Measuring Internet telephony quality: Where are we today?

Olof Hagsand[1], Ian Marsh[2] and Kjell Hanson[3]

[1] LCN Laboratory, IMIT, Royal Institute of Technology, Sweden
`olofh@kth.se`
[2] SICS AB, Stockholm, Sweden
`ianm@sics.se`
[3] Prosilient Software AB, Stockholm, Sweden
`kjell@prosilient.com`

Abstract. Users of Internet telephony applications demand good quality audio playback. This quality is largely dependent on the instantaneous network conditions as well as the time of day. In this paper we describe a scheme for measuring network connections as well as a motivation for including a new metric when judging quality. Our tests included a wide range of geographically distributed sites. Our results give useful feedback to users and operators of IP Telephony networks and important information for developers of Voice over IP applications.

1 Introduction

The Internet makes no *guarantee* about the delivery of data in real time to applications such as Internet telephony. Due to the shared nature of many of the resources as well as propagation delays this is a difficult, if not impossible, task. All is not lost however, under good conditions in a well dimensioned network timely delivery of packets is achievable.

Due to the number of simultaneous connections and relative times across connected sites, the instantaneous quality of connections can vary dramatically. The goal of this paper is to show the benefit of taking and measuring quality so action can be taken immediately (if possible) or in the future. Also we wanted to produce a report on "how well are we doing today".

Internet telephony and voice over IP applications are already being used on the Internet. Most transport data using the Real Time Protocol (RTP)[SCFJ96] report statistics using RTCP, the Real Time Control Protocol. One function of RTCP is to report statistics back to the sender on the received quality of its receiver(s). In combination, the sender and receiver can obtain information on the following: the packet loss since the last message, the total packet loss in this session, the variance in packet arrival (jitter) and an estimate of the round trip time.

Many studies and measurements have been conducted on the Internet. Two of the most recent cited (and complete) works were done by Vern Paxson in [Pax96,Pax97]. Work dedicated to the measurement of audio data included Jean

Bolot et. al in 1995 in [BCG95] where the authors performed loss measurements and developed a model indicating that forward error correction (FEC) would rectify most loss situations. In 1997 work done by Mexemchuck and Lo in [ML97] defines the quality of a connection is as the fraction of the time that a channel is free of distortion for intervals that are long enough to transmit active speech segments.

In this work we look at both loss and delay and give measurements including these metrics, additionally we look closer at the interarrival variance or *jitter*, in particular the statistical distributions for a number of Internet sites. The results not only impact on quality but also include useful hints for application writers when designing adaptive playout schemes.

2 Measurements

2.1 Motivation

In addition to the simple quality argument, we have give a number of motivating reasons for performing wide area measurements:

- To perform an up-to-date series of measurements on today's Internet as far as voice is concerned.
- To dismiss hard limits for delay regarding Internet telephony quality [Int93]. We agree that delay should be minimized as much as possible without sacrificing playout quality, however it should not be a *definition* of the quality.
- The number of users and their access to the Internet, particularly through wireless connections, will affect the traffic considerably, we wanted to make measurements before this next quantum jump occurs.
- On well provisioned links we can carry packet audio without additional QoS mechanisms such as Integrated or Differentiated services. Finding the limit of what a provisioned link is, is one goal of this work.
- Clearly the time which one sends audio packets clearly influences the quality one receives. Traffic varies according to a daily rhythm. The respective times of both ends of the communication should be taken into account when presenting the measurement results.
- Work done in [Pax97] cites asymmetric routes. Therefore propagation delays, packet losses and possibly quality might be affected. This is expanded further in Section 2.4.

2.2 Delay

As most telephony users can be sensitive to high delays, we include a delay estimation. Measuring the real end-to-end delay is non trivial. Problems with synchronizing clocks and the best we can achieve is an estimation of the delay. We can assume the end-to-end delay is equal to the sum of some small packetization delay at the sender, the propagation delay and the receiver induced delay. Artificially introduced delay at the receiver is necessary to achieve smooth playout of packets due to interarrival differences contributed by the statistical multiplexing of packet switched networks. This is a well known artifact and is explained in more detail in [RKTS94] [MKT95] bounds and [Sch92]

The packetization delay on most operating systems is relatively small approximately 20ms according to [SSO83]. The propagation delay we can approximate from RTT timer in RTCP and dividing by two. A point to note is if the route is asymmetric then this may be an approximation, however other methods exist to calculate the one-way delay. The delay added by buffering can be estimated by counting the number of bytes that the packet must wait in a buffer [Jac94].

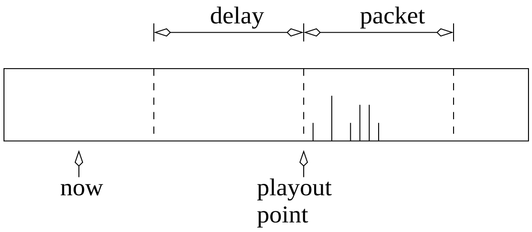

Fig. 1. Delaying packets for playout

The time the packet must wait (in milliseconds) is dependent on the sample rate and the number of bits per sample:

$$\text{delay} = \frac{sampling\ \ rate\ \cdot\ bytes}{samples\ \ per\ \ sec\ \cdot\ bits\ \ per\ \ sample}$$

for a typical 160 byte payload at 8000 Hz this becomes:

$$\text{delay} = \frac{8000 \cdot 160}{8000 \cdot 8} = 20\,ms$$

2.3 Interarrival times

Interarrival times can be obtained by finding the difference of consecutive arrivals. The variance of interarrival times is found by continually updating the differences from previous packets. The mean variation is usually calculated for each packet and this value is modified by a smoothing constant normally less than 1. The smoothing constant can be tuned to give more or less weight to recent arrivals. This mean variance or *jitter* gives a good indication of the connection conditions. Plots of some interrarival times for connections are shown in Figure 5. The jitter value shown in these figures is sent by the receiver to the sender as part of the RTCP receiver report.

2.4 Network asymmetry

Asymmetry is a problem in today's Internet. Paxson reported that in 1995, half of the 40,000 measurements he took had at least one different city in each direction of a bidirectional flow. There are implications for IP telephony, particularly loss, where one receiver may report much poorer quality than the other. We show one example of an asymmetric plot of the interarrival times for a trace in Figure 9.

2.5 Software tool and measurements

We have implemented an IP telephony tool capable of sending RTP and RTCP packets. We can record and play PCM encoded files. The textual log files are in the following format:

```
I 919339994 729295
T 919339995 732069
E 919339996 38883 172 32869 0 160
E 919339996 68883 172 32869 1 320
---------------------------------
1 2         3     4   5     6 7
(Header: V=2 P=0 X=0 CC=0 M=0 PT=101)
```

Fig. 2. Log file format

3 Measurement setup

Our selected test sites consisted of one national (Luleå, Northern Sweden), one continental (Cambridge, UK), one trans-Atlantic site in the US, (Amherst, Cambridge) and finally one site in Buenos Aires in Argentina. For all measurements, a recorded session was sent from the remote site to a local site in Sweden. In the bidirectional tests the executions where performed simultaneously. Table 3 gives the details of the recorded session.

Trace File	Value
File Size	584480 bytes
Duration	70 secs
Payload	160 bytes (20 ms)
With Silence Suppression.	3643 packets
W/o Silence Suppression	2064 packets
% Transmitted	56.6 %

Fig. 3. Trace file properties

4 Results and Discussion

Figure 4 shows the delay as reported by RTCP at the receiver for hosts in Argentina and the US. It is quite evident that the connection to the US is much better. Table 6 shows a summary of the delay for the four sites. Figure 5 shows

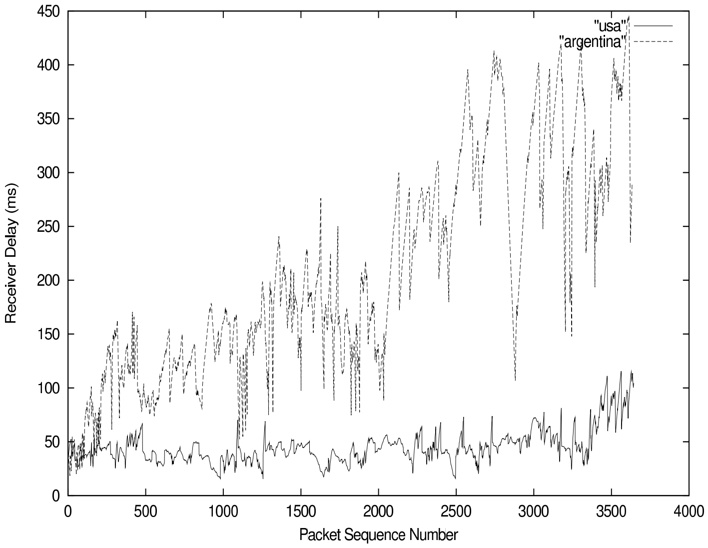

Fig. 4. Receiver Delay for 3 Remote Sites

the variance in the packet interarrivals. Packets from the national site (Luleå) form a peak which would result in less jitter and consequently less frequent changes in the playout buffer size. The converse is true for the connection to Argentina, where the variance typically ranges from 20ms to 150ms.

Figure 6 shows an estimate of the propagation delay (based on the RTCP reporting) and an estimation of the delay incurred by the end system. Within RTCP it is possible to extract the time a report was held at a particular node. Figure 7 shows the number of hops and the measured jitter for the same four sites. Finally figures 8 and 9 show losses in both directions for a conversation from Stockholm to Buenos Aires, in the first figure it is clear to see the higher loss rate from Argentina to Sweden. Similarly so for the interarrival times, which show much more variance in the South America to Europe direction.

5 Conclusions/Future Work

In this short paper we have presented some criteria for evaluating the quality of Internet Telephony connections. On average, with the exception of the South American site, reasonable quality Internet telephony can be supported. Obviously more connections need to be tested, but can be done with the tools we

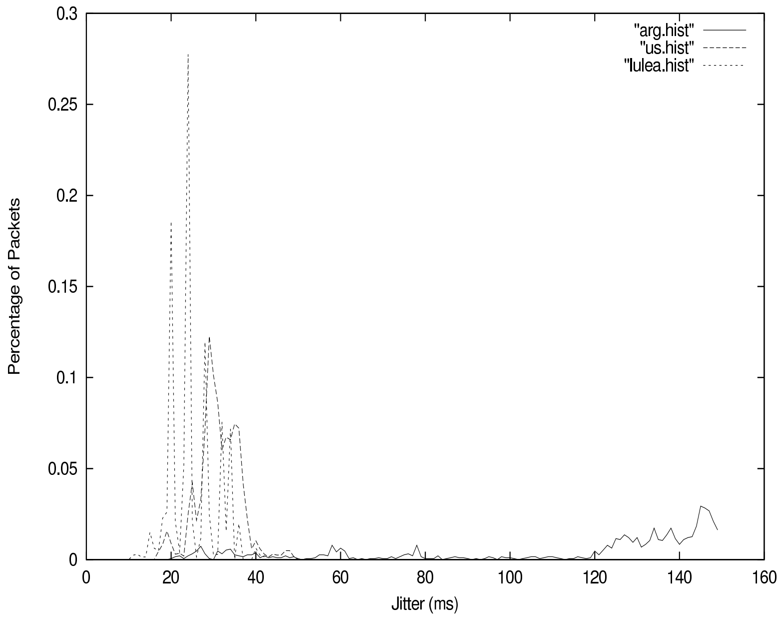

Fig. 5. The interarrival histogram for two sites

Source	Propagation	Receiver	Total
Luleå (Sweden)	12.13	42.14	54.27 ms
Cambridge (UK)	33.3	54.39	87.68 ms
Amherst (US)	57.3	50.82	108.12 ms
Buenos Aires (Arg)	273.0	119.12	392.12 ms

Fig. 6. Mean Delay Values

Source	Hops	Jitter
Luleå (Sweden)	9	5.12 ms
Cambridge (UK)	15	7.38 ms
Amherst (US)	15	11.25 ms
Buenos Aires (Arg)	18	117.69 ms

Fig. 7. Mean Jitter Values

have started to develop. One proviso, the sites used were located on academic networks, further tests would have to be done on commercial networks.

We have developed a tool that probes, gathers and produces statistics based on the RTP and RTCP protocols. Additionally we have included time-driven

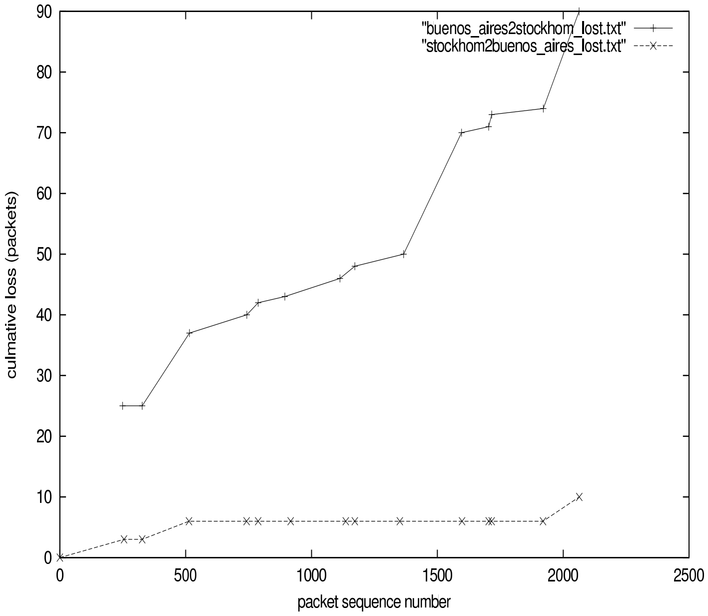

Fig. 8. Bidirectional Loss (Cumulative)

tracing as well as looked at the bidirectional (and hence asymmetric) quality. We plan to extend the tool to operate as a daemon in which the tests can run without intervention. This also allows any site to connect to any other therefore producing a full mesh of connections rather than our current centralized system.

We would like to thank Thiemo Voigt and Bengt Ahlgren for their comments on this paper, Ericsson and Telia for their financial support. Also thanks to Steve Pink, Jim Kurose and Pablo Giambiagi for the accounts that we have used.

References

[BCG95] J. Bolot, H. Crepin, and A. Garcia. Analysis of audio packet loss in the internet. In *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Lecture Notes in Computer Science, pages 163–174, Durham, New Hampshire, April 1995.

[Int93] International Telecommunication Union (ITU). Transmission systems and media, general recommendation on the transmission quality for an entire international telephone connection; one-way transmission time. Recommendation G.114, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, March 1993.

[Jac94] Van Jacobson. Multimedia conferencing on the Internet. In *SIGCOMM Symposium on Communications Architectures and Protocols*, London, England, August 1994. Tutorial slides.

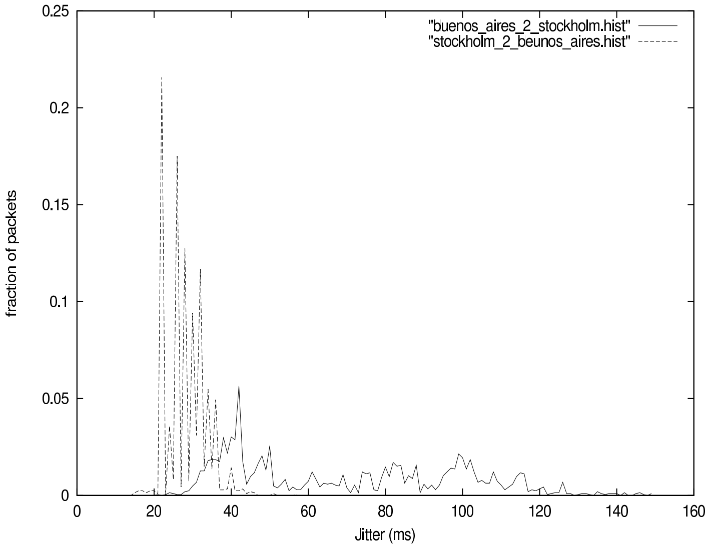

Fig. 9. Histogram of the bidirectional interarrival times

[MKT95] Sue B. Moon, Jim Kurose, and Don Towsley. Packet audio playout delay adjustment algorithms: performance bounds and algorithms. Research report, Department of Computer Science, University of Massachusetts at Amherst, Amherst, Massachusetts, August 1995.

[ML97] N. F. Maxemchuk and S. Lo. Measurement and interpretation of voice traffic on the internet. In *Conference Record of the International Conference on Communications (ICC)*, Montreal, Canada, June 1997.

[Pax96] Vern Paxson. End-to-end routing behavior in the internet. In *SIGCOMM Symposium on Communications Architectures and Protocols*, Stanford, California, August 1996.

[Pax97] Vern Paxson. End-to-end internet packet dynamics. In *SIGCOMM Symposium on Communications Architectures and Protocols*, Cannes, France, September 1997.

[RKTS94] Ramachandran Ramjee, Jim Kurose, Don Towsley, and Henning Schulzrinne. Adaptive playout mechanisms for packetized audio applications in wide-area networks. In *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, pages 680–688, Toronto, Canada, June 1994. IEEE Computer Society Press, Los Alamitos, California.

[SCFJ96] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: a transport protocol for real-time applications. Request for Comments (Proposed Standard) 1889, Internet Engineering Task Force, January 1996.

[Sch92] Henning Schulzrinne. Voice communication across the Internet: A network voice terminal. Technical Report TR 92-50, Dept. of Computer Science, University of Massachusetts, Amherst, Massachusetts, July 1992.

[SSO83] Daniel C. Swinehart, L. C. Stewart, and S. M. Ornstein. Adding voice to an office computer network. In *Proceedings of the IEEE Conference on*

Global Communications (GLOBECOM), pages 392–398 (11.4), San Diego, California, November 1983. IEEE. also Xerox Report CSL-83-8.