

# Sicsophone: A Low-Delay Internet Telephony Tool

Olof Hagsand, Ian Marsh and Kejll Hanson

Swedish Institute of Computer Science  
Box 1263, SE-164 29 Kista, Sweden

16th December, 2002

## **Abstract**

The end to end delay is a critical factor in the perceived quality of service for Voice over IP applications. The described solution is a complete system-level platform and complements QoS work in the network and application areas. We describe a VoIP system that couples the low level features of audio hardware with a jitter buffer playout algorithm. Using the sound card directly eliminates intermediate buffering as well as providing fine control over timers needed by a soft real-time application such as VoIP. A statistical based approach for inserting packets into audio buffers is used in conjunction with a scheme for inhibiting unnecessary fluctuations in the system. We give comparisons for the performance of the playout algorithm against idealised playout conditions. We also present mouth to ear delay measurements for selected VoIP applications and show that several hundreds of milliseconds can be saved by using the techniques described in this paper. A prototype for both UNIX and Windows platforms (NT and 9X) has been implemented, demonstrating that our system adapts to network conditions whilst maintaining low delays.

# 1 Introduction

Users of interactive VoIP applications demand low latency conversations. Replaying packetised audio requires that sufficient packets are available to the application in order to avoid gaps or glitches. The digital to analog conversion of sampled voice requires strict, synchronous timing despite the fact that the network and operating system may disrupt the process. The most common method to solve this problem is to introduce a small intermediary buffer between the decoded audio stream and the audio hardware which allows packets to be “available” for playout. Of course withholding packets instead of immediately playing them increases the total delay of a VoIP application. However, the longer packets can be delayed, the more resilient the receiver is to adverse network conditions. The contribution of this work is a description of the implementation of a buffer playout algorithm within a VoIP system. The playout algorithm implemented is almost identical to NeVoT [Sch92]. Better schemes in terms of playout delay have been documented [RQS00] but the simplicity needed for our scheme moved us to make this choice. We should point out that Sicsophone is a working implementation and this the algorithmic complexity is an important factor. We motivate this approach with real delay experiments and results. Hence the goal is not to compare the merits of various playout algorithms, this has been covered by many researchers, rather to give some insight what issues are important when realising these schemes and their effects.

In this paper we refer to the mouth to ear delay as the total one way delay experienced by two speakers including the analog-digital-analog conversion. By jitter we mean the variability in the packet delay. This variability is the reason we need to buffer packets, thus our work focuses on how to detect and compensate for packet jitter in an efficient manner. Our solution is to insert packets into the memory of sound cards relieving the need for data copying or context switching. This approach saves precious time, avoids scheduling problems but requires careful buffer management.

Figure 1 illustrates the complete path of audio samples from a microphone at a sender to the loudspeaker at a receiver. Traditionally, a sender writes voice samples to the operating system which are subsequently sent across the network to a receiving host. At the receiver, data is read from the operating system interface where it is the responsibility of the application to adjust the buffer size as required, this is shown by the solid lines in the illustration.

In our approach we use the buffering scheme in the operating system and copy the packets directly into the memory of the sound card. Therefore, we save copying the data to and from the application, plus not performing the de-jittering in the application. We describe our approach in the context of

DirectSound [BD98] on the Windows platform. It is important to point out that our approach is not confined to this architecture, a ring buffer with pointer support is sufficient to realise the ideas presented in this paper (alternatives for UNIX include [Riz97] or [Ree98]). However we describe the system using DirectSound as it is known to many developers and was used in our experimental evaluation.

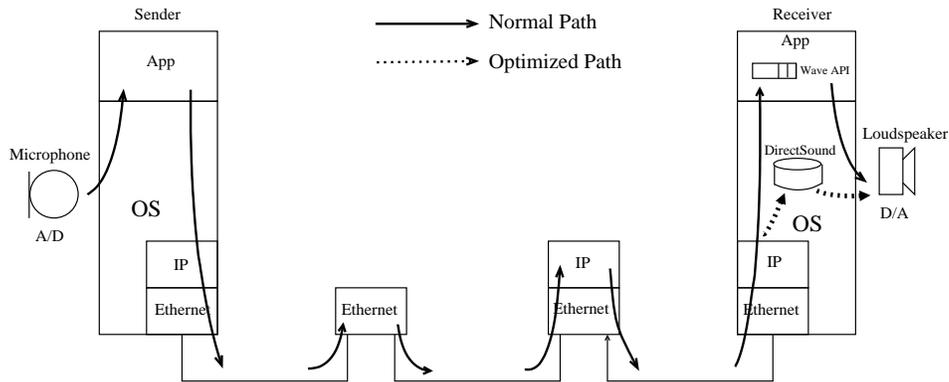


Figure 1: Audio Delivery Path within Sicsophone

If we now look at the steps a receiver must take to reply packetised audio from the network in more detail, Table 1 shows four such typical steps. Firstly, de-packetisation, removes the IP and UDP headers and passes the datagram together with a Real Time Protocol (RTP) payload to a VoIP application. This step takes a few milliseconds on most systems. Step two is to decode the sound samples, this is dependent on the compression scheme as well as the packet size used. Typically this takes from a few milliseconds to tens of milliseconds. Steps three and four are usually performed as distinct steps, absorption of network delays through buffering and delivery to the sound application. Our goal was to consolidate these steps into a single step, saving the time of inter-mediatory buffering and context switching. We refer to this approach, solely for definition by its software name, Sicsophone.

Table 1: Typical Receiver Incurred Delays

Step	Process	Overhead (ms)	Depends On
1	De-packetisation	10 - 50	Packet Size
2	Decoding	10 - 50	Coding Used
3	Buffer Delay	5 - 200	Network
4	Delivery	5 - 120	End System

The remainder of this paper is organised in the following fashion; Section 2 forms the main body of this work, low level adaption of playout buffers using ring buffers. Section 3 presents results of Sicsophone's performance of mouth to ear results for different VoIP tools. We also give comparisons of the playout delay with Sicsophone against the idealised case. We also give a brief description of measurements that can be made with tools such as Sicsophone. Section 4 is a description of related efforts with which this paper has commonalities, we round off the paper with some conclusions in Section 5.

## 2 End-system Adaption to Jitter

### 2.1 Buffering Issues

In this section we outline some issues associated with the data buffering scheme we have chosen. Our goal is to save time by avoiding data copying, setting up direct memory access (DMA) transfers ahead of time, using simple data structures and inserting de-jittered audio packets directly into the memory of the sound card. Using the sound card as a buffer has the advantage of not adding any extra buffering to the audio sample path. It also avoids copying data from a kernel to an application and back again. Direct memory access is used to move data from memory to the sound card and vice versa without intervention of the CPU. Using DMA efficiently is not trivial, as it can take some time to set up the transfer. However once it is done, the transfer can be done much quicker and more efficiently. This offers significant time savings over posting an interrupt for every packet, particularly in the older (and non-DirectX) versions of the Windows operating systems.

One potential problem of using the sound card memory as a buffer is it could be overrun by packets arriving too quickly, for example on a fast connection. Modern sound cards however are equipped with megabytes of RAM to store down-loadable sound samples, DirectSound can allocate buffers up to this physical size when a hardware buffer is initialised. Another potential cause of overrun or underrun is misaligned or drifting clocks, there is no mechanism in Sicsophone to detect this. A solution to detect this problem is presented in [MST99].

Another important but often overlooked issue is mixing. For an application like VoIP where the voice channel is stopped and started continuously, we would like to minimise this setup time. Valuable time can be lost by setting up mixers for software and hardware buffers where we normally do not want to mix audio from different sources, e.g. the VoIP and a MP3

song. Therefore we allocate a DirectSound primary buffer to give better delay characteristics, as it does not need to be mixed before outputting to the D/A conversion.

The coding scheme used is another issue. Packets have to be decoded before insertion into the buffer if they are not in PCM format. Using PCM allows us to DMA the payload into the sound card memory without any audio format conversion. The general design tradeoffs now become clearer, the more generality and flexibility we demand, the greater the price in processing delay.

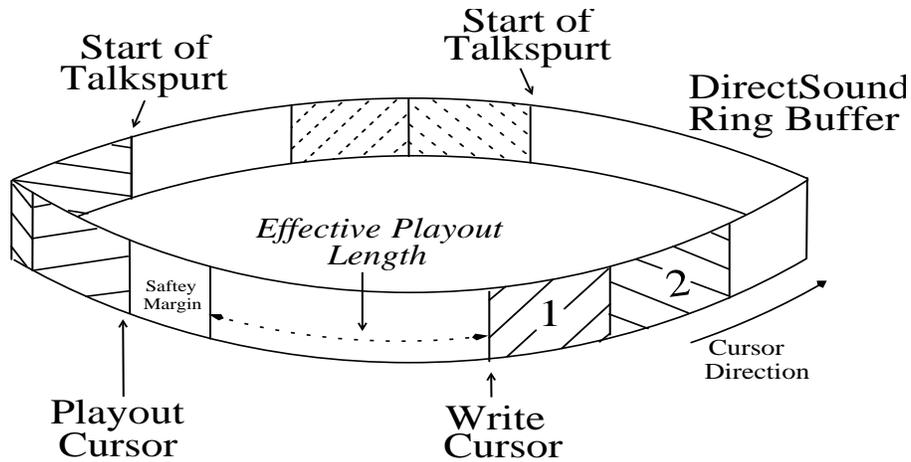


Figure 2: DirectSound Buffer Structure

Figure 2 shows the interface offered by DirectSound. Data is written at the write cursor and replayed by the trailing read pointer. The read and write pointers are updated by the system, and continuously encircle the buffer. In Sicsophone they are used as **both** timers and pointers. They function as a timer by indicating if a packet is too late. If the read pointer has already passed the point where a packet should be, and it has not been written, then we know that this packet is late. Insertion is simply a modulo operation and a pointer copy. Using a ring buffer in this manner is identical to that described in [Sch92] where the authors motivate their choice of using a circular buffer for performance reasons. In order not to replay old data from the buffer when no packets are being sent we write “empty” samples that sound like audible background noise into the buffer so that the listener is aware the connection is still open.

Used as pointers, the read and write cursors give memory locations where data is read or written to depending on the operation to be performed. Given these pointers it is easy to adjust the buffer length, it is simply where packets

are chosen to be read from. The closer the read pointer is to the write pointer the smaller the effective buffer length will be. Note there is a small margin of 15ms in front of the read pointer to allow data that has been written to be “ready” for playback. Use of this safety margin is recommended by the developers of DirectSound.

To give a concrete example, using an estimate of the network delay and its variation, described previously, we insert packets at a specified “distance” ahead of the read pointer. Therefore a translation from milliseconds to bytes is needed;  $bytes = (samples\_sec \cdot bits\_sample \cdot P_i) / 8000$ . For example, if one substitutes 8000 for *samples\_sec*, 8 bits per sample and 200ms for the playout point this equals 1600 bytes. This means that the write cursor can simply be set 1600 bytes in front of the read cursor. The safety margin is also included in the length of playout buffer but it is possible to simply subtract the value (120 bytes in this case) from the calculation. To re-iterate once the playout point has been calculated it is trivial to insert packets into the buffer, no complex data operations are needed.

## 2.2 Fast Startup Adaption

In an adaptive VoIP application we normally consider changing the buffer size during a silence period so as not to introduce audible glitches in the analog audio stream. Since the goal of this work is to produce a low-level VoIP tool we would like to keep the buffer length close to optimal. However in the startup phase we have no idea of the network condition and therefore have to use default values for the network delay (Sicsophone uses  $min = 20ms$ ,  $max = 60ms$ ). We therefore adjust the buffer length after monitoring only a few packets to settle to an estimate quickly.

Figures 3 a) and b) show packet delay in the jitter buffer during the start up phase of Sicsophone as an example. The y-axis shows the waiting time in the buffer (in ms) and the x-axis shows the number of packets received, sorted by the time spent in the buffer, note this is **not** the sequence number. It shows the number of packets and their respective waiting times. Figure 3a) shows the buffer state after ten packets have been received and 3b) after an additional 40 packets have arrived, the original ten are shown with bolder lines. After ten packets were stored, the time spent in the buffer varied between 14 and 30 ms whereas after 50 packets the median delay incurred is around 20 ms. This is not surprising as packets are sent with a 20 millisecond separation. Although this trace was taken on a local area network we have seen that the majority of packets arrive with the same inter-packet spacing even on wide area networks [HHM99].

Fast adaption is worthwhile during the start up phase of a VoIP session.

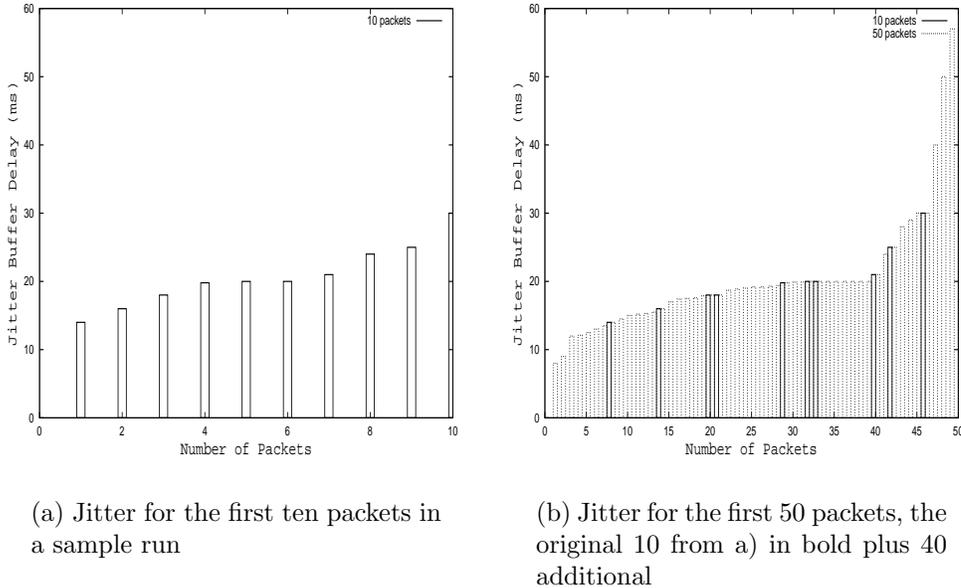


Figure 3: Arrival of 10 and 50 Packets

The alternative approach is to be conservative in the start up phase and have long playout buffers until a value for the playout point can be calculated. In the presence of spikes [RKTS94] we can re-estimate the jitter value quickly. Since the goal of this paper is a low delay VoIP tool we chose to adapt quickly. Furthermore, usually there are sufficient silence periods during the startup phase of a conversation to perform fast adaption. In the case where there are none (such as call waiting, i.e. music playing) we adjust the buffer when a packet is excessively delayed or if there is a loss, failing these possibilities we adjust the buffer length and tolerate an audio glitch.

### 2.3 Bounding the Estimated Network Delay

Sudden increases in the network delay<sup>1</sup> can cause VoIP applications problems. A spike is referred to a sudden and rapid increase in the network delay which is typically short lived, often less than one round trip time. One solution is to follow the increase in the delay and adjust the buffer length accordingly. The alternative is to include the values of the jitter estimate but not adapt the buffer size. It is because of this temporal property plus we are dealing with real hardware, has led us to be more conservative to network conditions and

<sup>1</sup>referred to as spikes in [RKTS94]

not to adapt the DirectSound buffer length to sudden increases in network delay.

We have implemented a system which bounds the delay jitter estimate. As stated, the spikes are not completely ignored but we do not react immediately to their presence. The estimated jitter value should vary between an upper  $Q_{max_i}$  and a lower  $Q_{min_i}$  (see the key in Figure 4) bound in a “corridor”, where  $Q_{min_i} < d_i < Q_{max_i}$ . If the running estimate breaks either of the boundaries we re-calculate the new buffer length, taking into account the value of the spike, but reset the mean estimate to the middle value of this new  $(Q_{min_i}, Q_{max_i})$ . Figure 4 shows an example of a receiver jitter buffer during a conversation between two machines on a local network. The y-

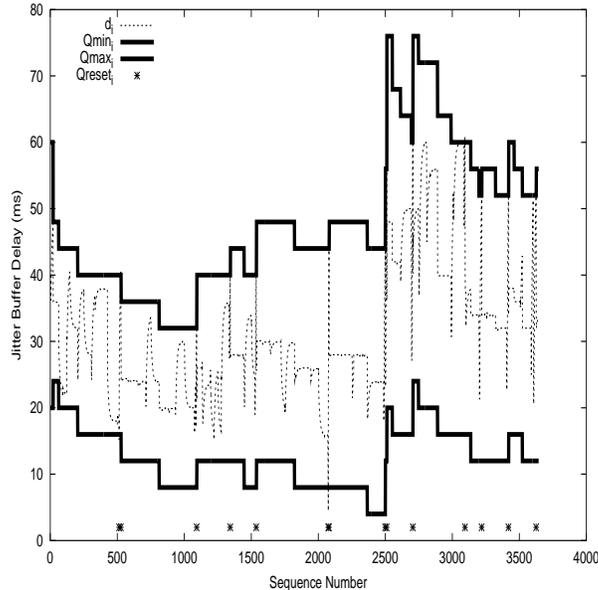


Figure 4: Jitter Buffer Delays with Bounded Mean

axis shows the jitter buffer length and the x-axis the sequence number. The system starts with  $Q_{min}$  and  $Q_{max}$  set to the default values 20ms and 60ms for the minimum and maximum bounds respectively. At the bottom of the figure we show the corridor breaks as stars to highlight the breaks. We have found this scheme to work well, in the given trace there were only 14 breaches of the corridor from over 3600 packets (less than 0.5%). More importantly we did not make costly, unnecessary changes to the DirectSound primary buffer.

### 3 Results

We divide this section into two sections, the first gives the total delay of popular VoIP tools compared to Sicsophone in a laboratory environment, with basically no, or little, network delay. Secondly, we show the performance of the playout algorithm using trace files taken over the Internet (two of the ones used in [RKTS94]). This shows a typical WAN component due to the jitter on the Internet plus the best possible playout that could have been achieved by post-processing trace files. We chose to give the results in this manner to estimate the real mouth to ear delay by including both the WAN and LAN components. It can be seen as a sum of these two quantities. Essentially the first set includes the delay due to coping with the operating system and the second with the network conditions. Using these trace files we can compare our implementation with those published. We must point out these trace files are old so we have taken much more recent ones and describe this a little in the third subsection of the results.

#### 3.1 Mouth to Ear Measurements

The delay contributed by the end systems is the main result of this paper. We performed one-way mouth to ear measurements with a range of VoIP tools and the results are summarised in Table 2. It's important to state that no parameter tweaking of these tools was done, we used their default installation values. The experimental setup used was as shown in Figure 1. The measurements were done using a signal generator feeding a sender and an

Table 2: Mouth to Ear Latency Measurements

Audio Tool	Latency (ms)
Sicsophone prototype	25-100
Ericsson Lanphone	300
VocalTech Internet Phone 4.5 (S. Blaster)	450-550
VocalTech Internet Phone 4.5 (PhoneJack)	580-620
Microsoft NetMeeting 2.1 (S. Blaster)	620
Microsoft NetMeeting 2.1 (PhoneJack)	750
VAT 3.4 (Solaris)	1200
RAT 3 (Solaris)	1500

oscilloscope to measure the time difference between the sender and receiver. We can see that there are large variations between the various applications. One important result of this paper is to highlight the design of end systems

for VoIP applications. Our goal is not simply to state Sicsophone is superior to other tools, rather to show the considerable time savings, 10's to 100's of milliseconds, can be saved by using the approach described or a similar one.

### 3.2 Comparison with Ideal Playout Conditions

In the introduction we mentioned that a jitter buffer playout algorithm essentially has to tradeoff low delay or packet loss due to their arrivals being too late. Low delay implies a short playout buffer, incurring higher packet loss due to late arrivals. When comparing the performance of algorithms it makes sense, therefore, to consider loss and delay.

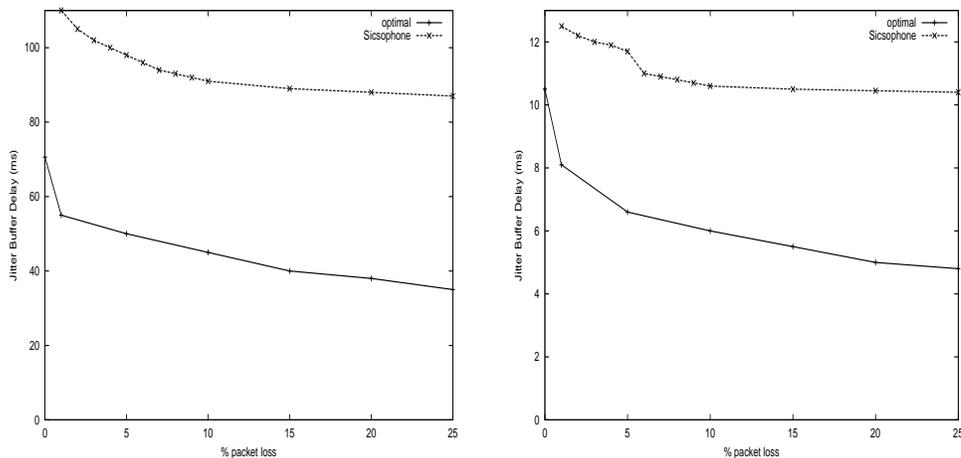
As stated in the introduction to this section, we include this figures to give an estimation of the delay incurred by the buffering we employ. They should be taken as the delays incurred for a given trace at a particular time and are used to show the *network* delay which would be added to the *system* delay, i.e. those in in Table 2.

Figure 5 shows the results for two Internet trace files<sup>2</sup>. To calculate the optimal playout point we order all the packets and remove the 1% with the highest delay. We then calculate the delay needed to play the remaining 99% of the packets resulting in the delay for 1% packet loss, this process is repeated up to 25% packet loss (although in practice more than 10% would be deemed unacceptable). Figures 5a) and 5b) show Sicsophone's playout delay performance in comparison.

In Figure 5a) Sicsophone is about 50ms from the ideal playout point and remains more or less constant as the packet loss increases. For a given loss rate, e.g. 5% Pinto and Christensen ([PC99]) quote a slightly lower delay than we do, 72ms compared to our 98ms and similarly so for other loss rates. We should re-iterate the focus of this paper is the implementation and the absolute/measurable delays rather than the playout algorithm itself. Nevertheless, the result in this case is due to the large variation of jitter ( $\pm 20ms$ ), which makes it hard to settle to a constant value for the minimum buffer length, this can be verified by looking at the absolute jitter. The ITU G.114 document recommends that the one-way delay should not exceed 150ms, so this is about one third of the recommended delay used in the buffer playout algorithm. A second test is shown in Figure 5b) which shows a trace from the University of Massachusetts in Amherst to GMD in Berlin. In this case the jitter is much better and the difference between the optimal and Sicsophone is only 5ms. We have included both the worst and best test cases of the five traces available. The traces are seven years old so we took some

---

<sup>2</sup><http://gaia.cs.umass.edu/~sbmoon/traces.html>



(a) Playout Delays for a Trace from UCI, California to INRIA in France

(b) Playout Delays for a Trace from Amherst, Mass. to GMD, Berlin

Figure 5: Playout Delays for Two WAN Traces

new ones using Sicsophone invoked automatically by scripts to obtain more and newer data for test traces, this is described next.

### 3.3 Accurate VoIP Measurements

Using Sicsophone we have been able to gather delay, jitter and loss statistics for a number of Internet paths. Since we intercept the packets at a low level we are able to record the arrival of packets with no or little application delay. This enables us to obtain accurate measurements, particularly the jitter. Loss is relatively easy without low level access to the data but delay is harder to calculate accurately without tracing both the sender and the receiver and allowing for unsynchronised clocks that do not drift.

In 1999 we performed WAN measurements for VoIP sessions which is documented in [HHM99]. We performed the same measurements in 2002 and compared the results, these are shown in Table 3. Sicsophone is invoked automatically, both the sender and the receiver and a sample conversation is sent between two sites. Connecting the sites available to us in a full-mesh topology enabled us to gather more than 60,000<sup>3</sup> VoIP trace files. We used more than ten test sites around the globe (Figure 3.3) to gather data about the delay, loss and jitter. Of the sites which remained the same see Figure

<sup>3</sup>Available from <http://www.sics.se/~ianm/COST263/cost263.html>.

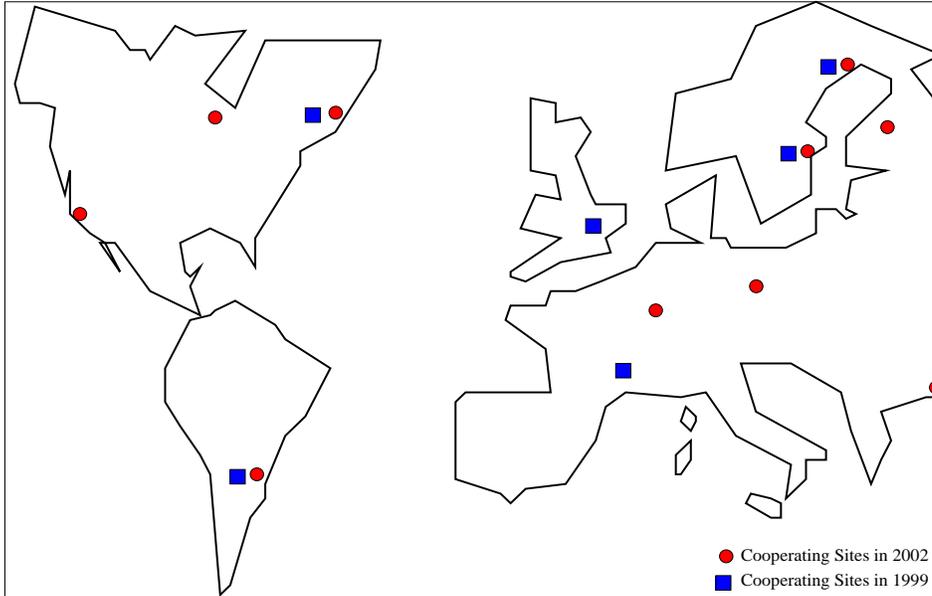


Figure 6: Locations of Test Sites

3.3 (a square and a circle together) we were able to make an approximate comparison of VoIP quality given the three year separation. Using a tool like Sicsophone makes VoIP measurements very straightforward and gives us a very accurate timestamp of the arrival of the packets. More results from these measurements can be found in [Fen02].

Quality	1999	2002	% Diff. (+/-)
Jitter	45.1ms	22.6ms ( $\pm 13.7$ )	-50.0%
Loss	1.2%	0.5% ( $\pm 2.6$ )	-58.3%
Delay	115 ms	84.95ms ( $\pm 44.85$ )	-26.1%

Table 3: 1999 and 2002 VoIP Quality Differences

## 4 Related Work

The early 90's produced a surge in packet audio playout research. One of the first efforts to implement a voice application on an IP network with an adaptive buffer playout strategy was NeVoT [Sch92]. They use a variation estimate similar to the one given earlier, however they make a slight distinction for the first packet in a talkspurt and subsequent ones. The playout for the first packet is delayed longer due to lack of information on the network

state after the silence period. Our work shares theirs in the choice of a ring buffer for buffering packets, only we perform this within the operating system and not as part of the application. VAT (Visual Audio Tool) [JM92] is a well known VoIP tool that implements a playout buffer similar to the one described, including a circular buffer to hold the packets before playout. We use an additional scheme to prevent the jitter estimates from varying too rapidly plus focus on the efficient insertion of packets into the playout buffer. Moon *et al.* [RKTS94] present four different playout algorithms for packet audio. All calculate an estimate of the network delay and jitter as an average from all the packets measured. The authors study jitter spikes in traces and also do not adapt the buffer size to these spikes. Pinto and Christensen [PC99] describe an algorithm for jitter compensation based on the target packet loss rate. Their “gap based” approach compares the current playout time with the arrival time and calculate a gap for both early and late packets. They compare the current playout delay, for any particular talkspurt in progress, with an optimal playout delay. This optimal theoretical delay is defined as minimum amount of delay to be added to the creation time of each packet which would result in a playout of a talkspurt at the given loss rate. Our calculation of the optimal playout is similar to the one described in this paper. Luigi Rizzo describes a generic sound card driver for FreeBSD [Riz97]. Aspects of it resemble our work, in particular, handling of timers, DMA transfer and buffer size allocation. They include hooks to use the driver for VoIP applications, one such example is a `select()` call which can be scheduled to return only when a certain amount of data is ready for consumption. As stated earlier, it is possible to implement Sicsophone on this device driver. [RQS00] looked at combining target-based playout algorithms in conjunction with FEC schemes, and propose a number of new playout algorithms based on this coupling. [KH97] keep the flow of audio constant during operating system load by using buffering in the audio hardware. They also look at reducing the amount of buffering in the application by keeping the buffers in the application as small as possible. In our case we try and totally eliminate it by only using the hardware buffers.

## 5 Conclusions

In this paper we have shown how careful buffer management combined with a simple statistical playout scheme can reduce mouth to ear delay for VoIP applications. As stated at the start of this paper, delay is one of the most important factors in the perceived QoS and this has been the focus of this work. The results are encouraging as the mouth to ear delay of Sicsophone

on a LAN is around 50ms on a Windows NT system with DirectX 8.0.

We also include an estimate of the delay induced by network conditions using the playout algorithm, but include it to help estimate the network delay of VoIP tool such as Sicsophone. The total delay can be seen as an approximate summation on these two delays, since a real end to end measurement over a WAN is very difficult.

Future work in this direction is to measure the delay with operating systems such as Windows XP or UNIX platforms. The tool will be released publicly after these updates have been made.

## References

- [BD98] Bradley Bargaen and Peter Donnelly. *Inside DirectX*. Microsoft Press, 1998.
- [Fen02] Li Fengyi. Wide area measurements of voip quality. Master's thesis, Royal Institute of Technology (KTH), 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. IEEE.
- [JM92] Van Jacobson and Steve McCanne. vat - LBNL audio conferencing tool, July 1992. Available at <http://www-nrg.ee.lbl.gov/vat/>.
- [KH97] Isidor Kouvelas and Vicky Hardman. Overcoming workstation scheduling problems in a real-time audio tool. In *Proc. of Usenix Winter Conference*, Anaheim, California, January 1997.
- [MST99] Sue Moon, Paul Skelly, and Don Towsley. Estimation and removal of clock skew from network delay measurements. In *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, New York, March 1999.
- [PC99] J Pinto and K Christensen. An algorithm for playout of packet voice based on adaptive adjustment of talkspurt silence periods. In *Proceedings of the IEEE 24th Conference on Local Computer Networks*, pages 224–231. ACM, October 1999.
- [Ree98] Dicken Reed. A new audio device driver abstraction. In *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, 1998.

- [Riz97] L. Rizzo. The FreeBSD audio driver. *Lecture Notes in Computer Science*, 1356, 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.
- [RQS00] Jonathan Rosenberg, Lili Qiu, and Henning Schulzrinne. Integrating packet FEC into adaptive voice playout buffer algorithms on the internet. In *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, Tel Aviv, Israel, March 2000.
- [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.