QoS considerations on IP multicast services

Falko Dressler
University of Erlangen-Nuremberg
Regional Computing Center (RRZE) /
Department of Computer Science IV (Operating Systems)
Martensstr. 1, 91058 Erlangen, Germany
falko.dressler@rrze.uni-erlangen.de
http://bsd.rrze.uni-erlangen.de/~fd/

ABSTRACT

A lot of applications do use IP multicast to communicate between several clients. Typical multicast services are multimedia transmissions from one server to many clients (broadcast) or between several clients (conference). The same infrastructure is also used for non-multimedia applications like NTP (network time protocol) or the delivery of network news.
This paper should summarize current approaches to provide reliability as well as a higher quality of service (QoS). There are at least different ideas and tools to check the functionality and to measure some basic QoS parameters. Also, this paper should provide an overview on what has been done within the last years to help providers as well as end users to get an overview on which QoS they can expect from the network at the moment. Last but not least, there is a look to the current work at the IETF and to activities of some individuals.

I INTRODUCTION

Since the early beginnings of the MBone (Multicast Backbone) a lot of applications have been developed which use the possibilities of IP multicast. Most of these applications use this technique for various multimedia transmissions. Due to the low resource requirements when transmitting one multimedia stream to more than just one client, IP multicast is deployed in most backbone networks as well as in a high number of campus networks. Typical services are video broadcasts or video conferences.

IP multicast like IP does not support to transport realtime data with a specified QoS. It is a best effort transport network, but there are some approaches to put some quality to IP to provide a better quality of service. Also, the network administrators have to find out information about the reliability of their network. For both measurements, there are a lot of tools available and the number of such tools is still growing.

So, one problem is, that you cannot measure the whole internet to apply these data to all possible services (applications).

There are a lot of approaches to measure the reliability of IP multicast networks such as the MRM (Multicast Reachability Monitor, [4]) and the Multicast Beacon. Also, there are ideas to include QoS measurements as well. Such ideas have been introduced as enhancements to the MRM and the Multicast Beacon. Unfortunately, all these tools do not allow the measurement of all required data and also fail to scale for measurements of larger networks due to the idea to measure just the network not to measure data for particular services.

The next sections should give an overview to the some basic principles of IP multicast and IP multicast services as well as to current implementations of measurement tools. Another section shows some ideas on how to provide a higher quality of service. Finally, there is a section which should provide a look to the future.

II TYPICAL SERVICES

The term service means for this work, which application the users do need, which multicast groups this application uses and which end systems do participate. The most interesting types of service are conferences and broadcasts.

A. Conferences

A conference is an good example for the any-to-many transmission model. The basic idea is, that every communication end point works as a sender and as a receiver. This means that every system is participating at the service equally. This scenario is shown in Fig. 1.

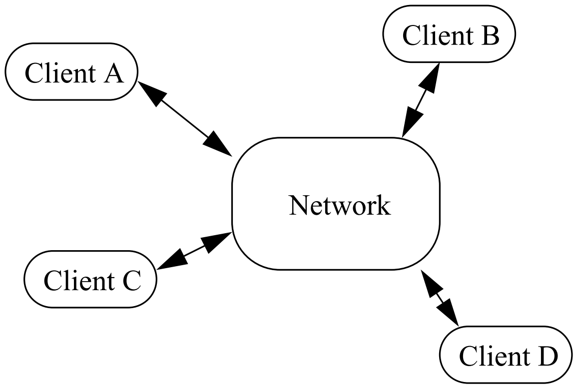

Fig. 1. Multicast services: conference

One example of a multicast conference is the TKBRZL [1]. About 8 to 10 participants from several universities in Bavaria talk about their current problems. They are using audio and video tools and a shared text editor to communicate. So every workstation in fact has up to 9 connections to other machines. If you want to make sure that the service is running fine, you have to check each of the 10 workstations against every other one for a proper multicast connectivity. If you also want to test the actual QoS, you have to check it the same way. Because such conferences appear to last only a short time, you need to do some measurement before they start. First, because the users want to know if the network has the capabilities to start the session and second, because the short time the communication takes does not allow some passive measurements while the service is running.

B. Broadcasts

The other common type of a multicast service is the broadcast. It looks like the conference, if you assume, that only one person is talking and all others are only listening. The model is also known as the one-to-many transmission model. An example of such a scenario is shown in Fig. 2.

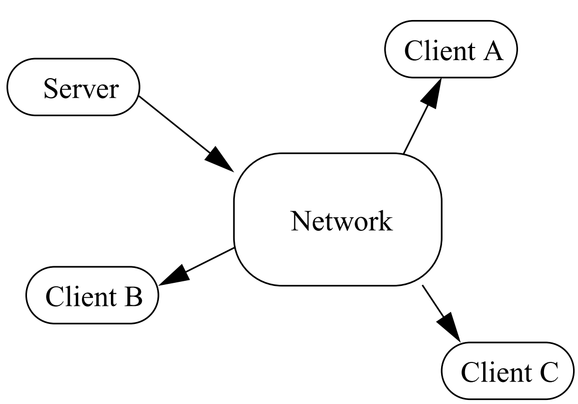

Fig. 2. Multicast services: broadcast

To verify the functionality and the QoS, you have to check the connectivity from the sender (broadcaster) to all the clients (participants, receiver). One example for such a broadcast are the channels from Uni-TV [2], which stream on a near-Video-on-Demand scheme recorded lectures to students of the University of Erlangen-Nuremberg. Since such broadcasts usually last for a longer time, also passive measurements are possible, based on the service itself.

III Reliability and QoS measurements

Currently, there exist two different approaches to measure 'some' reliability of IP multicast networks. Both operating according to the same principle. The network administrator is required to locate proper places within the network to place some monitoring probes which interact using IP multicast. The result is a matrix which allows the administrator to identify parts of the network which do not allow a proper IP multicast transmission.

Due to the principles of IP multicast routing it is required to do a lot of measurements. Especially, this applies to any-to-many transmissions. It is not possible to use the information A reaches B and B reaches C to conclude A reaches C (this statement is also true for IP unicast).

The next paragraphs show two current approaches to measure the reliability and the quality of an existing IP multicast network. Both are capable of providing a lot of information about a particular network. Unfortunately, none of these ideas examines the problem where to place the measurement probes. An approach to solve this problem is shown in chapter IV.

A. MRM - Multicast Reachability Monitor

Formerly, the MRM has been developed as an IETF draft. The IETF decided to stop its work on the MRM because overlaps with other ongoing works such as SNMPv3.

The MRM defines three different processes: the mrm manager, the test sender and the test receiver. The first implementation of the MRM has been done by Cisco Systems who included the functionality into the operation system (IOS) of their routers [8]. A second implementation has been done by a different group for Sun Solaris systems [9]. The following picture (Fig. 3.) should provide an overview to the interoperating processes. The figure shows a simple network consisting of 4 routers and 4 clients. The MRM manager process is running on one of these clients and there are several test sender and test receiver on different clients and router.

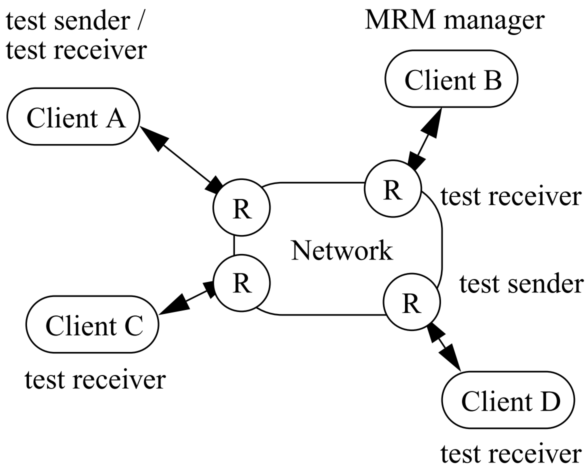

Fig. 3. MRM processes

The MRM is able to create a configurable flow of packets at each test sender which is being sent to the test receivers. Out of the received packets, the test receiver is able to compute results such as the packet loss ratio which provides a good value for the reliability of the network. Finally, the MRM manager process gets informed by the MRM clients and can provide the measured data to the network administrator for further processing.

The MRM uses RTP (Real-time Transport Protocol, [5]) as the transport protocol for the packets sent for measureing. This protocol includes field for sequence numbers and time stamps and is also used by nearly every multimedia networking application.

Using packet flows originated by the test senders as well as other RTP streams originated by any active IP multicast service, the test receiver can also compute values such as the delay and the delay variation (jitter). These parameters are the most important QoS values for multimedia transmissions.

B. Multicast beacon

The multicast beacon [10] is the result of a research project from the NLANR (National Laboratory for Applied Network Research). Currently, there is an implementation in JAVA available which should run on every end system with a JVM (Java virtual machine). The principle functions of the multicast beacon and the MRM are very similar.

The definition of the multicast beacon includes a server which computes the QoS parameters out of measured data and so called beacons which do the real measurement. Every beacon interacts directly with each other by constantly sending IP multicast packets using the RTP protocol to a specified multicast group.

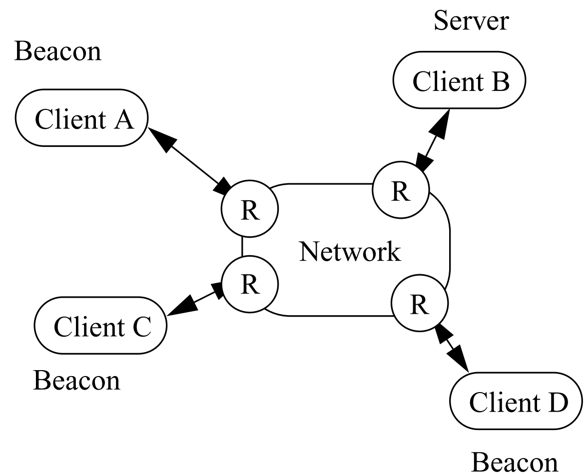

Fig. 4. Multicast beacon

Each beacon reports its measured data, i.e. the results of received packets (beacons) to the server. The server calculates a matrix including each active beacon and allows to access these results via a web gateway.

So the main difference between the MRM and the multicast beacon is the capability of the multicast beacon to provide a direct access to the measured QoS parameters and the possibility of the MRM to distinguish between a test sender and a test receiver which results in a much lower network usage for measurements only. Also, the MRM allows to use other already active RTP streams which reduces the influence on the network also.

IV A MODEL FOR IP MULTICAST SERVICES

As shown in the previous chapters, it is very difficult to find proper places to deploy measurement probes and to use the computed results to provide better reliability and quality to these services.

A. Modeling IP multicast networks and services

One idea to solve this problem is to generate a model of the network including the overlying services. This chapter should give an overview to such a model [11].

The model should be able to include important functions from OSI Layer 1/2 (Link), Layer 3 (IP) and Layer 7 (Application, the services). The primary result of such a model is to find out which parts of a network are required for a particular service. This can be done by attaching various routing algorithms. A first implementation called MRT (Multicast Routing Tool) has been done by Juan Meija, a student at the University of Erlangen-Nuremberg.

The basic objects for the model of the network are shown in Fig. 5.

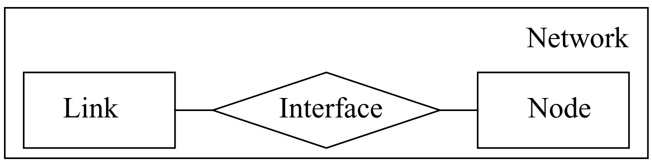

Fig. 5. Basic objects to model a network

Based on these objects, it is possible to generate models of complex networks and to calculate a best route through the network. The objects allow to include parameters such as CPU load, bandwidth of an interface or a loss ratio of a link. Using these values, it is possible to recalculate the routes through the network to find better ways for a particular transmission. The current implementation includes the Dijkstra algorithm for route calculations.

Besides the model of the network, the representations for the services have to be modeled. Each object of class service stands for one multicast transmission which may use more than one multicast group (Fig. 6.).

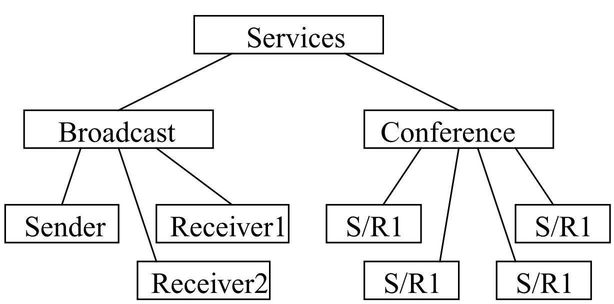

Fig. 6. Example of service objects

Based on the concept of analyzing the most important services within the network first and applying these information to a detailed model of the complete IP multicast network allows to extract the interesting parts of the network. This means, that it enables the network administrator to deploy measurement tools not only based on his wisdom but also on the requirements of the used applications.

Summarizing the already presented objects and their capabilities, the following picture (Fig. 7.) shows the class diagram which has been used to implement the model in JAVA. This implementation has been done by Juan Ceballos-Mejia at the University of Erlangen-Nuremberg [3] as a part of his masters thesis.

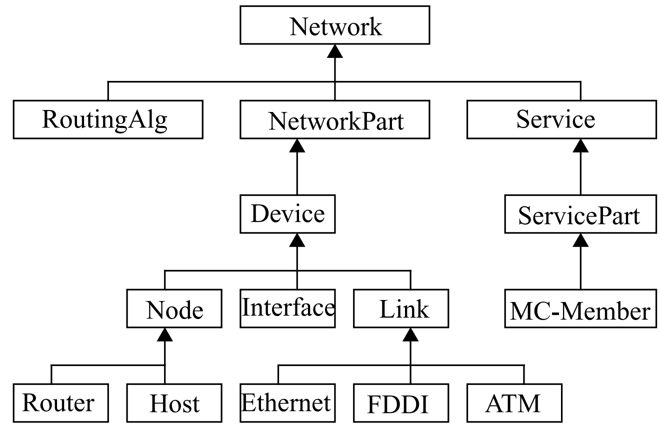

Fig. 7. Object hierarchy

B. Using the model for measurements

In the last chapters an overview of a method to model an IP multicast network including the network itself, the applications / services and the participating hosts has been presented. Also, some real measurement tools have been introduced. The final question is 'How do I use this model to measure reliability and quality of IP multicast services?'.

The current implementation of the model allows you to model your network and check for optimum paths for IP multicast transmissions using the attached routing algorithm. To find the best way, the algorithm uses the constants out of the modeled objects such as bandwidth of a particular interface or the hop count. The following figures show the mechanism. Using the information about the network and the service you can find out which parts of the network are used for this service.

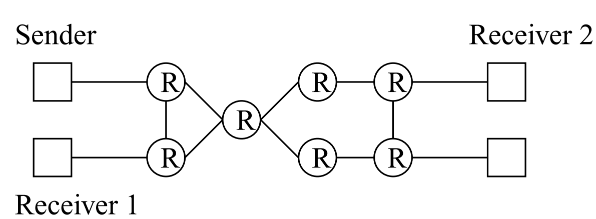

Fig. 8. physical network

Please note, this is only a first step to find all the required parts of your network for a particular service. Our implementation allows already to incorporate dynamic information about the current state of the network. The most interesting values are the state of a node, the state of an interface, the packet loss ratio from an interface and the load of a node. Also, the routing tables of your routers (IP unicast and IP multicast) have to be examined to get closer to the real behavior of your IP network.

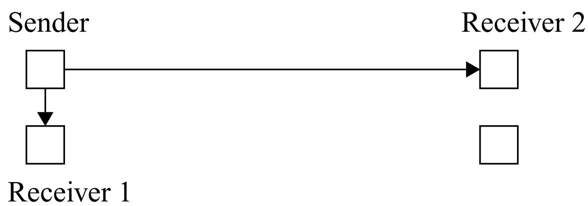

Fig. 9. logical data flow

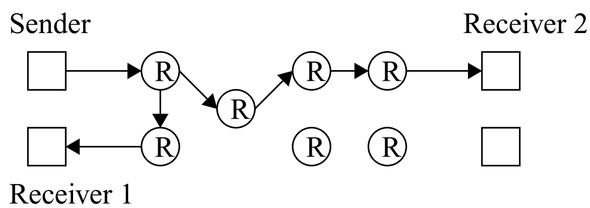

Fig. 10. real packet flow

Based on these dynamic data and the knowledge about the network, the tool allows to find the used components and paths for the current situation and a particular service (Fig. 10.). Using this wisdom, you can adapt your measurement tools to prevent unnecessary tests. This is very important for two reasons. First, you cannot measure the whole internet. You do not have access to all the routers and hosts and you cannot deploy probes nearly everywhere in the internet. Second, also for smaller networks, you need to deploy your measurement tools very carefully since most measures are active tests. This means that you need to create test packets and send them around the network. Using IP multicast, you need to test between every sender and all of its receivers. This does not scale very well, so you have to identify the parts of the network where the probes have to be placed.

V Providing higher reliability and QoS

In the last chapters, methods to measure the current QoS of an IP multicast network have been shown. The next question is, how to use these information to provide a higher reliability and also a higher quality for a particular multicast service of a particular multicast network.

A. Design change

Using the results of the described measurements, it is possible to identify weak points in the network. Mostly, it would be possible to change the design of the network or to manipulate the current routing decisions to modify the load of parts of the network in a way which results in a better utilization of the whole network.

B. Implementation of QoS mechanisms

Also, the results allow the administrator to implement some QoS mechanisms. These mechanisms can be divided into to main groups which are called integrated services and differentiated services. Not all of these technics are available in current routers [12].

The most famous protocol of the integrated services architecture is RSVP (Resource Reservation Protocol, [13]). The idea is to have an end to end communication path with special QoS requirements. Since this is an connection oriented service, every connection is separated into three parts. First, the communication has to be initialized. Within this phase, every router along the path checks for the availability of the required resources. Only if the QoS is available, the communication path is set up. Second, data can be exchanged between the end systems with a guaranteed QoS. Finally, the communication path will be terminated. One problem with RSVP is the high resource requirement. Another one is justified by the inter ISP routing. Every ISP would have to configure its routers in the same way.

The differentiated services approach does not require an end to end signalling. The basic idea is to identify packets of different classes of service. This can be done using the destination IP address, the TCP/UDP port number or the TOS (type of service) field out of the IP header. Finally, it is possible to implement different queueing strategies such as priority queueing or class based queueing. So it became possible to forward packets of different classes based on their priority.

VI Summary

As shown in the paper, currently, there are different tools available to measure some QoS parameters in IP multicast networks. These approaches result in implementations which are under development.

Unfortunately, most of these approaches do not include ideas how to scale for a really large network up to the whole internet. Using the shown model, it becomes possible to reduce the required measurements to the parts of the network which are used by the particular service. So the deployment of already existing measurement tools becomes much easier.

Another result is the requirement of much more research work to provide more efficient and also more powerful approaches for the reliability and quality measurements.

The last chapter introduced some ideas how to improve the quality of different IP multicast services. It is to represent only an introduction, since this field of research activity is very global.

VII References

[1] TKBRZL: http://tkbrzl.bhn.de/index.html

[2] Uni-TV: http://www.university-tv.de/index_en.html

[3] J.-F. Ceballos-Mejia, "Design and implementation of a modeling tool for multicast networks," Master-Thesis, University of Erlangen-Nuremberg, 2000.

[4] K. Almeroth, L. Wei, D. Farinacci, "Multicast reachability monitor (MRM)," IETF draft, January 2001.

[5] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, "RTP: a transport protocol for real-time applications," RFC 1889, January 1996.

[6] J. Postel, "User datagram protocol," RFC 768, August 1980.

[7] K. Almeroth, L. Wei, D. Farinacci, "Multicast reachability monitor (MRM)," IETF draft, January 2001.

[8] Cisco MRM: ftp://ftpeng.cisco.com/ipmulticast.html#MRM & ftp://ftpeng.cisco.com/ipmulticast/mrm/mrm.guide

[9] Kamil Sarac, Kevin Almeroth, "MRM end host implementation," http://steamboat.cs.ucsb.edu/mrm/

[10] Multicast beacon: http://dast.nlanr.net/projects/beacon/

[11] Falko Dressler, "How to Measure Reliability and Quality of IP Multicast Services?" In: Proceedings to 2001 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, PACRIM'01, Victoria, B.C., Canada, August, 26.-28., 2001.

[12] F. Dressler, U. Hilgers, "Routing mit QoS-Eigenschaften unter Linux," In: Echtzeitbetriebssysteme und LINUX: Fachtagung der GI-Fachgruppe 4.4.2 Echtzeitprogrammierung, PEARL, Boppard, 23./24. November 2000 / PEARL 2000, Workshop über Realzeitsysteme. Peter Holleczek (Hrsg.). Springer, Berlin, 2000 (Informatik aktuell).

[13] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification," RFC 2205, September 1997.