

# Wide Area Measurements of Voice Over IP Quality

Ian Marsh<sup>1</sup>, Fengyi Li<sup>2</sup> and Gunnar Karlsson<sup>2</sup>

<sup>1</sup> SICS AB, Kista S-164 29, Sweden  
ianm@sics.se

<sup>2</sup> LCN Laboratory, IMIT, Royal Institute of Technology, Sweden  
d97-fli@nada.kth.se, gk@imit.kth.se

**Abstract.** Time, day, location and instantaneous network conditions largely dictate the quality of Voice over IP calls. In this paper we present the results of over 18000 VoIP measurements, taken from nine sites connected in a full-mesh configuration. We sample the quality of the routes on a hourly basis by transmitting a pre-recorded call between a given pair of sites. We repeat the procedure for all nine sites during the one hour interval. Based on the obtained jitter, delay and loss values as defined in RFC 1889 (RTP) we conclude that the VoIP quality is acceptable for all but one of the nine sites involved. We also conclude that VoIP quality has improved marginally since we last conducted a similar study in 1998.

## 1 Introduction

It is well known that the users of real-time voice services are sensitive and susceptible to variable audio quality. If the quality deteriorates below an acceptable level, or is too variable, users often abandon their calls and retry later. Since the Internet is increasingly being used to carry real-time voice traffic, the quality provided has become, and will remain an important issue. The aim of this work is therefore to disclose the current quality of voice communication at selected end-points on the Internet.

It is intended that the results of this work will be useful to many different communities involved with real-time voice communication. We now list some groups to whom this work might have relevance. First end users can determine which destinations are likely to yield sufficient quality. When deemed insufficient they can take preventative measures such as adding robustness, for example in the form of forward error correction to the outgoing packets. Operators can use findings such as these to motivate upgrading links, adding control admission schemes, rerouting traffic or even implementing QoS mechanisms where poor quality is being reported. Network regulators can use these kind of measurements to verify the quality level that has been agreed upon or even to settle disputes. Speech coder designers can utilise the data as input for a new class of codecs, of particular interest are designs which yield good quality in the case of bursty packet loss. Finally, researchers could use the data for their own investigations.

The structure of this paper is as follows: Section 2 begins with some background on the quality measures we have used in this work namely, loss, delay and jitter. Following on from the quality measures, section 3 gives a description of the methodology used to measure the parameters. In section 4 the results are presented (and condensed) into a single table. In section 5 the related work is given, comparing results obtained in this study with other researchers' work. This is considered important as it indicates whether quality has improved or deteriorated since those studies. Section 6 rounds off with some conclusions and a pointer to the data we have collated.

## 2 What Do We Mean by Voice over IP Quality?

Ultimately, users judge the quality of voice transmissions. Organisations such as ETSI, ITU, TIA, RCR plus many others have detailed methods to assess voice quality. Assigning a quality measure involves replaying coded voice to both experienced and novice listeners and asking them to adjudge the perceived quality. Measuring the quality of voice data that has been transmitted across a wide area network is more difficult primarily due to the delay. The network inflicts its own impairment on the quality of the voice stream. By measuring the delay, jitter and loss of the incoming data stream at the receiver, we can provide some indication on how suitable the *network* is for real-time voice communication.

It is important to point out we did not include the quality contribution of the end systems in this study. This is because the hardware was different at each site, even with our own software Sicsophone [HMH03]. But also, to avoid including the delays of our own application into the measurements. In order to assess the delay contribution of each end system it would be difficult without isolation tests. We chose to use simple A-law PCM coding to maintain a theoretically known coding/decoding delay.


The quality of VoIP sessions can be quantified by the network delay, packet loss and packet jitter. We emphasise that these three quantities are the major contributors to the perceived quality as far as the *network* is concerned. The G.114 ITU standard states that the end-to-end one way delay should not exceed 150ms [RG98]. Delays over this value adversely effect the quality of the conversation. An alternative study by Cole and Rosenbluth state that users perceive a linear degradation in the quality up to 177ms [CR02]. Above this figure the degradation is also linear although markedly worse. As far as the packet loss is concerned, using speech coding schemes such as A-law or  $\mu$ -law coding, tests have shown that the mean packet loss should not exceed 10% before an appreciable loss in quality. Note that a loss rate such as this does not say anything about the distribution of the losses. As far as the authors are aware of, no results exist that state how jitter solely can affect the quality of voice communication. Work on jitter and quality are often combined with loss or delay factors. When de-jittering mechanisms are employed, the network *jitter* is manifested into application *delay or loss*. The application must hold back a sufficient number of packets in order to ensure smooth, uninterrupted playback of speech. To sum-

marise, we refer to the quality as a combination of delay, jitter and loss. It is important to mention we explicitly do not state how these values should be combined. The ITU E-model is one approach but others exist, therefore we refer the interested reader to the references such as [LE01] and [KKI91].

### 3 Simulating and Measuring Voice over IP Sessions

Our method to measure VoIP quality is to send pre-recorded calls between globally distributed sites. Through the modification of our own VoIP tool, Sicsophone, the intervening network paths are probed by a 70 second pre-recorded ‘test signal’. The goal of this work is therefore to report in what state the signal emerges after traversing the network paths.

Nine sites have been carefully chosen with different hop counts, geographic distances and time-zones to obtain a diverse selection of distributed sites. One important limitation of the available sites was they were all located at academic institutions, which are typically associated with well provisioned networks. Their locations are shown in Figure 1. The sites were connected as a full mesh allowing



**Fig. 1.** The nine sites used in the 2002 measurements are shown with circles. The six depicted with squares show those that were available to us in 1998, three remained unchanged during the past four years.

us, in theory, to measure the quality of 72 different Internet paths. In practice, some of the combinations were not usable to port blocking there were four such cases. Bi-directional sessions were scheduled on a hourly basis between any two given sites. Calls were only transferred once per hour due to load considerations on remote machines.

In Table 1 we list the characteristics of the call we used to probe the Internet paths. As stated, the call is essentially a fixed length PCM coded file which can be sent. The file consisted of a PCM coded recording of a read passage of text. Over a 15 week period we gathered just over 18,000 recorded sessions. The number of sessions between the nine sites is not evenly distributed due to outages at some sites, however we attempted to ensure an even number of measurements per site.

<i>Test "signal"</i>	
Call duration	70 seconds
Payload size	160 bytes
Packetisation time (ms)	20ms
Data rate	64kbits/sec
With silence suppression	2043 packets
Without silence suppression	3653 packets
Coding	8 bit PCM
Recorded call size	584480 bytes
<i>Obtained data</i>	
Number of hosts used (2003)	9
Number of traces obtained	18054
Number of data packets	32,771,021
Total data size (compressed)	411 Megabytes
Measurement duration	15 weeks

**Table 1.** The upper half of the table gives details of the call used to measure the quality of links between the sites. The lower half provides information about the data which was gathered.

### 3.1 A Networking Definition of Delay

We refer to the delay as the *one way network* delay. One way delay is important in voice communication, particularly if it is not the same in each direction. Measuring the one way delay of network connections without synchronised clocks however, is a non-trivial task. Hence many methods rely on round-trip measurements and halve the values, hence estimating the one way delay. We measured the network delay using the RTCP protocol which is part of the RTP standard [SCFJ96]. A brief description follows. At given intervals the sender transmits a so called "report" containing the time the report was sent. On reception of this report the receiver records the current time. Therefore two times are recorded within the report. When returning the report to the sender, the receiver subtracts the time it initially put in the report, therefore accounting for the time it held the report. Using this information the sender can calculate the round-trip delay and importantly, discount the time spent processing the reports at the receiver. This can be done in both directions to see if any significant anomalies exist. We quote the network delay in the results section as it explicitly does not

include any contribution from the end hosts. Therefore it is important to state the delay is not the mouth-to-ear delay, rather only the network delay.

### 3.2 Jitter - An IETF Definition

Jitter is the statistical variance of the packet interarrival time. The IETF in RFC 1889 defines the jitter to be the mean deviation (the smoothed absolute value) of the packet spacing difference between the sender and the receiver [SCFJ96]. Sicsophone sends packets of identical size at constant intervals which implies that  $S_j - S_i$  (the sending times of two consecutive packets) is constant. The difference of the packet spacing, denoted  $D$ , is used to calculate the interarrival jitter. According to the RFC, the interarrival jitter should be calculated continuously as each packet  $i$  is received. The interarrival jitter  $J_i$  for packet  $i$  is calculated using the previous packet  $J_{i-1}$  thus:

$$J_i = J_{i-1} + (|D(i-1, i)| - J_{i-1})/16.$$

According to the RFC “the gain parameter 1/16 gives a good noise reduction ratio while maintaining a reasonable rate of convergence”. As stated earlier, buffering due to jitter adds to the delay of the application. This delay is accounted for in the results we present. The real time needed for de-jittering depends on how the original time spacing of the packets should be restored. For example if a single packet buffer is employed it would result in an extra 20ms (the packetisation time) being added to the total delay. Note that packets arriving with a spacing greater than 20ms should be discarded by the application as being too late for replay. Multiples of 20ms can thus be allocated for every packet held before playout in this simple example. To summarise, the delay due to de-jittering the arriving stream is implementation dependent, thus we do not include it in our results.

### 3.3 Counting Ones Losses in the Network

We calculate the lost packets as is exactly defined in RFC 1889. It defines the number of lost packets as the expected number of packets subtracted by the number actually received. The loss is calculated using expected values so as to allow more significance for the number of packets received. For example 20 lost packets from 100 packets has a higher significance than 1 lost from 5. For simple measures the percentage of lost packets from the total number of packets expected is stated. As stated, the losses in this work *do not* include those incurred by late arrivals, as knowledge of the buffer playout algorithm is needed. Detailed analysis of the loss patterns is not given in the results section, we simply state the percentages of single, double and triplicate losses.

## 4 Results

The results of 15 weeks of measurements are condensed into Figure 2. The table

should be interpreted as an 11 by 11 matrix. The locations listed horizontally across the top of the table are the locations used as receivers. Listed vertically they are configured as senders. The values in the rightmost column and bottom row are the statistical means for all the connections **from** the host in the same row and **to** the host in the same column respectively. For example the last column of the first row (directly under “Mean”) is the average delay to all destinations from Massachusetts (112.8ms).

Each cell includes the delay, jitter, loss, number of hops and the time difference prefixed by the letters D, J, L, H and T for each of the connections. The units for each quantity are the delay in milliseconds, the jitter in milliseconds, the loss as a percentage, the hops as reported by traceroute and time differences in hours. A ‘+’ indicates that the local time from a site is ahead of the one in the corresponding cell and behind for a ‘-’. The values in parenthesis are the standard deviations. A NA signifies “Not Available” for this particular combination of hosts. The bottom rightmost cell contains the mean for all 18054 calls made, both to and from all the nine hosts involved.

The most general observation is the quality of the paths is generally good. The average delay is just below the ITU’s G.114 recommendation for the end-to-end delay. Nevertheless at 136ms it does not leave much time for the end systems to encode, dejitter, decode and replay the voice stream. A one packet buffer would easily absorb the 4ms jitter. However the 150ms threshold can be seen as the lowest possible, some extra delay can be tolerated.

There are two clear groupings from these results, those within the EU and the US and those outside. The connections in Europe and the United States (and between them) are very good. The average delay between the US/EU hosts is 105ms, the jitter is 3.76ms and the loss 1.16%. Those outside fair less well. The Turkish site suffers from large delays, which is not surprising as the Turkish research network is connected via a satellite link to Belgium (using the Geant network). The jitter and loss figures however are low, 5.7ms and 4% respectively. The Argentinian site suffers from asymmetry problems. The quality when sending data to it is significantly worse than when receiving data from it. The delay is 1/3 higher, the jitter is more than twice it in the opposite direction and the loss is nearly four times higher than when sending to it. Unfortunately we could not perform a traceroute from the host in Buenos Aires, so we cannot say how the route contributed to these values.

We now turn our attention to results which are not related to any particular site. As far as loss is concerned the majority of losses are single losses. 78% of all the losses counted in all trace files were single losses whereas 13% were duplicate losses and 4% triplicate losses. For some connections (22 from 68), some form of packet loss concealment would be useful, as the loss is over 1% but nearly always under 10%.

Generally the jitter is low relative to the delay of the link, approximately 3-4%. This is not totally unexpected as the loss rates are also low. With the exception of the Argentinian site, the sites did not exhibit large differences in asymmetry and were normally within 5% of each other in both directions. It is

interesting to note that the number of hops could vary somewhat under the 15 week measurement period. Only very few ( $< 0.001\%$ ) out of sequence packets were observed. Within [Li02] there are details of further tests, such as the effect of using silence suppression, using larger payload sizes and daytime effects. In summary no significant differences were observed in these tests. We can attribute this (and the good quality results) to generally well-provisioned academic networks.

## 5 Related Work

Similar but less extensive measurements were performed in 1998 [HHM99]. Only three of the hosts remain from four years ago, so comparisons can only be made for these routes (and towards Stockholm). An improvement, in the order of 5-10% has been observed for these routes. We should point out though, the number of sessions recorded four years ago numbered only tens per host, whereas on this occasion we performed hundreds of calls from each host. Bolot et. al. looked at consecutive loss for designing an FEC scheme [BCG95]. They concluded that the number of consecutive losses is quite low and stated that most losses are one to five losses at 8am and between one to ten at 4pm. This is in broad agreement with the findings in this work. Maxemchuk and Lo measured both loss and delay variation for intra-state connections within the USA and international links [ML97]. Their conclusion was the quality depends on the length of the connection and the time of day. We did not try different connection durations but saw much smaller variations (almost negligible) during a 24 hour cycle (see [Li02]). We attribute this to the small 64kbits per second VoIP session on well dimensioned academic networks. It is worthy to point out our loss rates were considerably less than Maxemchuks (3-4%). Dong Lin had similar conclusions [Lin99], stating that in fact even calls within the USA could suffer from large jitter delays. Her results on packet loss also agree with those in [BCG95], which is interesting, as the measurements were taken some four years later.

## 6 Conclusions

We have presented the results of 15 weeks of voice over IP measurements consisting of over 18000 recorded VoIP sessions. We conclude that the quality of VoIP is good, and in most cases is over the requirements of many speech quality recommendations. Recall that all of the sites were at academic institutions which is an important factor when interpreting these results as most universities have well provisioned links, especially to other academic sites. Therefore this work should not be over generalised. Nevertheless, the loss, delay and jitter values are very low and from previous measurements the quality trend is improving. We can only attribute this to more capacity and better managed networks than those four years ago. However some caution should be expressed as the sample period was only 15 weeks, the bandwidth of the flows is low compared to the available capacities and only sampled once per hour. We have a large number of sample sessions

so can be confident the findings are representative of the state of the network at this time. One conclusion is that VoIP is obviously dependent on the IP network infra-structure and not only on the geographic distance. This can be clearly seen in the differences between the Argentinian and Turkish hosts. Concerning the actual measurement methodology, we have found performing measurements on this scale is not an easy task. Different access mechanisms, firewalls, NATs and not having permissions on all machines, complicates the work in obtaining (and validating later) the measurements. Since it is not possible to envisage all the possible uses for this data we have made it available for further investigation.

Future work involves further improvements in collecting and analysing the data. During these measurements we did not save (but sent) the audio samples at the receiver, however future measurements will do so in order to capture the quality degradation to lost packets. Extending the measurement infra-structure to non-academic sites is also a natural progression of this work. Performing quality measures that include the end systems should also be considered, although how to include the heterogeneity of the end systems still remains unresolved.

## 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.
- [CR02] R.G Cole and J.H Rosenbluth. Voice over IP Performance Monitoring. *ACM Computer Communication Review*, 2002.
- [HHM99] Olof Hagsand, Kjell Hansson, and Ian Marsh. Measuring Internet Telephone Quality: Where are we today ? In *Proceedings of the IEEE Conference on Global Communications (GLOBECOM)*, Rio, Brazil, November 1999.
- [HMH03] Olof Hagsand, Ian Marsh, and Kjell Hanson. Sicsophone: A Low-delay Internet Telephony Tool. In *29th Euromicro Conference*, Belek, Turkey, September 2003.
- [KKI91] Nobuhiko Kitawaki, Takaaki Kurita, and Kenzo Itoh. Effects of Delay on Speech Quality. *NTT Review*, 3(5):88–94, September 1991.
- [LE01] B.M.Lines L.F.Sun, G.Wade and E.C.Ifeachor. Impact of Packet Loss Location on Perceived Speech Quality. In *Proceedings of 2nd IP-Telephony Workshop (IPTEL '01)*, pages 114–122, Columbia University, New York, April 2001.
- [Li02] Fengyi Li. Measurements of Voice over IP Quality. Master’s thesis, KTH, Royal Institute of Technology, Sweden, 2002.
- [Lin99] Dong Lin. Real-time voice transmissions over the Internet. Master’s thesis, Univ. of Illinois at Urbana-Champaign, 1999.
- [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.
- [RG98] ITU-T Recommendation G.114. General Characteristics of International Telephone Connections and International Telephone Circuits: One-Way Transmission Time, Feb. 1998.

- [SCFJ96] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A Transport Protocol for Real-Time Applications. RFC 1889, Internet Engineering Task Force, January 1996.  
<http://www.rfc-editor.org/rfc/rfc1889.txt>.

sender	receiver	Massachusetts	Michigan	California	Belgium	Finland	Sweden	Germany	Turkey	Argentina	Mean
Massachusetts	*	D:38.0 (17.1) J:2.4 (1.7) L:0.1 (0.6) H:14 (+1) T:0	D:54.2 (15.8) J:2.4 (1.5) L:0.1 (0.9) H:19 T:3	D:67.1 (15.5) J:3.6 (1.5) L:0.1 (0.8) H:11 T:6	D:97.1 (2.6) J:2.5 (1.5) L:0.1 (0.8) H:15 T:7	D:99.5 (8.5) J:3.2 (1.7) L:0.04 (0.2) H:21 T:6	D:58.4 (5.0) J:4.5 (1.4) L:0.0 (0.0) H:17 (+3) T:6	D:388.2 (43.2) J:10.4 (4.9) L:8.9 (7.2) H:25 T:1	D:112.8 J:6.1 L:1.2 H:17		
Michigan	*	D:36.4 (15.4) J:4.7 (0.8) L:0.0(0.2) H:14 (+1) T:0	D:40.4 (4.5) J:4.4 (0.8) L:0.2 (1.1) H:20 (+1) T:3	D:63.5 (4.2) J:4.3 (0.7) L:0.0 (0.1) H:11 T:6	D:88.2 (8.0) J:4.1 (0.7) L:0.1 (1.1) H:17 T:7	D:86.7 (4.7) J:5.2 (0.6) L:0.1 (2.2) H:23 T:6	D:63.6 (8.2) J:7.3 (1.9) L:0.2 (0.9) H:16 (+1) T:8	D:112.1 (10.6) J:18.7 (7.9) L:6.5 (7.0) H:25 T:4	D:106.2 J:6.8 L:1.3 H:18		
California	*	D:54.5 (16.7) J:2.0 (1.0) L:0.1 (0.36) H:18 (+1) T:3	D:40.6 (5.1) J:1.2 (0.6) L:0.1 (1.9) H:21 T:3	*	D:106.0 (3.0) J:1.4 (0.8) L:0.6 (1.4) H:25 (+1) T:10	D:108.0 (2.4) J:2.1 (0.9) L:0.2 (0.3) H:30 (+2) T:9	D:81.5 (1.8) J:4.9 (1.5) L:2.8 (3.0) H:23 T:9	D:386.9 (60.5) J:5.3 (1.7) L:4.4 (2.4) H:23 T:10	D:122.2 J:4.6 L:2.2 H:23		
Belgium	*	D:65.2 (10.1) J:1.6 (0.6) L:0.0 (0.0) H:16 T:6	D:63.4 (3.3) J:0.6 (0.1) L:0.0 (0.0) H:17 T:6	D:84.0 (1.3) J:0.9 (0.8) L:1.2 (1.0) H:23 T:9	*	D:33.4 (0.2) J:1.6 (0.9) L:0.0(0.0) H:22 T:1	D:15.8 (10.4) J:3.4 (1.5) L:0.21 (0.7) H:13 T:0	D:341.1 (24.7) J:6.9 (2.0) L:3.8 (2.7) H:16 (+2) T:1	D:96.4 J:2.0 L:0.6 H:17		
Finland	*	D:97.8 (4.2) J:1.7 (0.8) L:0.0 (0.1) H:15 (+1) T:7	D:86.8 (1.9) J:1.1 (0.6) L:0.0 (0.3) H:17 (+1) T:7	D:109.9 (4.7) J:1.4 (0.8) L:0.7 (1.4) H:24 (+2) T:10	D:30.7 (0.3) J:1.4 (0.6) L:0.1 (0.3) H:16 T:1	D:13.6 (1.0) J:1.9 (0.9) L:0.0 (0.0) H:20 T:1	D:26.3 (7.3) J:3.9 (1.1) L:0.0(0.0) H:20 (+1) T:1	D:321.2 (39.3) J:3.4 (1.7) L:3.2 (1.7) H:17 (+2) T:0	D:106.3 J:4.1 L:1.4 H:18		
Sweden	*	D:99.3 (8.8) J:3.0 (1.9) L:0.0 (0.0) H:22 (+1) T:6	D:84.9 (1.9) J:2.5 (2.0) L:0.03 (0.4) H:25 T:6	D:105.6 (2.1) J:3.2 (1.96) L:0.1 (0.1) H:30 T:9	D:33.3 (0.4) J:2.8 (1.6) L:0.1 (0.3) H:24 T:0	D:13.5 (0.5) J:2.4 (1.8) L:0.0 (0.01) H:21 T:1	D:29.8 (12.8) J:4.8 (2.5) L:0.0 (0.0) H:25 T:0	D:322.2 (30.3) J:3.2 (1.49) L:2.9 (1.0) H:41 T:5	D:107.8 J:2.8 L:0.4 H:26		
Germany	*	D:63.5 (9.6) J:1.72 (0.7) L:0.0(0.0) H:15 T:6	D:60.4 (0.5) J:0.7 (0.3) L:0.0 (0.0) H:16 T:6	D:84.4 (1.0) J:1.8 (0.7) L:2.5 (1.9) H:22 T:9	D:11.1 (0.2) J:0.8 (0.3) L:0.0 (0.0) H:12 T:0	D:27.8 (7.3) J:1.0 (0.5) L:0.0 (0.0) H:17 T:1	D:29.2 (7.6) J:1.5 (0.6) L:0.0 (0.0) H:22 T:0	D:300.7 (39.7) J:4.8 (2.1) L:3.7 (2.5) H:16 T:1	D:90.9 J:1.6 L:0.8 H:17		
Turkey	*	D:379.1 (47.1) J:8.6 (0.7) L:8.1 (2.8) H:16 (+1) T:7	D:387.9 (35.5) J:8.9 (1.2) L:8.0 (2.9) H:20 T:7	D:410.9 (43.9) J:8.8 (2.5) L:7.6 (6.8) H:19 T:10	D:330.2 (28.6) J:9.2 (2.0) L:7.10 (4.0) H:17 T:1	D:311.1 (6.3) J:9.1 (0.7) L:8.4 (3.1) H:25 T:1	D:378.2 (49.3) J:10.7 (1.2) L:8.0 (3.1) H:16 T:6	D:490.8 (26.0) J:NA L:NA H:18 T:5	D:375.9 J:8.0 L:6.9 H:19		
Argentina	*	D:117.0 (30.8) J:4.2 (2.0) L:0.5 (1.4) H:NA T:1	D:146.7 (44.2) J:4.3 (2.3) L:0.5 (1.5) H:NA T:1	D:152.0 (47.8) J:3.1 (2.4) L:0.6 (1.8) H:NA T:4	D:NA J:4.2 (2.0) L:0.5 (1.4) H:NA T:5	D:160.9 (47.7) J:2.9 (0.8) L:0.0 (0.1) H:NA T:5	D:180.5 (50.5) J:4.7 (1.5) L:0.1 (0.1) H:NA T:5	D:NA J:6.0(1.2) L:5.8 (3.0) H:NA T:6	D:115.2 J:4.2 L:1.1 H:NA		
Mean		D:114.1 J:3.4 L:1.1 H:14	D:113.6 J:3.6 L:1.1 H:16	D:115.7 J:3.5 L:1.6 H:19	D:77.1 J:3.5 L:1.0 H:13	D:105.8 J:3.1 L:1.1 H:16	D:105.2 J:3.4 L:1.1 H:20	D:104.4 J:5.5 L:1.4 H:16	D:346.6 J:5.7 L:4.0 H:17	D:180.0 J:9.3 L:4.00 H:23	D:136.2 J:4.1 L:1.8 H:18

**Fig. 2.** A summary of 18000 VoIP sessions. The delay (D), jitter (J) and loss (L) for the nine sites. The delay and jitter are in milliseconds, the losses are in percentages. The number of hops (H) and time zones (T) in hours are also given. The means for each and all sites are given plus the standard deviations (in parentheses). An NA signifies 'Not Available'.