and delay requirements, and content specific information, e.g., data format and compression scheme. Moreover, the path the session announcement took from the session sender to a specific receiver is included as well in order to reconstruct the path (see below). Additionally, the description contains the multicast address of the original data stream and the multicast address of the service announcement group. Based on the session description receivers decide whether to join a session.

Service Announcement. Whereas the session control group provides information of all available AMnet sessions, the *service announcement group* forms a data base of the available AMnet services for a given session. This group is used for the signalling of heterogeneous service capabilities and demands, i.e., all available service level groups are announced within this group. All session participants are member of the service announcement group to be able to learn about available services. Because if a new service is established by an AMnode, the node advertises the appropriate service description within this group.

Service Module Repository. The service module repository is a distributed data base which contains service modules and their description. The purpose of the repository is to make service modules available to an AMnode in case that a service module is not already cached at that node. Service modules can be stored in the service repository by AMnet users and network management procedures.

Service Access. The mechanism for service access follows the same principles as the current practice of the MBone. If a participant wants to use an AMnet service, e.g., media translation or enhanced error control, the participant checks by joining the session's service announcement group whether the desired service is already available. If the service exists, the participant simply joins the associated group of the best matching AMnode that provides that service. If the service does not exist, the participant may ask an appropriate AMnode to provide and announce that service.

4. SIGNALLING AMNET SERVICES

This chapter describes first the concepts of the active signalling protocol, afterwards the prototype implementation is presented.

4.1. GENERAL CONCEPT

An important aspect of the AMnet service control is the selection of an appropriate AMnode for service provision. This selection depends on the type of the requested service and on the metrics propagated by the receiver of the session towards the AMnode. For example, one receiver may choose to use a media translation service which has to be placed as close to the session sender as possible, whereas another receiver may choose a protocol adaptation service which has to be placed near by the receiver itself. For the purpose of selecting an appropriate AMnode, each session receiver comprises a *performance meter.* According to the metrics and according to the description of the desired service the performance meter rates the service announcements the receiver gets from the joined service announcement group in order to determine whether an appropriate service is already in operation.

For the selection procedure, two cases need to be distinguished:

- the desired service is already provided at the time a receiver joins the session
- no appropriate service is available

In the first case the selection takes place within the set of AMnodes that already provide the desired service, in the second case it is selected within the set of AMnodes that fulfill the propagated metrics and, therefore, potentially can provide the requested service.

In both cases active packets – so-called *evaluation packets* – are used for the selection. These packets contain the following data:

- the path the session announcement took from the session sender to the receiver; this information is included in the session announcement and is needed because every AMnode has to know its parent AMnode to correctly forward the evaluation packet (see below). Caching this information in the AMnode is considered as too much overhead.
- the address and the result of the best matching AMnode so far resp. the address and the result of the actually evaluated AMnode (see below)
- the description of the desired service

- the evaluation program to be executed on each AMnode which receives such an evaluation packet to determine whether it satisfies the requirements of being a service provider

The way the packets are forwarded and the AMnodes are evaluated slightly differ if there are already providing AMnodes or if an AMnode has to be found which fits into the metrics and can provide the desired service:
If one or more AMnodes currently provide the desired service, the performance meter of the session receiver issues the evaluation packet to each of these AMnodes in order to evaluate them according to the propagated metrics. This is done per unicast since it is assumed that there is only a very limited number of corresponding AMnodes available. The AMnodes attach the evaluation results to the evaluation packet and return it to the receiver. After the performance meter has collected all responses, it selects the best matching AMnode among the queried ones and makes the receiver join the appropriate service level group.

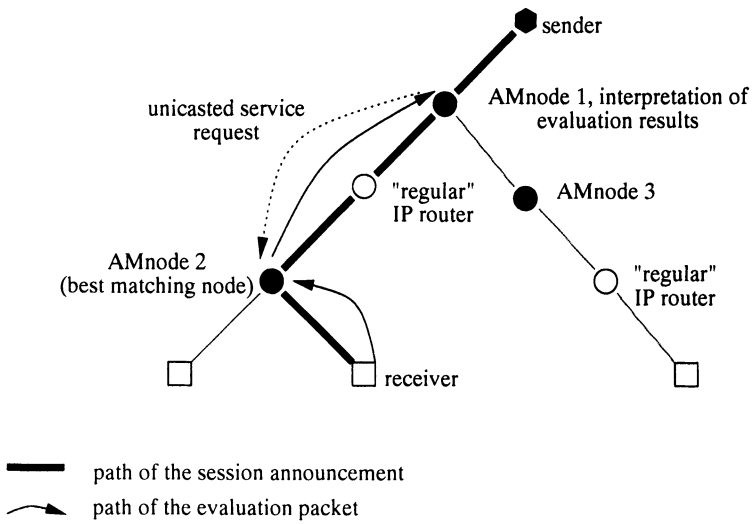

Figure 4 Scheme of the basic tree search

In case no AMnode provides the requested service yet or no providing AMnode fits to the propagated metrics a slightly different selection procedure is used. That means in the worst case, both selection procedures take place one after another. This selection procedure uses a different mechanism – the so-called *basic tree search* – to determine a suited AMnode. Again, the performance meter of the session receiver issues an evaluation packet. This is forwarded to the closest AMnode

(AMnode 2, Figure 4). This node can be identified by the path-entry included in the evaluation packet and known from the session announcement. According to the description of the desired service contained in the received evaluation packet and according to the available node resources, the AMnode management decides whether the node is capable of providing the service. In the positive case, it executes the evaluation program and incorporates the result into the evaluation packet. This, however, is only performed if the actual AMnode fits better than the AMnode previously noted in the packet – if there is any. Therefore, the evaluation packet contains at most one address: the address of the best matching AMnode so far. If the service can not be provided, because of limited resources or similar reasons, or if a better matching AMnode is already cached in the evaluation packet, the packet remains unchanged.

With the basic tree search the actual evaluated AMnode does not directly respond back to the session receiver. In contrast, it forwards the evaluation packet to its parent AMnode on its way to the session sender (AMnode 1, Figure 4). At this node the admission test and evaluation takes place again. This is repeated until the evaluation packet arrives at the last AMnode in front of the session sender (AMnode 1, Figure 4). This AMnode is responsible for the interpretation of the evaluation result carried by the evaluation packet. This last AMnode, then, unicasts a service request for the desired service – known from the description in the evaluation packet – to the best matching AMnode (e.g., AMnode 2, Figure 4).

After the selected AMnode received the request to provide the service, it downloads the appropriate service from the service module repository and announces the establishment in the service announcement group. Since the session receiver which started the whole procedure is a member of the service announcement group it learns about the availability of the desired service and subsequently can join the corresponding service level group.

The session announcements are sent periodically. Therefore, it is possible to notice newly available AMnodes – they will be included in the path-description of the session announcement – and it is as well possible to include them dynamically into the evaluation.

Since the basic tree search mechanism is limited to AMnodes being located on the multicast tree, an *extended tree search mechanism* is under development. AMnodes beyond the path might as well be able to provide the service and fulfill the metrics even in a better way than the AMnodes on the path. The extended tree search can possibly use the concepts of IP-multicast and forward the evaluation packets to the neighbor AMnodes of the AMnodes on the multicast tree. Therefore, the

AMnodes have to be reachable over a known multicast address. Then, for example, a ring search evaluation can be processed on the AMnodes beyond the multicast tree limited through a TTL. The value of the TTL has to be decreased with increasing distance from the receiver.

It is obvious that the evaluation packets are capsules as introduced in the context of Active Networking. Therefore, overall it can be noticed that AMnet is based on Programmable Networking technologies with service creation through loadable service modules, but as well on Active Networking technologies (i.e., capsules) for the location of probable service providers.

4.2. PROTOTYPE IMPLEMENTATION

Currently, an implementation of the signalling protocol is in progress. This subsection gives a short overview about the basic design considerations.

Java is the chosen programming language with the Java Development Kit 1.2 of SUN because of its platform-independent facilities. The session announcement and the service announcement group are realized through IP-multicast groups. The signalling components on the session sender, on the receivers and on the AMnodes are implemented as Java classes.

The analysis of the requirements for the service module repository showed that they can be fulfilled widely by an LDAP server. Therefore, we selected LDAP for the design and implementation of the repository. Other reasons for utilizing LDAP were its easy use and its extensibility as well as the availability of Java implementations (e.g., from Netscape). In [16] a suggestion for recording Java objects into an LDAP directory information tree is given. Table 1 gives the attributes which describe the service modules. With this description the modules can be stored and recorded directly in the LDAP server.

5. SIMULATION EXPERIMENTS

In order to evaluate the signalling protocol outlined above, simulation experiments were carried out in [17] using the simulator **ns-2** [19] from U.C.Berkeley/LBNL. Issues of interest throughout the simulations include the frequency of AMnode setups, the delay between service request and provision and the like.

The network topology used for the simulation experiments consisted of 250 nodes. It was automatically generated by GT-ITM [11]. The backbone is constructed out of 10 nodes with up to 3 belonging subnets with 8 nodes at an average (see Figure 5).

Attribute	*Meaning*
moduleName	name of the service
serviceType	e.g., QoS filter or error control
description	description of the service
os	operating system needed by the service
osVersion	version number restriction for the operating system
format	data format in which the service is stored within the directory information tree
moduleType	mode of the service (e.g., Java, shared library)
moduleData	service data
objectClass	if the service is a Java module, objectClass has to give the name of the executing class
className	if the service is not just Java code, className has to give the name of the class which starts that service

Table 1 Description attributes for service modules

The simulations described in [17] were all performed on the same network topology (see Figure 5) with the same receiver distribution and requirements but with different input parameters. As routing protocol, a DVMRP-like protocol was used. The metrics propagated by all receivers was the minimization of the distance between the receiver and the AMnode which provides the desired service.

The variable configuration parameters were:

- the lifetime of a service module on an AMnode,
- the time for loading and installation of service modules on AMnodes,
- the maximum allowed distance between the AMnode with the appropriate service and the receiver,

In the following, selected results of the simulations are presented. In every simulation 42 AMnodes are available. They are located on edge routers interconnecting the backbone with the local subnetworks (e.g., AMnode 162 in Figure 5). Moreover, every node in the backbone can operate as an AMnode.

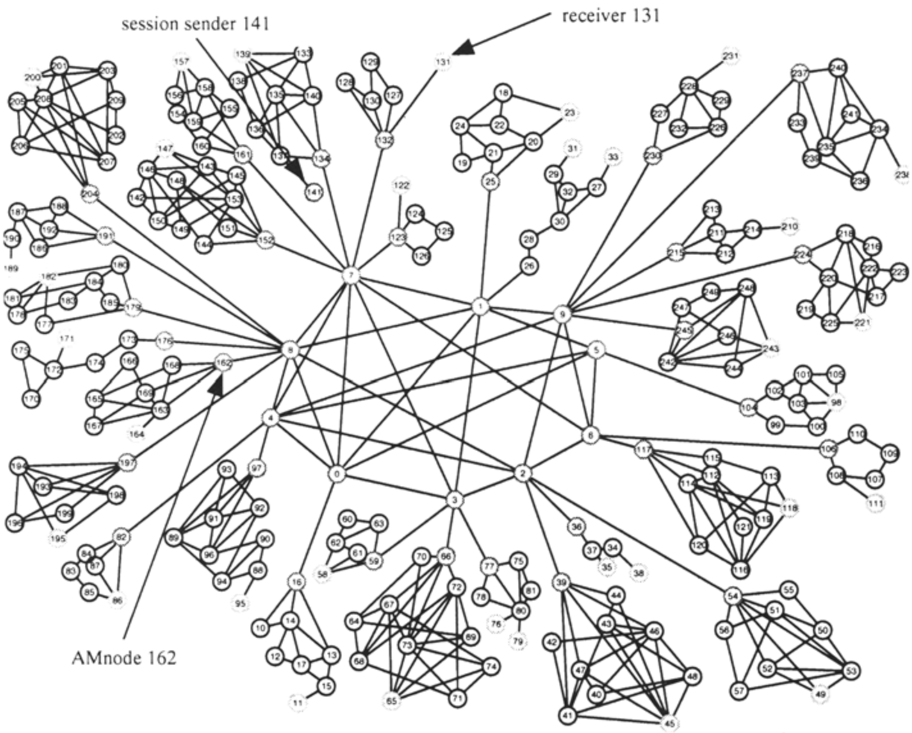

Figure 5 Simulation topology

In some arbitrary chosen subnets additional AMnodes are placed. 33 receivers with 10 different requests are distributed across the 30 subnets. They join the session successively every 50 ms. However, every second receiver just announces its desired service and leaves the session right afterwards. This is done to simulate dynamic scenarios and to simulate situations in which service modules start to operate without having an actual user. The session announcement is repeated every 10 ms as well as the service announcement of every available service.

Before the first receiver joins the session no service module on an AMnode is available except for the session sender where the original data stream is produced. Therefore, the first ten receivers with different requests have to start the basic tree search to find an appropriate AMnode capable to provide the desired service. Thus, every simulation has a starting phase of about 5 seconds until the first ten different receivers have joined. Whether the next joining receiver has to start the basic tree search, too, depends on the lifetime of the services – if the appropriate services are still in operation, the receiver can join the corresponding service level group – and on the distance between the AMnode which already provides the desired service and the receiver – the distance has to be smaller than the maximum allowed distance.

For every simulation the following items were observed:

- the delay at the receiver in receiving the desired data
- the number and distribution of AMnodes with active service modules
- the cumulative number of signalling packets sent/received from/at the nodes
- the number of signalling packets sent from two distinct nodes
- the number of signalling packets received at two distinct nodes

As an example, the observed results for receiver 131 are presented (see Figures 6 to 8). The observed period of time is 17 seconds, the first receiver joins the session at 0.5 seconds. The number of packets it received during different simulation experiments – dependent on the variation of the maximum allowed distance between the AMnode with the appropriate service and the receiver – is depicted in Figure 6. The influence of the lifetime of a service module on an AMnode is shown in Figure 7. The loading and installation overhead of service modules on AMnodes is shown in Figure 8.

Figure 6 shows that the reduction of the maximum allowed distance between the receiver and its corresponding AMnode increases the number of packets received at node 131. This is due to the fact that similar service requests of AMnodes which have a distance of less than twice the allowed distance from each other can be combined. Therefore, less AMnodes have to provide the desired service and less service announcements are sent. If the maximum allowed distance is very small, just a few similar requests can be combined. That means, another AMnode has to provide the service and additional service announcements have to be sent, so the number of received packets at the receiver increases. The results of the simulations with a distance of 10 respectively 15 AMnodes are very similar because of the diameter of the chosen network. The diameter is between 10 and 18 nodes. Thus, 10 is a good choice for the distance because then the best combination of similar service requests is achieved. The delay time between the service request and the receiving of the data is equal in the simulation experiments with 10 and 15 nodes as maximum distance. However, the delay times decrease with decreasing distance because the service providers are closer to the receiver.

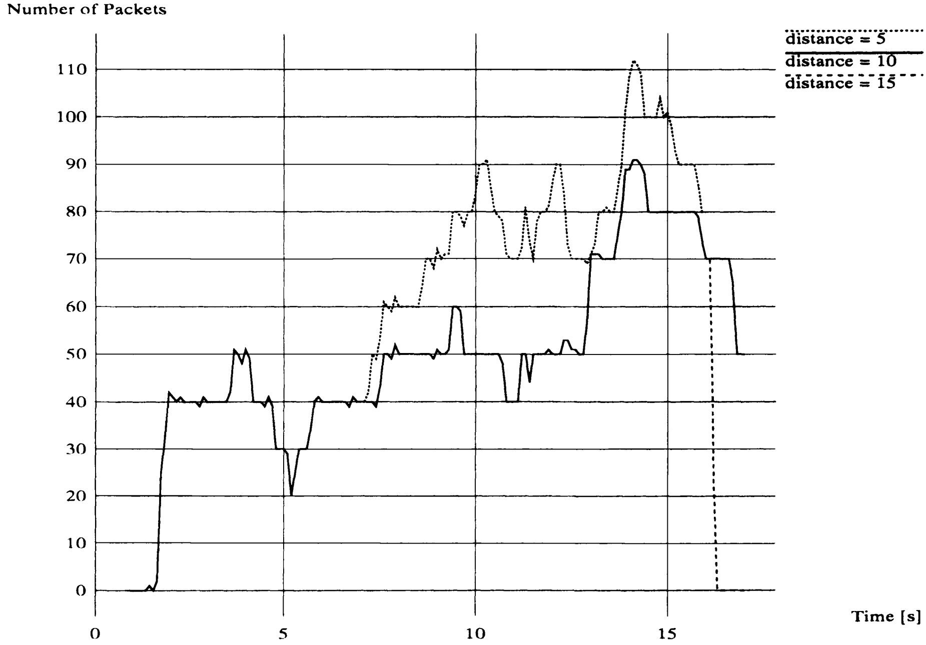

Figure 6 Influence of the maximum distance

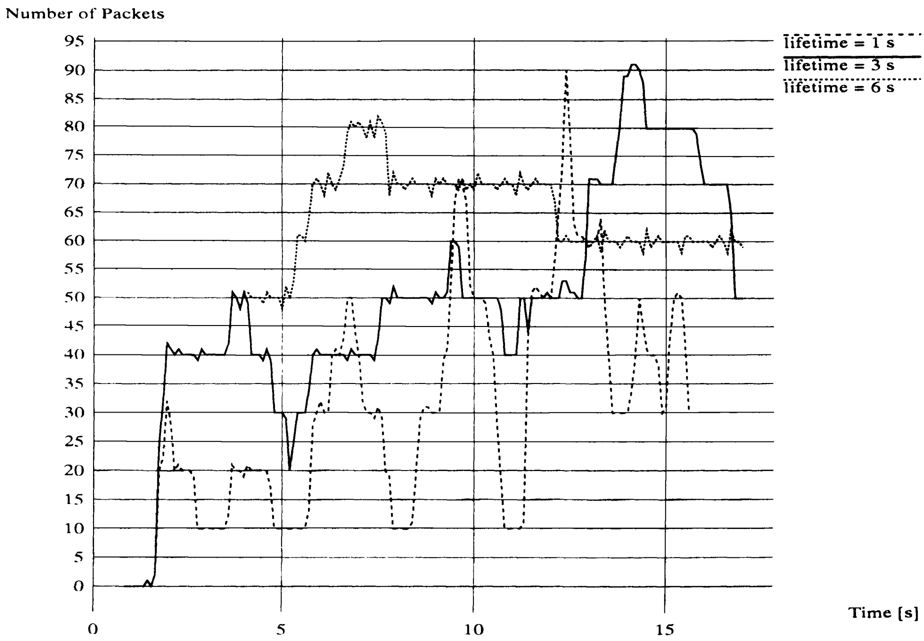

Figure 7 Influence of the lifetime of the service modules

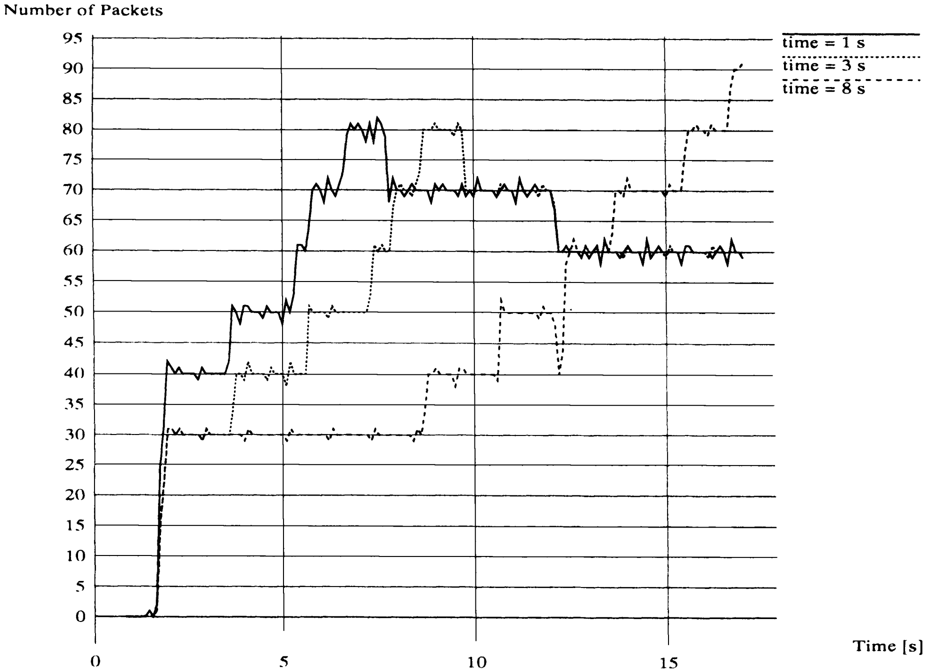

Figure 8 Influence of loading and installation overhead

Additionally, the tree search needs not to look further than 5 nodes up in the multicast tree.

Figure 7 shows the effect of a longer lifetime of the service modules. If a service module has no users, it operates just for as long as its lifetime is. If a receiver starts to use this service during its active time, the service is refreshed and can again operate for its lifetime. Because even modules which do not have any users operate for the whole lifetime of 6 seconds, receiver 131 gets a lot more packets than if the lifetime is just 3 respectively 1 seconds. The disadvantage with short lifetimes is that the receiver may have to require its desired service frequently. The delay times for the receiving of the required data are very similar because the lifetime of the service modules do not have any influence on that.

Figure 8 underlines that the loading and installation overhead of service modules has strong effects on the number of received packets. The graph belonging to the simulation with 1 second loading time firstly is very steep. The number of received packets increases very fast due to the fact that service requests are fulfilled rapidly and, therefore, many AMnodes start sending right after the receivers' request. The graph belonging to the simulation with a loading time of 8 seconds is considerably less steep. For the first 8 seconds just the session announcement packets were measured. Only after 8 seconds when the first service starts to operate the number of packets increases. Clearly, the delay times for receiving the required data increase similar to the growth of the loading and installation overhead.

The simulations resulted in the following conclusions:

- The longer the lifetime of a service module the rarer do the receivers have to look for a new service provision. Thereby, the time while a receiver has to wait for its data and the signalling effort which is necessary for the search of a new service provider decreases. However, long lifetimes of service modules produce high overload on the AMnodes: If the appropriate group stays without any participants, the resources of the AMnode are unnecessarily wasted.

- The greater the allowed distance between sending AMnode and receiver the fewer sending AMnodes are necessary because distributed but similar requirements can be combined and fulfilled by the same AMnode.

- The longer the time for loading and installation of service modules on AMnodes the more identically service provisions are requested. There is still no mechanism which indicates to a new joining receiver that its desired service was already requested and that the search is still in progress.

6. SUMMARY AND OUTLOOK

In this paper an outline of the AMnet approach for flexible and rapid service creation in IP based networks was presented. The AMnet signalling protocol for locating and announcing services provided by active intermediate nodes (AMnodes) was discussed in some detail. Programmable Networking technologies are used with service creation through loadable service modules. This mechanism helps to enhance the functionality of the network without long global standardization processes. Active Networking technologies are utilized for signalling purposes by active packets – the so-called evaluation packets. Simulation experiments showed the influence of the propagated metrics, the lifetime and performance of service modules and the influence of loading and installation overhead.

Future work concerning the signalling protocol will focus on further experiments, followed by the development of merging and reconfiguration mechanisms for service modules and their localization inside the network as well as improvements to the presented mechanisms. For example, service requests in progress should be cached to avoid any unnecessary requests and, thus, to reduce signalling overhead. Beyond it, the implementation of the signalling protocol will be followed up and a larger testbed initially spawning four German universities will be set up.

References

[1] A. Albanese, J. Bloemer, J. Edmonds, M. Luby, and M. Sudan. Priority Encoding Transmission. *IEEE Transactions on Information Theory*, 42(6):1737–1744, November 1996.

[2] D. S. Alexander, W. A. Arbaugh, M. W. Hicks, P. Kakkar, A. D. Keromytis, J. T. Moore, C. A. Gunter, S. M. Nettles, and J. M. Smith. The SwitchWare Active Network Architecture. *IEEE Network*, 12(3):29–36, 1998.

[3] E. Amir, S. McCanne, and H. Zhang. An Application Level Video Gateway. In *Proc. of ACM Multimedia '95*, pages 255–265, San Francisco, California, November 1995.

[4] H. Balakrishnan, S. Seshan, E. Amir, and R. Katz. Improving TCP/IP Performance over Wireless Networks. In *Proc. of 1st ACM Conference on Mobile Computing and Networking*, pages 2–11, Berkeley, California, November 1995.

[5] R. Braden, D. Clark, and S. Shenker. Integrated Services in the Internet Architecture: An Overview. RFC 1633, ISI, June 1994.

[6] M. Calderon, M. Sedano, A. Azcorra, and C. Alonso. Active Network Support for Multicast Applications. *IEEE Network Magazine*, 12(3):46–52, May 1998.

[7] Y. Chawathe, S. Fink, S. McCanne, and E. Brewer. A Proxy Architecture for Reliable Multicast in Heterogeneous Environments. In *Proc. of ACM Multimedia '98*, pages 151–159, Bristol, UK, September 1998.

[8] J. V. Guttag D. J. Wetherall and D. L. Tennenhouse. ANTS: A Toolkit for Building and Dynamically Deploying Network Protocols. In *Proceedings of the IEEE OPENARCH*, pages 117–129, April 1998.

[9] Theodore Faber. ACC: Using Active Networking to Enhance Feedback Congestion Control Mechanisms. *IEEE Network*, 12(3):61–65, 1998.

[10] B. Fenner. Internet Group Management Protocol, Version 2. RFC 2236, IETF, November 1997.

[11] GT-ITM. Georgia Tech Internetwork Topology Models. http://www.cc.gatech.edu/fac/Ellen.Zegura/graphs.html.

[12] Mark Handley. SAP: Session Announcement Protocol. Internet Draft, Internet Engineering Task Force, November 1996. Work in progress.

[13] S. McCanne, V. Jacobson, and M. Vetterli. Receiver-driven Layered Multicast. In *ACM SIGCOMM '96*, pages 1–14, San Francisco, USA, August 1996.

[14] Bernard Metzler. *A Flexible and Efficient System Support for Open Communication Services Platforms.* PhD thesis, Technische Universität Braunschweig, December 1999. in German.

[15] Bernard Metzler, Till Harbaum, Ralph Wittmann, and Martina Zitterbart. AMnet: Heterogeneous Multicast Services based on Active Networking. In *Proc. of the 2nd Workshop on Open Architectures and Network Programming (OPENARCH99)*, pages 98–107, New York, NY, USA, March 1999. IEEE.

[16] V. Ryan, S. Seligman, and R. Lee. Schema for Representing Java Objects in an LDAP Directory. RFC 2713, IETF, October 1999.

[17] A. Speer. A flexible Signalling Protocol for AMnet. Diploma thesis, TU Braunschweig, July 1999. in German.

[18] D. L. Tennenhouse, J. S. Smith, W. D. Sincoskie, D. J. Wetherall, and G. J. Minden. A Survey of Active Network Research. *IEEE Communications Magazine*, 35(1):80–86, January 1997.

[19] UCB/LBNL/VINT. Network Simulator ns-2. http://www-mash.cs.berkeley.edu/ns/.

[20] D. Wetherall, U. Legedza, and J. Guttag. Introducing New Internet Services: Why and How. *IEEE Network Magazine Special Issue on Active and Programmable Networks*, July 1998.

[21] R. Wittmann and M. Zitterbart. AMnet: Active Multicasting Network. In *Proc. of Intern. Conf. on Communications (ICC'98)*, pages 896–900, Atlanta, GA, USA, June 1998. IEEE.

[22] Y. Yemini and S. daSilva. Towards Programmable Networks. In *Proc. of IFIP/IEEE International Workshop on Distributed Systems: Operations and Management*, L'Aquila, Italy, October 1996. http://www.cs.columbia.edu/dcc/netscript/Publications/publications.html.