

A Study of Packet Losses in the EuQoS Network ^{*}

Loránd Jakab, René Serral-Gracià, and Jordi Domingo-Pascual

Universitat Politècnica de Catalunya
{ljakab,rserral,jordid}@ac.upc.edu

Abstract. As Internet usage grows, more efforts are put into analyzing its internal performance, usually such analysis comes through simulation using different models. While simulation can provide a good approximation of network behavior, modeling such a complex network as the Internet is very difficult if not impossible. This paper studies the network's performance from an experimental point of view using the EuQoS project's overlay network as a testbed.

In the framework of the EuQoS project, many performance tests have been done for proving the reliability of the data transmission. The tests show some rough edges which need further analysis, the most important being random packet losses in UDP flows, and a great amount of out of order packets. This paper focuses on the study of such packet losses, searching for their causes, and more importantly, to show their effects on real-time traffic such as VoIP. As a basis for comparison, the paper also uses TCP traffic to relate the performance of bulk data transfer versus the sustained rate of UDP/RTP flows used for real-time applications.

To achieve this goal, several applications are used to generate and capture such traffic and measure its behavior at network level.

1 Introduction

The European academic network is extensively used by several research projects of the many connected universities which rely on its availability and good performance. One such project is EuQoS [1], which operates an overlay network using the resources offered by this academic network. This overlay network, from now on referred to as the EuQoS network, consists of different access technology specific testbeds, interconnected by each country's NREN, which in turn connects to the Géant paneuropean network. Access technologies include Ethernet, WiFi, xDSL, UMTS, while the testbeds (a total of 12) are spread across different European countries, such as France, Switzerland, Italy, Poland, Portugal and Spain.

The project's main goal is to develop the technologies necessary for ISPs to operate a QoS enabled network infrastructure over the Internet. As such, the project needs to develop tools to monitor and test parameters that define QoS. For this reason, a number of different tests were performed in the initial phase of the project to assess the raw performance of the network. These tests revealed some strange behavior of the network, apparently random packet losses, as well as a considerable amount of out of order packets in UDP flows.

The main contribution of this paper is to perform a deep analysis of the packet losses on the academic network in an experimental environment. It will focus on the study of the packet loss behavior, its possible relation with the available bandwidth, and the effect it may have on

^{*} This work was partially funded by IST under contract 6FP-004503 (*IST-EuQoS*), MCyT (Spanish Ministry of Science and Technology) under contract FEDER-TIC2002-04531-C04-02 and the CIRIT (Catalan Research Council) under contract 2001-SGR00226.

interactive multimedia applications. Using artificially generated traffic resembling multimedia applications it will analyze the traces, looking for packet loss patterns.

The rest of the paper is structured as follows: the next section will present a short overview of the testbed used to perform the measurements and the details of the methodology of the measurements performed. After that we present the results of the measurements and section 4 will conclude the paper.

2 Measurements

This section describes the testbed and the measurements performed on the EuQoS network. First we'll give an overview of the testbed in which the measurements were performed and then detail the methodology of the performed tests.

2.1 Testbed Description

The testbed used for the measurements presented in this paper is that of the EuQoS project, consisting of the access network specific testbeds of the partners which are interconnected through NRENs and Géant. Figure 1 presents a simplified version of this network, by reducing the number of partners. It is organized as an overlay network, where each connected partner is an Autonomous System of its own, with an A class reserved address space. The border routers of the partners establish GRE tunnels with the border routers of all other partners, establishing this way a full mesh type network between the ASs. This is necessary for the first phase of the project. However, as each testbed is a distinct AS, the BGP protocol is deployed as well for the second phase of the project, when multi-AS links will also be tested.

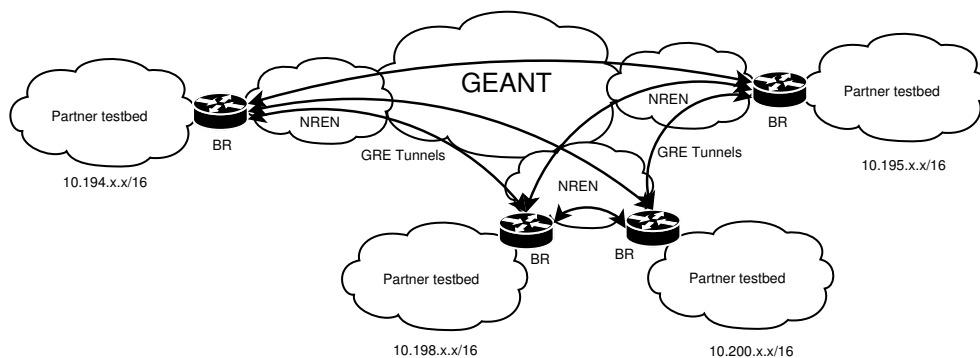


Fig. 1. Simplified EuQoS Network

The overlay network hides completely the NREN and Géant network nodes and these are not directly addressable from within the EuQoS network. But this part of the network is exactly the one under test in this paper.

One way delay measurements are based on the difference between the transmission and reception timestamps of the packets. This requires a good synchronization of the measurement endpoints, especially on today's high-speed links. In order to provide the precision necessary for the measurements, some of the partners have GPS time sources, which then offer synchronization to other machines through NTP. To improve synchronization between machines, each measurement endpoint machine is synchronized to more than one source, so with time the time convergence of the whole network is very good, if no serious disruptions in network service or uptime occur.

A typical LAN testbed is presented in Figure 2.




Fig. 2. Example LAN testbed

From the testbed point of view we define two types of measurements:

Site-to-site measurements: the endpoints involved are the outermost measurement points of each partner, which means the border router (if software based, e.g. Linux/BSD) or a machine directly connected to the border router (hardware, e.g. Cisco/Juniper) through a high-speed link. In the example testbed from Figure 2 that would be either one of the machines 10.198.0.2 or 10.198.0.3.

End-to-end measurements: the endpoints involved are located inside the access network specific testbed of each partner, so the effect of each access network technology is considered. In the example testbed from Figure 2 that would be either one of the machines 10.198.0.10 or 10.198.0.11.

This paper deals with site-to-site measurements only, because it would be very hard to control the environment of the internal testbeds of the other partners and since they can have several applications with different activities on their networks they could interfere with the measurements. On the other hand some access network types have a relatively limited bandwidth and the purpose of the paper is to highlight some issues observed on the academic network. They also have activity outside their testbed, but rarely on the overlay network. All tests are scheduled for

specific timeframes to ensure that tests will not overlap in time. The only constant background traffic on the overlay network is that of the NTP protocol, which is negligible compared to the test traffic, described in the following. Some background traffic is also produced by the *ssh* connections needed to perform the tests, these however occur only before the test is started and do not interfere with the test.

2.2 Active Measurements

Most of the testing was done by means of active measurements, i.e. by injecting test traffic onto the network and analyzing the response. For this purpose the NetMeter tool [2], developed inside the Advanced Broadband Communication Center of UPC, was selected. This tool is a graphical frontend combining several open source network measurement tools, from which the most important is MGen [3], developed by the U.S. Naval Research Laboratory.

We performed two major types of active measurements: UDP tests, which simulate multimedia traffic and TCP bandwidth tests, which try to determine the available network bandwidth between each combination of partner/testbed pair.

UDP Tests. The UDP tests were performed in EuQoS in order to simulate different profiles of multimedia traffic and see the behavior of the network under test in terms of QoS parameters. The measurements targeted all well-known parameters, however this paper will consider only the results related to packet loss.

In order to perform active measurements we injected three different traffic profiles into the network, specifically site-to-site tests were performed with all possible combinations of partner pairs. As previously mentioned, the NetMeter tool was used for this purpose, which started the sender and receiver processes on the two endpoint machines, and generated constant bitrate, periodic traffic. The two most important parameters that MGen takes to generate the traffic are the UDP payload size in bytes, and the number of generated packets per second. Table 1 presents the three traffic profiles used for the UDP tests. The bandwidth column of the table considers the UDP header and IP header size in addition to the UDP payload, but not the layer 2 overhead, neither the overhead produced by the overlay network tunnels. Each test ran for 10 minutes.

Table 1. Traffic profiles

Profile	Rate	Payload Size	Bandwidth	Packets/test
	<i>[pps]</i>	<i>[bytes]</i>	<i>[bps]</i>	
VoIP	20	60	14.08 K	12000
UPD1	96	1420	1.11 M	57600
UPD2	897	160	1.34 M	538200

TCP Tests. One important goal of the TCP tests was to estimate the available bandwidth between partners to see if the minimum requirements of the EuQoS project are satisfied. The UDP1 and UDP2 tests had an approximate bandwidth of 1Mbps, therefore it was necessary to verify if this bandwidth was actually available, so that packet losses would not be the result of lack of bandwidth. The length of each test between the pairs was 10 minutes, in order to obtain reliable results, avoiding possible “noise” produced by short bursts of traffic. Since these tests were announced, it is expected that no other traffic was present on the EuQoS network except NTP, and some sporadic *ssh* connections, which don’t bias the results due to small traffic.

To perform TCP tests the tool Netperf [4] was chosen. This tool is able to generate a TCP stream that will fill up the available bandwidth between the test endpoints, thus measuring it. It consists of a server process called *netserver* listening on port 12865 for incoming connections and the client application (*netperf*) which will set up a connection, generate packets and calculate the bandwidth.

The tests were automated using a small Perl script to check for netserver availability, generate the test matrix (tests were performed in both directions for each combination pair), then issue the test on the source endpoint (through *ssh*) and finally to get the result.

The *netperf* command line specified a message size of 1420 bytes. This value was chosen in order to have messages fill a whole packet, yet not to be fragmented.

3 Results

This section will present what is the main contribution of the paper, an overview and analysis of the results obtained with the performed tests. First we will comment on the obtained TCP bandwidth matrix between partners and then analyze the losses observed during the UDP tests.

3.1 TCP Bandwidth Tests

As explained before, the role of the TCP tests was to determine the average available bandwidth between each partner’s edge router (from the EuQoS overlay network’s point of view).

The results presented in the following pertain to a successful test between 10 of 12 partners during out-of-office hours. All possible pair combinations yielded 90 tests, each of which lasted 10 minutes. For this reason the total time of the tests was rather long (a little more than 15 hours, considering inter-test delays) so the timeframe chosen was between 16:30 and 7:30.

The results were as follows:

- One of the partners had a maximum 0.24 Mbps bandwidth with all partners in both directions. This was normal behavior since they only have a 256 kbps internet connection.
- Between two Polish partners the throughput was extremely high: 35 Mbps in one direction and 42 Mbps in the other. Since the traffic doesn’t leave the MAN, this is normal behavior.
- after disregarding the above values from computations, here are the statistical parameters of the throughput:
- Most of the links are acceptably symmetrical but it’s very typical to have about doubled throughput in one direction compared to the other.
- There is one very asymmetrical link: 4.55 Mbps one way, 18.95 Mbps the other

Table 2. TCP test results

Parameter	Value [Mbps]
Minimum	0.41
Average	4.68
Maximum	18.95
Std. deviation	2.76

It must be noted that the average throughput of links is well below the 8 Mbps expectation of the EuQoS project.

3.2 UDP Tests

In the following paragraphs we will go through the results obtained through the three types of UDP tests – VoIP, UDP1, UDP2 – performed in the EuQoS network.

We must point out that the results of these tests were obtained during a larger timeframe, not one day as with the TCP tests. Setting up scripts for the UDP tests was more error prone, and some tests had to be repeated.

As the main goal of these tests was to verify the reliability of the overlay network’s links, some irregularities were observed, which shouldn’t be present on a production network. One such example was a consistent high packet loss rate of one partner with all others. As further investigation revealed, the cause was a half-duplex connection on an Ethernet link consisting of a few cascaded switches. Site-to-site tests shouldn’t involve the testbeds, however this machine was test endpoint outside of the testbed, connected directly over the Ethernet link to the border router, so test requirements were met.

VoIP Overall, the VoIP tests performed very good from the packet loss point of view, experiencing only very reduced packet losses (0 or 1 packets) and only in a few cases, mainly with partners that had problems, such as the case presented in the previous paragraph. Table 3 presents data separately for the partners with connection/tested problems and partners without (obvious) problems.

Table 3. Packet loss statistics

	VoIP		UDP1		UDP2	
	<i>Avg</i>	<i>Max</i>	<i>Avg</i>	<i>Max</i>	<i>Avg</i>	<i>Max</i>
Best case	1.095	34	2.86	27	24.856	292
Worst case	59.346	1330	53.96	294	1019	1235

UDP1 In general the results were good, with several tests without any packet loss. These tests were performed after the VoIP test, and one partner solved his issues, so only one partner remained with high packet loss ratio. But the test revealed another partner with a problem only on the outgoing link, which was not obvious from the VoIP tests. Refer to Table 3 for the details. Note however that partners classified as normal/problem are different between the different measurement sets, depending on the results.

UDP2 This test emphasized to the problem of the previously mentioned outgoing link, with constant losses of over 1000 packets/test. The rest of the partners had behavior consistent with the previous results, with a slightly increased average, as expected.

3.3 Packet Loss Patterns

In order to identify repeating or any type of recognizable patterns, the traces with high packet loss rates were analyzed further. We are interested in correlation of packet loss with out of order packets and their burstiness.

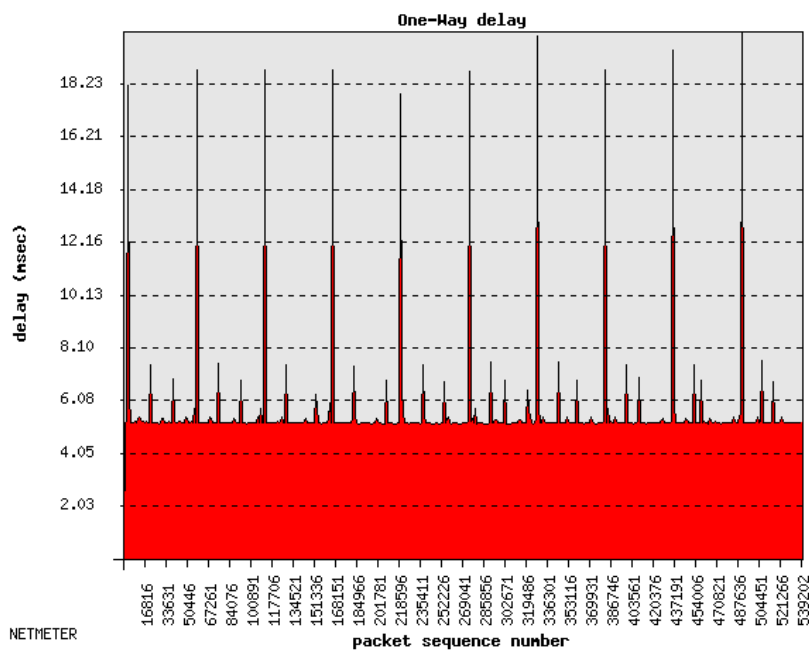


Fig. 3. One way delay

Taking the problem partner from the UDP2 tests with an average of over 1000 packets lost/test a clear pattern was observed: about every 54000 packets a burst of 100 consecutive packets were lost, representing a period of about 112 ms. Figure 3 shows the one way delay

graph of the whole test, while Figure 4 the first part of the graph magnified, for the lost packets to become visible. The spikes on the first graph come right before each group burst of lost packets, and they are also periodic, at about 54000 packets (every 60 seconds). The spikes represent a group of delayed packets, possibly due to buffering: after the buffers are full, we start losing packets. Further work is needed to determine the exact cause of this behavior.

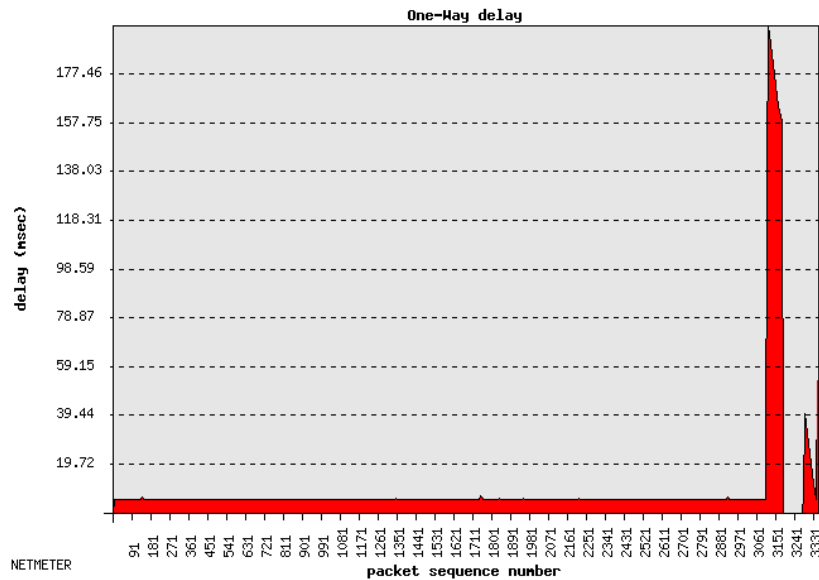


Fig. 4. One way delay magnified

These tests had no out of order packets except with one partner with whom the ratio was 29%, but it is known to us that they have load balancing and had a high amount of out of order packets in all tests. This is consistent with our previous conclusion.

Analyzing the UDP1 traces with uncommonly high number of lost packets it was revealed that all losses consisted of a single burst of continuous packets and in some cases some isolated lost packets, which fit into normal random loss profile.

As a general conclusion for UDP1 tests, if packet loss was less than 10 packets, the distribution was totally random, and for packet loss over 100 packets there is always a long burst of consecutive packets, combined eventually with some isolated packet loss. The interval in-between doesn't have such a deterministic behavior and varies from test to test.

As the main problem of packet losses in this environment are buffer overflows, an easy solution to decrease the number of losses is to grant some degree of priority of the UDP/RTP flows over TCP flows along the path. Anyway, to have a clear path towards packet loss reduction, a deeper analysis of the path of the packets is in order. This way, it should be possible to point the exact reason of such losses: either because of excessive UDP traffic, or because the deprecated treatment of UDP (or IP in IP) flows.

4 Conclusions

The paper presented a study of packet losses on the EuQoS overlay network, analyzing data from a broad range of tests on a network that extends over several European countries. We presented an overview of the overlay network and presented our test methodology.

Running TCP tests was needed to assess the approximate bandwidth between each pair of partners, in order to verify if the chosen 1 Mbps rate for UDP tests is feasible. Preliminary tests helped detect and solve problems and the results of the final tests gave a positive acknowledge for performing tests. Average bandwidth was found to be over 4 Mbps not considering test results of partners with a known limited Internet connection and partners in the same MAN with a very high bandwidth connection between them.

The UDP tests tried to simulate VoIP traffic and tests were performed with high bandwidth (1 Mbps) stream also to simulate other types of multimedia traffic. The latter test had a combination of lower packet rate with increased packet size and vice-versa.

The VoIP tests revealed that on links without apparent problems there is only 1 packet loss on average for a 10 minute one-way stream. This is very good quality of service and it proves that VoIP shouldn't be a problem in today's Internet.

The higher throughput tests show that for low packet rate flows we have a good behavior too on the average. However, when packet losses occur, they tend to be bursty in nature, leading to temporary service interruptions. They also tend to be preceded by a group of packets which have a delay much higher than average, indicating the presence of buffers along the path, which cannot cope with interruption of service for the given packet rate.

High packet rate flows show a higher average packet loss rate, but it scales well with the increased rate: each one is about 9 times higher for UDP2 tests. This results in the same amount of time for service interruptions, since UDP2 losses are also bursty.

It is however difficult to determine the causes of lost packets, so future work is needed to see the effects of route changes or other bursty traffic on packet losses.

Acknowledgments

We would like to thank the EuQoS partners for making their networks and machines available for the tests and making this paper possible.

References

1. [IST] EuQoS - End-to-end Quality of Service support over heterogeneous networks (2004)
<http://www.euqos.org>
2. Serral-Gracià, R., Cabellos-Aparicio, A., Julian-Bertomeu, H., Domingo-Pascual, J.: Active measurement tool for the EuQoS project. In: MOME 3rd International Workshop on Internet Performance, Simulation, Monitoring and Measurements. IPS-MoMe 2005. (2005)
3. Hervella, J.F.R.: MGEN: The Multi-Generator Toolset (2002)
<http://matrix.it.uc3m.es/long/software/mgen6/mgen6/>
4. Jones, R.: Public netperf homepage (2003)
<http://www.netperf.org/netperf/NetperfPage.html>