

Measurement of IP Transport Parameters for IP Telephony

B.V.Ghita[†], S.M.Furnell[†], B.M.Lines[†], E.C.Ifeachor[†]

[†]Centre for Communications, Networks and Information Systems, Department of Communication and Electronic Engineering, University of Plymouth, Plymouth, United Kingdom

Abstract

This paper presents an overview of transport performance measurement for Voice over IP (VoIP, Internet Telephony, or IP Telephony) and a non-intrusive method of determining network performance parameters for voice packet flows within a VoIP call. The method allows both end-to-end performance monitoring of voice flows, but also inspects the transport parameters of a specific part of the network when VoIP traffic transits through it. The method is may be used for accurate fault location within a specified link.

Keywords

Voice over IP, transport parameters, non-intrusive monitoring, fault location.

1 Introduction

Voice over IP (VoIP) represents an area of significant interest in Internet world, and it changes the traditional concept in telecommunications from ‘data over voice’ to ‘voice over data’. Using the Internet instead of the Public Switched Telephone Network (PSTN) as a voice transport carrier has a number of advantages, the most obvious being the low cost for long-distance phone calls.

An important barrier in the development of VoIP is the Internet Protocol (IP) itself. IP works as a best-effort connectionless protocol, meaning that there are no guarantees about the delivery time or the reliability of a packet being transferred. The most important aspects, when considering an audio conference are exactly those that Internet cannot guarantee: time and bandwidth. The quality of the resulting conference depends upon the satisfaction of these requirements. Within this context, the concept of Quality of Service (QoS) was introduced. Although Internet represents an environment in which the QoS cannot be guaranteed, there are measurable parameters for a specific service, as presented in a QoS overview study [1].

The QoS is the overall rating for a service. Measurement of QoS essentially includes the assessment of a number of application dependent parameters and then combining them in a weighted sum. If we consider QoS for VoIP, the object of the analysis is voice quality at the receiving end, in terms of its two main characteristics, sound and interactivity. There are two main sources of impairments for the voice heard by the receiver: the codec, which compresses the speech flow in order to send it over the network at a lower bandwidth than original and the transport. The audio flow, after encoding, is packetised and sent over the Internet. Because of the Internet structure, the arrival of the packets at the destination cannot be guaranteed (see below). The paper is focused on the impairments introduced by the transport. From here on, we will neglect the impairments generated by the codec, and consider the transport QoS for equivalent with the overall QoS of the VoIP.

This paper identifies the performance issues associated with VoIP and then proceeds to present an offline method of determining network performance parameters for voice packet flows within a VoIP call. An advantage of the method is that it allows not only end-to-end performance monitoring of the flows, but also makes it possible to inspect the behaviour of the network when faced with delay sensitive traffic.

2 Transport performance considerations for VoIP

Building a list of performance parameters for a service should start by identifying the application that requires that specific service. For example, if the targeted application is a file transfer, then the delay or jitter parameters are almost irrelevant when compared to throughput or packet loss. In a similar manner, for a real-time application, delay is far more important than the other parameters. When considering QoS for VoIP applications, a network-related view of the performance should include the following parameters:

- delay - the time elapsed between the sending of a packet and its arrival at the destination
- jitter - the variance of the delay value
- packet loss - the number of lost packets, reported in the time elapsed
- throughput - the amount of data transferred from one place to another or processed in a specified amount of time (e.g. 1 second)

There are several suggested methods that can improve or guarantee the QoS for transport, such as DiffServ (Differentiated Services) [2], Tenet [3], or QoS Routing combined with RSVP (Reservation Protocol) [4]. Unfortunately, none of them have been applied on global basis because of the large dimension and complexity of the Internet. Therefore, it is vital to determine in such an environment whether or not a specific connection meets the requirements of a VoIP call.

2.1 Current state for voice flows performance measurement

VoIP is a relative new concept, therefore most of the work performed in this area is still on a development stage. From the large range of standards for VoIP, the H.323 protocol stack [6], developed by the ITU, was considered for the work presented in this paper.

Transport QoS has two main areas: end-to-end measurements and, in case there are changes in the level of parameters, fault localisation. An example is given in Figure 1. This shows, for an arbitrary division of the entire route of the packets, the end-to-end parameters, and two sets of parameters, 'East' and 'West', which can be used to localise a fault in either 'East' or 'West' sub-network, by comparison with the end-to-end parameters.

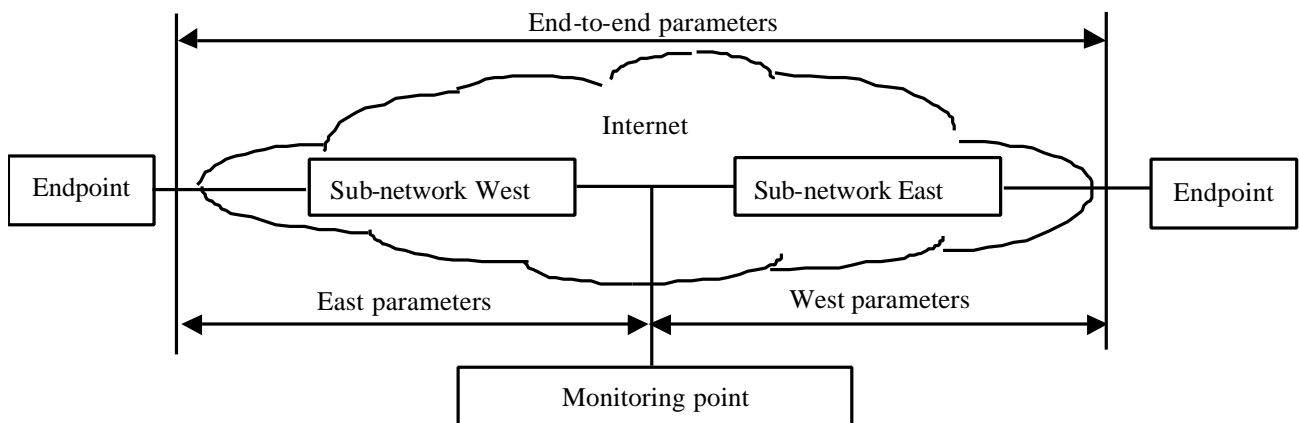


Figure 1: The Performance parameters for a general example of monitoring

In a classical approach, the two aims would require a 3-tool configuration. For end-to-end measurements, testing clients should be put at both ends, and, for fault location, a testing server should be placed at the monitoring point. After that, traffic should be collected by the end stations, then sent to the server, in order to be analysed and compared with the data collected by it. This description has two main disadvantages:

- it is intrusive; in the best case, even if the endpoint clients are just monitoring, they have to send the data to the server in order to be analysed;
- it requires placement of monitoring devices at both ends.

Current tools (e.g. Hammer VoIP Analysis System, HP Internet Advisor, rtpmon [7]) can be used to determine the QoS for transport from the audio flows (which run on RTP, Real Time Protocol) within a call. The calculations are based on parsing both the RTP and/or the accompanying control flows (running on RTCP, Real Time Control Protocol) and displaying the available data; both these flows will be described later. The main disadvantages are as follows:

- none of the tools can establish fault location without using the traditional approach mentioned above;
- they do not build any relation between the end-to-end parameters, obtained from the RTCP flows, and the end-to-monitoring-point parameters, obtained from the RTP monitoring.

Considering all these limitations, we aim to obtain a better view of the network performance, without using several devices and, if possible, without injecting additional traffic into the network. This paper presents a non-intrusive method of determining the transport performance parameters for the real-time traffic within a VoIP call, using a single point of monitoring, as well as a possible pseudo-non-intrusive improvement that can raise its performance. The proposed method can reveal both the end-to-end performance and the fault localisation, if the monitored parameters change their value along the route, and also avoids both of the disadvantages identified.

3 Method and implementation

The monitoring procedure comprises three steps. First, the voice flows (RTP) are identified and then captured using one of the capture programs. In the monitoring phase, the RTP header fields and the RTCP packets are used to determine the performance parameters. Then correlation of RTP and RTCP is used to establish the location of the problem area. The stages are described in more detail in the following paragraphs.

3.1 Identification of the audio flows

The analysis is targeted on the audio streams. The ports on which the audio streams run can be determined only by capturing the connection establishment phase, then parsing the setup and control messages, which contain the audio stream ports as parameters.

3.2 Parameters measurement using RTP monitoring and RTCP parsing

The analysis uses as input the header fields of RTP packets [5] together with the timestamp of the packet arrival, given by the capture program. The following types of parameters can be determined using the RTP header fields and the arrival timestamp of each packet, from the packet capture program:

- a. delay-related parameters: inter-arrival delay, inter-arrival jitter, one-way delay jitter;
- b. packet-accounting parameters: lost packets and out of order packets;
- c. flow speed parameters: throughput.

The RTCP packets can be used as an instrument for end-to-end measurements. Their fields provide the values for inter-arrival jitter and lost packets. The following issues appear when we use RTCP to analyse the flows:

- it runs on UDP (User Datagram Protocol), therefore it is possible that a number of packets will not arrive, so no data will be available for that period of time;
- it has scalability problems [8]. The RTCP messages are limited to 5% of the whole traffic. In the case of a many-to-many conference, on normal behaviour, there would be a low number of RTP messages per-terminal (in order to maintain the 5% limit) [5];
- it returns only end-to-end parameters, therefore cannot locate the cause of parameter changes (this problem exists regardless of the conference characteristics)

It should be noted that the analysis is performed on a 'per-flow' basis. Prior to performing the analysis, the incoming packets (from several audio channels) are split into flows (each flow representing a channel). When saying successive packets, we refer to packets belonging to the same flow.

3.3 Correlating RTP analysis with RTCP content

By correlating the two sets of parameters, obtained from RTP and RTCP, it can be determined whether or not a specific problem (e.g. a high number of lost packets) is caused by a problem which exists in the East sub-network or the West sub-network. Figure 2 presents the captured flows.

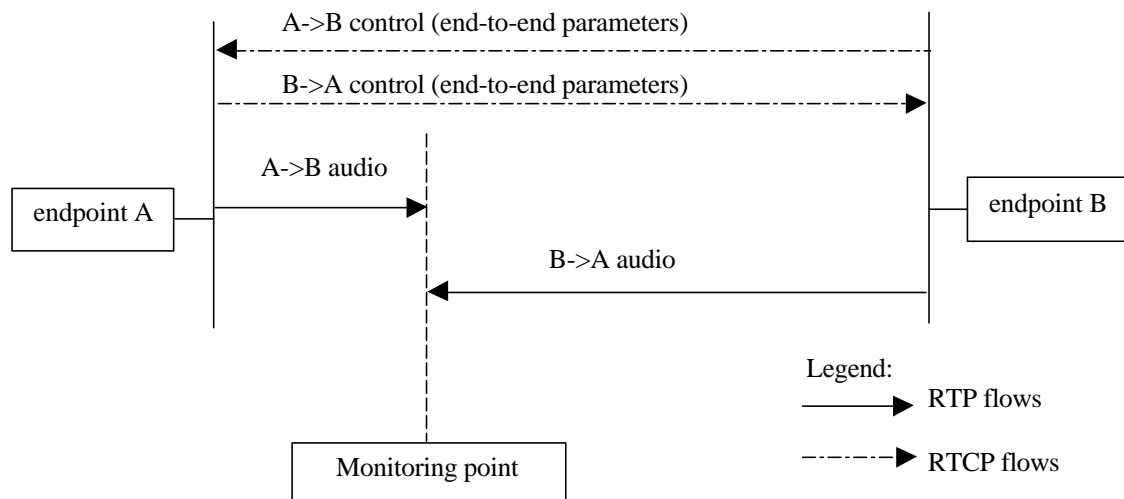


Figure 2 RTP and RTCP flows monitoring

The RTP streams, as captured on the monitoring point, are:

- A->B: after passing through the West sub-network;
- B->A: after passing through the East sub-network.

Therefore, by measuring the parameters of these flows, we can determine the performance of:

- the West sub-network, from the A->B flow;
- the East sub-network, from the B->A flow.

We have to bear in mind that the A->B direction does not fully characterise the behaviour of the network, as it can be very good for one direction and bad for the other (it does not have to be symmetrical in terms of performance).

Meanwhile, as mentioned, RTCP provides the end-to-end parameters, i.e. the performance of the entire A->B and B->A routes, but it has no indication about how these parameters change on the

route, i.e. cannot establish where a faulty behaviour of the network determined a change in the values of the parameters.

Putting together the two sets we obtain parameters for the following segments:

- A->B and B->A, end-to-end – from the RTCP flows;
- A->monitoring point and B->monitoring point – from the RTP flows;
- monitoring point->B and monitoring point->A – by subtracting the RTP obtained values from RTCP end-to-end parameters.

In conclusion, by using both RTP and RTCP, it is possible to obtain both the end-to-end and the end-to-monitoring point parameters for the monitored flows. A network testbed has been constructed in order to validate the proposed method. The testbed includes two networks, connected through an emulated faulty link, and the monitoring point is placed on the route, at the exit point (after the router) of one of the networks. A detailed description of the method validation procedure is given in [9].

Although the monitoring tool was built, and these preliminary tests were performed, for a proper tuning, further analysis is required in a real or simulated VoIP environment. Such an environment would include several simultaneous conferences, running between endpoints situated at different location, over various routes.

4 Fault location improvements

The presented method has two main advantages:

- non-intrusive. All the measurements are performed without generating any additional traffic. Given that they do not require any supplementary capacity, the measurements can be performed continuously without affecting the traffic in any way
- single point. A single monitoring station/device is enough to provide the user with all the presented data.

In spite of these advantages, the method presented also has several problems. One of them is the fault location: fault location cannot be determined more precisely than the presented ‘East network’ and ‘West network’. If we assume a generic case, a segment with parameters we want to measure, such a measure will not fully determine if a parameter change is due to that segment behaviour or some other parts of the network. Therefore, for further information, additional monitoring must be performed. The most convenient configuration for such an enhancement is to place two monitoring stations, at the end of the segment to be analysed, as illustrated in Figure 3.




Figure 3 - Enhanced fault location - Segment monitoring

The two monitoring stations are performing the same tasks as in the one in Figure 2, but they have to exchange information with each other. The secondary monitoring station has to transmit its set of measured parameters (East/West) to the main monitoring station. The main monitoring station uses

these parameters and compares them with the ones measured by it, in order to determine if there is any degradation on the monitored segment. In this way, four different sets of parameters can be obtained, for the following segments:

- end-to-end
- sub-network West
- sub-network East
- monitored segment

For achieving the full measurement of the monitored segment's parameters there is a compromise that has to be made: the proposed method is no longer non-intrusive (as mentioned, the two stations will exchange data). However, the method can be considered pseudo-non-intrusive, as the data sent between the two stations is not used as input for parameter measurement (as it happens in an intrusive method), but is only to inform the main station about the results of the measurements. Therefore, it will not increase (and, as a result, it will not affect) the overall traffic.

5 Conclusions

In this paper an off-line method to measure the QoS transport parameters for an H.323 VoIP call from a single point, by non-intrusive monitoring, has been described. A possible enhancement to it, in order to improve the fault location function was also introduced. The jitter and packet loss analysis seem promising, but further work is required to determine, monitor and analyse the other parameters. Also, a specific change in the performance parameters group can be related with a specific network event (e.g. a congested router). Therefore, analysis of the dynamics of the calculated parameters is required.

References

- [1] Stiller B., 'Quality of Service Issues in Networking Environments', internal report, www.cl.cam.ac.uk:80/ftp/papers/reports/TR380-bs201-qos_issues.ps.gz
- [2] RFC 2474 – 'Differentiated Services Field', December 1998.
- [3] Ferrari D., Banerjea A., Zhang H. 1994. 'Network Support for Multimedia – A Discussion of the Tenet Approach', Computer Networks and ISDN Systems, December 1994.
- [4] RFC 2386 – 'A Framework for QoS-based Routing', August 1998.
- [5] RFC 1889 – 'RTP – A Transport Protocol for Real-Time Applications', January 1996.
- [6] ITU. 1998. 'Packet based multimedia communication systems', H.323 ITU Recommendation, February 1998.
- [7] Bacher D., Swan A. 1996. 'rtpmon: A Third-Party RTCP Monitor', *ACM Multimedia '96*.
- [8] Rosenberg J., Schulzrinne H. 1998. 'Timer Reconsideration for Enhanced RTP Scalability', *Proceedings of IEEE Infocom 1998*, March 29 - April 2, San Francisco
- [9] Ghita B.V., Furnell, S.M, Lines, B.M., Le-Foll, D., and Ifeachor, E.C. 2000. 'IP network performance monitoring of voice flows for IP telephony', to appear in *Proceedings of the Second International Network Conference (INC 2000)*, July 3-6, 2000, Plymouth, UK.