A Scalable Environment for Quality of Service Measurements in the Internet

Falko Dressler
University of Erlangen-Nuremberg
Department of Computer Science / Regional Computing Center
Martensstr. 1, 91058 Erlangen, Germany
Phone + 49 9131 85-27802, Fax +49 9131 302941
EMail: dressler@ieee.org

***Abstract* - The goal of this paper is to describe a scalable environment for quality of service (QoS) measurements in the internet. Since the early beginnings of communication networks, measurements of the QoS in the network have been developed and performed for several reasons. First, it was necessary to identify the capabilities of the different types of data networks and to compare them with analytical examinations. Secondly, with the evolving appearance of real-time applications such as multimedia transmissions or remote controls, the demands on the service quality of current network topologies and protocols have rapidly increased. This paper summarizes the available measurement techniques, especially in the case of IP multicast and proposes a measurement environment which builds a base for any kind of quality of service measurements in the internet, either unicast or multicast. The focus of this approach is on the scalability of such measurements which have to coexist with the standard applications in the same networking environment.**

***Keywords* - Quality of Service, Scalability, Network Analysis, Performance Measurements, IP Multicast**

1. Introduction

Quality of service (QoS) measurements in communication networks are required in different representations. Basically, the first measurement tools have been developed to test available network components and to perform checks of the throughput capabilities. With the increasing usage of the same network infrastructure for best effort services as well as for new real-time applications, it became obvious that more detailed quality of service measurements are required. It is not only a question of providing or even guaranteeing some minimum transmission quality, it is on the service providers and on the end users to employ performance tests of the application requirements and to check the conformance of guaranteed service levels.

The focus of this paper is to describe the basic scalability issues of most QoS measurements. Additionally, an approach is presented which is built on new concepts providing a high scalability. The proposed measurement environment was designed to work as a common framework for many kinds of quality of service determinations in IP networks. Therefore, it was important to build a measurement environment operating in multicast networks as well as for the unicast case.

The paper is organized as follows: a historical outlook to quality of service measurements is given in section 2 and common scalability considerations are discussed in section 3. An overview to existing measurement tools working in multicast environments is provided in section 4. The main section of this paper, section 5, introduces a scalable environment for quality of service measurements in the internet including a few sample measurement results and some usage scenarios. A conclusion (section 6) summarizes the work.

2. QoS Measurements - A Brief History

Measurements of the available quality of service have a long history. Since the early beginnings of communication networks, the capabilities of these networks have been questioned. First, as the term quality of service was not yet introduced, only parameters such as the available and the maximum bandwidth and the loss or error ratio have been investigated into. Even in "predictable" packet oriented networks such as X.25 and ATM (asynchronous transfer mode), the demand increased to create measurement environments to verify the promises of the vendors and the service providers.

A second era of QoS measurements started with the evolving requirement to measure a more widespread spectrum of quality of service parameters. Besides the throughput and the packet loss ratio, there was an increasing demand on measurements of the delay and the delay variation or "jitter".

Another step forward was the introduction of the one-way delay as one of the most important values for real-time multimedia communication. It is still one of the most interesting parameters but it is very difficult to measure due to the high synchronization requirements of the clocks of the involved measurement nodes.

Nowadays, IP is the protocol which is used for nearly every kind of data transmission - even if there are much more useful solutions. Therefore, the analysis and the test of the behavior of the IP networks in terms of the available quality of service become important.

The basic problems of the QoS measurements have not changed over the time even if the networks became much faster and the capabilities have grown to increase or even to guarantee some transmission quality. Formerly and today, one has to deal with the same difficulties. There are imprecise hardware and software components, losses introduced by the operating system and by inefficient programming, the problem of incorrect local clocks, and, last but not least, effects due to a low synchronization accuracy between the clocks of the measurement nodes.

Other challenges appeared with the evolution of multicast [8]. It became even more difficult to establish concepts for quality of service measurements in multicast environments. Some approaches which build the basis of the proposed measurement environment are discussed in section 4.

3. Scalability Considerations

Scalability is not only an issue of network protocols and the applications using the communication network. Most of the QoS measurements in multicast networks suffer from inefficient concepts and a low scalability. The most important scalability issues are discussed in the following.

3.1 Parallelism of single Measurements

A typical situation is the occurrence of multiple measurements which simultaneously take place in the same network environment. The problem is described in Fig 1.

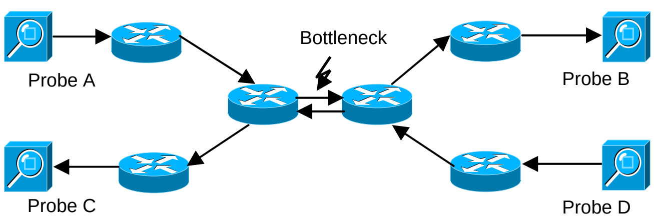

Fig 1. Problem of multiple simultaneous measurements. The probes A and B and the probes C and D are simultaneously initiating and analyzing different measurement data streams.

Even if there are still free resources in the network while multiple measurement data streams are active, the behavior of the network is not directly comparable to the one during a completely unused situation. The CPUs of the routers are loaded, the queues are filled and so on. Therefore, attention must be paid to the possibility of interfering measurements or, at least, the effects must be analyzed in order to correctly explain the situation. Nevertheless, it might be a good approach to create such a situation in order to test the network behavior during overload situations in comparison with idle times.

3.2 Impact on the Networks Behavior

Another very important issue is the influence of the measurement on the normal operation of the network. Especially, this is the case if the quality of high bandwidth data streams is examined. In most cases, it is not possible to build a testbed for all the required examinations and measurements. Additionally, there is an increasing demand to employ measurements in the network to predict the behavior and the quality of scheduled applications. Such tests will always run in the production network.

To solve this problem active and passive measurements are distinguished as shown in Fig 2. On the upper side, a probe is passively sniffing the traffic between two sites. In the case of multimedia transmissions, some QoS parameters can easily be determined using this gathered information due to the capabilities of the used transport protocol. RTP (real-time transport protocol, [13]) includes information such as time stamps and sequence numbers in the header of each packet allowing the receiving side, or in this case the measurement probe, to analyze the current QoS in the network between the sender and the receiver.

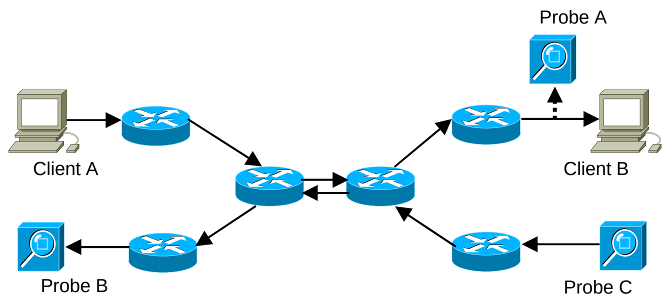

Fig 2. Active vs. passive measurement. There are two simultaneous data streams. The clients A and C are communicating using the same network environment which is used by the probes B and C for some measurements. Probe A is sniffing the traffic from client A to B in order to analyze it and to compute several QoS characteristics.

Obviously, the passive methods do not interfere with the normal operation of the network and, therefore, build the optimum basis for QoS measurements in production networks. Nevertheless, it is not very common to have an active transmission with the required service behavior available at the right time and at the right place. Therefore, active measurements must be employed as shown on the lower part of the figure.

3.3 Packet Explosion in Multicast Environments

The last issue on scalability questions to be discussed here is the multiplication effect in multicast environments. A message sent to a number of destinations requesting the receivers to answer it results in a number of response messages equal to the number of receivers. If the response is also sent via multicast, every participant has to receive and to analyze a number of packets equal to the product of senders and receivers.

Especially in measurement environments, the packet explosion problem occurs very often. Therefore, mechanisms have to be developed and implemented which deal with this problem.

4. Concepts for QoS Measurements, Related Work

The primary goal of this section is to provide an overview to the basic concepts of quality of service measurements. In the IP based internet, the IETF (internet engineering task force) established several working groups, e.g. the IPPM

WG (IP performance metrics working group), to work on this issue. Additionally, numerous groups are working on environments and tools for IP performance measurements. Most of them are only focused on unicast connections. Additionally, there is much work in progress on the synchronization of the clocks of the network nodes. Unfortunately, only a few people include multicast connections into their research and there is no common measurement environment for all kinds of quality of service measurements in the internet.

The goal of this section is to summarize some approaches which focus on QoS measurements in multicast networks due to the special scalability issues in multicast environments.

4.1 IP Performance Metrics

A first step towards a framework for defining internet performance metrics has been provided by Paxson [11]. He distinguishes between analytically and empirically specified metrics. Another approach by Awerbuch et al. [4] describes a cost-sensitive analysis of communication protocols. Several documents have been written by the IPPM WG of the IETF in order to define a generally accepted framework for IP performance measurements [12, 14]. These definitions can be divided into functional requirements: definitions of packet formats and common recommendations.

4.2 Multicast Reachability Monitor (MRM)

The MRM [3], started as an IETF draft [2], has been developed to allow a centralized reachability management based on probes located all over the multicast network. Even if the work on the MRM was cancelled, it is a very interesting tool, which was a starting point for further developments such as the MQM described later.

The MRM defines three different processes, the MRM manager, the test sender and the test receiver. Controlled by the manager, the test sender can insert a packet flow into the multicast network. Using the received packets, the test receivers are able to compute measurement results such as the packet loss ratio, which also provides a good value for the reliability of the network. The MRM clients send the results to the manager process. Therefore, the latter can provide the measured data to the network administrator for further processing [1].

4.3 Multicast beacon

The multicast beacon [5] is the result of a research project at the NLANR (National Laboratory for Applied Network Research). The principles of the multicast beacon and the MRM are very similar. The definition of the multicast beacon includes a server computing the QoS parameters out of measurement results and clients, so called beacons, originating and receiving measurement packets. Every beacon interacts directly with each other by constantly sending IP multicast packets using the RTP protocol to an administratively configured multicast group. Each beacon client reports its measured data, i.e. the results of received packets (beacons), to the server. The server calculates a matrix showing the current QoS between each active client and allows accessing these results via a web gateway.

The main differences between the MRM and the multicast beacon are the capability of the multicast beacon to provide a direct access to the measurement results and the wider range of QoS measurements (packet loss ratio, delay, and jitter).

4.4 Multicast Quality Monitor (MQM)

An approach initiated by Dressler is the MQM [6, 7]. It allows one to measure several quality of service parameters in an IP multicast network. The complete functionality of the MQM was split into two parts: the MQM ping mechanism and the MQM beacon mechanism. The first one was designed for availability tests in a large scale multicast network including the measurement of various QoS parameters such as the one-way delay or the packet loss ratio. The latter one was build for simulations of real application traffic as well as for passive measurements using currently occurring multicast transmissions.

The basis mechanism of the new multicast ping introduced by the MQM is described in Fig 3. It is suggested to repeat the ping process every 60 sec in order to refresh the state information in the network components (typical timeout values are 180 sec). The issues concerning the scalability of this mechanism are discussed in section V.

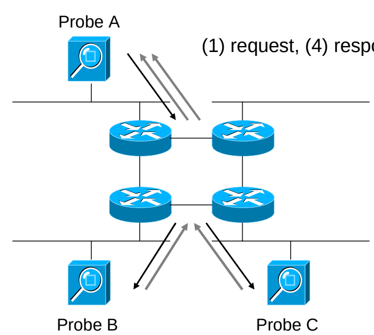

Fig 3. MQM ping mechanism. To test the reachability, a MQM ping request packet is sent (1) to all the others via a well-known multicast address. The other probes receive this request (2) and respond (3) by sending a MQM ping response towards the originator which receives and analyzes the response (4).

The second measurement mechanism, the MQM beacon mechanism is based on the analysis of received RTP packets. Such packet can be intercepted from an active multimedia transmission (passive measurement) or by simulating such a packet stream by itself (active measurement). Typically, the active measurements will have a noticeable impact on the network itself. Especially, if services like high quality TV broadcast with high bandwidth video streams are simulated.

5. A Scalable Measurement Environment

Based on the concepts of the multicast quality monitor, a common environment for any kind of quality of service measurements in IP networks was created. Within the following subsections, the basic concepts are presented and the scalability issues are discussed. Some measurement examples summarize this section.

The basic concepts of the MQM can be directly applied to any kind of QoS measurement in the internet regardless if unicast or multicast is involved.

5.1 Scalability Considerations

As already mentioned the scalability is always a critical issue in a measurement environment, therefore, the main focus of the presented approach is laid on it. The MQM includes concepts to prevent the typical scalability problems such as multiple uncontrolled concurrent measurements, the usage of active multimedia transmissions as a source for the measurements, and the packet explosion problem in multicast environments.

MQM ping mechanism. To prevent the packet explosion problem during the multicast ping process, this mechanism was designed with scalability in mind. The concept is shown in Fig 4. Using the information out of all the received packets, whether request messages or responses, the complete reachability information can be calculated:

Connection P_1-P_2: Request P_1 + Response P_2; ...

Connection P_2-P_3: Response P_2 + Response P_3; ...

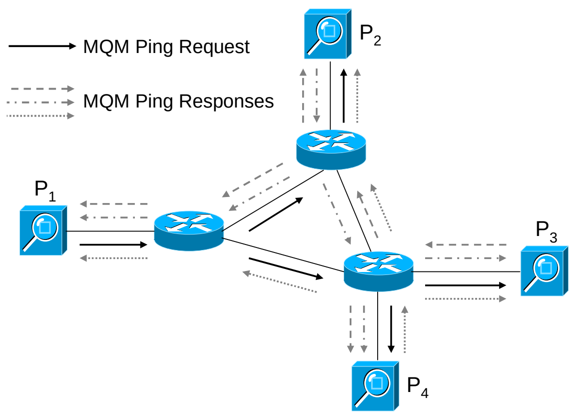

Fig 4. MQM ping process. A single ping packet and its replies is enough to calculate all the reachability information. In this example, P_1 is sending a request packet. All other probes answer by sending a response. A complete connectivity matrix can be calculated using the received messages.

It becomes obvious that two request packets in an interval of 60 sec (see section 4) is an optimum approach to prevent the network being congested by too many response packets. The implementation of this mechanism is straightforward. A probe is started at a random time and sends a request packet. Then it prevents to send any further request until it received less than two requests in the last interval.

MQM beacon mechanism. The MQM introduced a special mechanism to control the RTP based measurements, the MQM beacon mechanism. A RTP sender process and a RTP/RTCP receiver processes can be started using a centralized control instance which is able to verify that no, or at least no interfering, parallel measurements are started. All the standard MQM ping receivers are required to listen for this special type of MQM messages. A single beacon packet might contain a number of single beacon messages. Each contains a destination and a command sequence (see Fig 5).

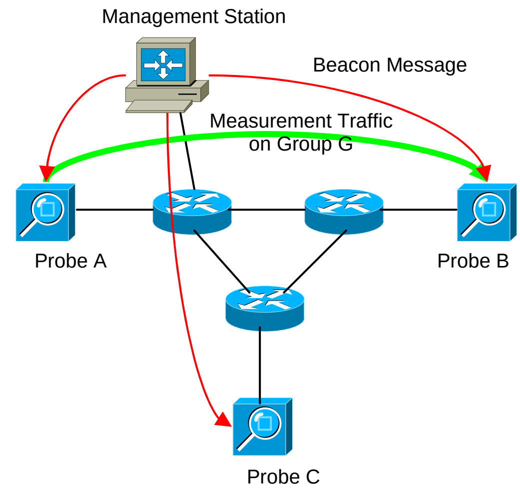

Fig 5. Beacon mechanism. A centralized management station is sending a beacon message to all the probes in the network containing two commands: (1) probe A is required to start a RTP sender process and probe B must initiate a RTP receiver. Shown is the state of the network after the reception of the commands. The management station continues to send the beacon messages until their time to live is reached.

For example, a number of probes can be introduced to start RTP receivers listening on a particular IP address (either unicast or multicast). The beacon message is continuously sent until the deadline of this particular beacon is reached. The RTP sender/receiver process on the probes is stopped either if a beacon message is received including a stop command, if the beacon deadline is reached, or if no new beacon message was received in the last time slice.

5.2 Measurement Results

Following the presentation of the basic concepts and the scalability considerations, the results of two sample measurements are shown in Fig 6 and Fig 2. A prototypical implementation on UNIX systems has been used to gather the results.

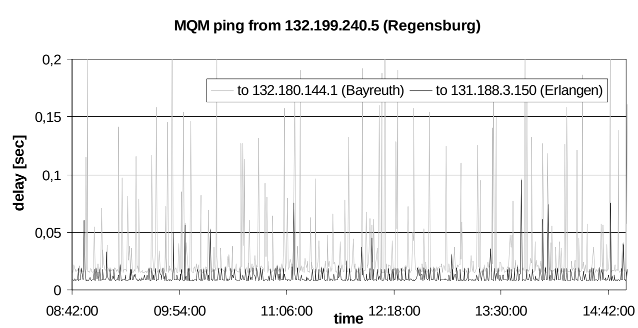

Fig 6. MQM ping example. A multicast ping is initiated at a host in Regensburg (R) and received and analyzed in Erlangen (E) and Bayreuth (B), respectively. The

connection from R to E is only slightly used and allows at least a throughput of 100 Mbps. The connection to B is using a 10 Mbps shared ethernet and a heavily loaded router. Therefore, the delay is a little higher and the variation of the delay, the high peaks in this figure, became very large.

The first example which is shown in Fig. 11. is the result of a multicast ping measurement. A single host was transmitting MQM ping requests to a preconfigured multicast group. Two other nodes were running MQM receiver processes and, therefore, receiving and processing the requests.

In the second example (Fig. 12.), an active multimedia application, a multicast video transmission, was used to passively measure the quality of service from the source of this packet stream towards the monitoring probes.

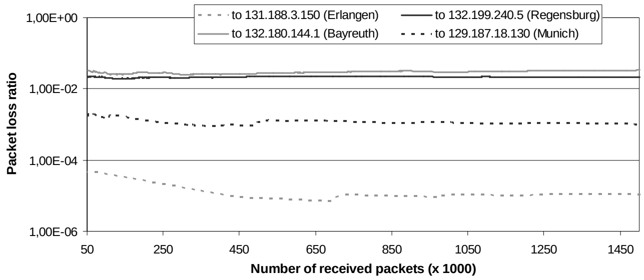

Fig 7. MQM RTP example. A second test was done using a RTP stream initiated in Erlangen (E) which was received and analyzed in Erlangen (E), Munich (M), Regensburg (R), and Bayreuth (B). The presented packet loss ratio, presented in a logarithmically scaled diagram, allows a characterization of the available transmission quality. From E to E and to M, the quality is very high (a loss ratio of less than 0.001). Admittedly, the connections from E to R and to B are less useful for a real-time multimedia transmission (a loss ratio of about 0.25).

The measurement results are conform with the expectations and they are only included to show the capabilities of the measurement environment.

5.3 Application Scenarios

Two typical usage scenarios of the common measurement environment summarize this section:

- selection of multimedia servers
- prediction of the service quality for forthcoming applications such as video conferences

A common approach to increase the availability of server systems is to deploy multiple instances with a replicated content [10]. Such mechanisms are also used to allow a load-sharing between all the instances. In the case of multimedia servers, the load of each server is not the primary selection criterion. New approaches include the current quality of service in the network from the server towards the client [9]. The test of the transmission quality can be easily performed using the proposed measurement environment.

Another application scenario is the check for available resources for forthcoming video conferences. It was shown that the acceptance of such tools strongly decreases with a low service quality. In order to predict the quality of service which can be achieved in a forthcoming session, the behavior of the session can be simulated for a short period of time using the multicast quality monitor. The measurement results allow the participants to decide whether to start or to reject a scheduled video conference based on the expected service behavior.

6. Conclusions

The main goal of this paper, the demonstration of a new scalable measurement environment for quality of service tests in the internet has been achieved. It was shown that the scalability is very important in current measurement environments. Especially this is the case for multicast networks. Due to network failures network partitioning may happen, which might be contra productive to high scalability solutions. Therefore, the robustness of a measurement approach is an issue. In this context several issues have to be considered.

The proposed approach was designed with these scalability questions in mind and, therefore, a framework was presented which not only allows one to initiate many kinds of QoS measurements such as the connectivity, the delay, the jitter, and the packet loss ratio, but also achieves the required data with a low impact on the network itself.

We already have a working prototype of the measurement tool active, which made it possible to achieve some first measurement results. The given sample results and the discussed application scenarios have proven the functionality and the applicability of the proposal.

The next steps toward a globally usable tool are to implement the processes on the mostly used operating systems and to establish a cooperating internet-wide heterogeneous measurement environment.

References

[1] K. Almeroth and L. Wei, "Justification for and use of the Multicast Routing Monitor (MRM) Protocol," draft-ietf-mboned-mrm-use-00.txt, February 1999.

[2] K. Almeroth, L. Wei, and D. Farinacci, "Multicast Reachability Monitor (MRM)," draft-ietf-mboned-mrm-01.txt, July 2000.

[3] K. C. Almeroth, K. Sarac, and L. Wei, "Supporting Multicast management Using the Multicast Reachability Monitor (MRM) Protocol," UCSB, Technical Report TR 2000-26, May 2000.

[4] B. Awerbuch, A. Baratz, and D. Peleg, "Cost-sensitive analysis of communication protocols," Proceedings of the ninth annual ACM symposium on Principles of distributed computing, Quebec City , Quebec , Canada, 1990, pp. 177-187.

[5] K. Chen, M. Kutzko, and T. Rimovsky, "Multicast Beacon Server v0.8.X (Perl)," January 2002. (http://dast.nlanr.net/projects/beacon)

[6] F. Dressler, "An Approach for QoS Measurements in IP Multicast Networks, MQM - Multicast

Quality Monitor," Proceedings of International Network Conference (INC 2002), Plymouth, UK, July 2002.

[7] F. Dressler, "MQM - Multicast Quality Monitor," Proceedings of 10th International Conference on Telecommunication Systems, Modeling and Analysis (ICSTM10), vol. 2, Monterey, CA, USA, October 2002, pp. 671-678.

[8] F. Dressler, "QoS considerations on IP multicast services," Proceedings of International Conference on Advances in Infrastructure for Electronic Business, Education, Science, and Medicine on the Internet (SSGRR 2002w), L'Aquila, Italy, January 2002.

[9] F. Dressler, "Considerations on Selection Criteria for Sources of Multimedia Traffic," Proceedings of 7th World Multiconference on Systematics, Cybernetics and Informatics (SCI 2003), vol. III, Orlando, Florida, USA, July 2003, pp. 34-38.

[10] M. Guo, M. H. Ammar, and E. W. Zegura, "Selecting among replicated batching video-on-demand servers," Proceedings of 12th International Workshop on Network and Operating Systems Support for Digital Audio and Video, Miami, Florida, USA, 2002, pp. 155-163.

[11] V. Paxson, "Towards a Framework for Defining Internet Performance Metrics," Proceedings of INET '96, 1996.

[12] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis, "Framework for IP Performance Metrics," RFC 2330, May 1998.

[13] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacabson, "RTP: A Transport Protocol for Real-Time Applications," RFC 1889, January 1996.

[14] S. Shalunov and B. Teitelbaum, "A One-way Active Measurement Protocol Requirements," draft-ietf-ippm-owdp-reqs-03.txt, August 2002.